1 Stochastic Methods for Modeling Decision-making

Adele Diederich and Keivan Mallahi-Karai

1.1 Introduction

Decision-making situations, small or large, are everywhere: from simple and fast perceptual decisions to those with long-term effects, such as accepting or declining a job offer – we have to make decisions on a daily basis. The choice alternatives may be few or many; the outcomes may be certain or probable; the decision-maker may or may not be under time constraints. They all call for an attempt to find ways to model the underlying choice processes that can account for the data that can be collected during these processes. The goal is to come up with models that are, on the one hand, accurate and with explanatory power, and on the other hand, enjoy simplicity and robustness.

Mathematics has a role to play here. During the last 60 years mathematical models of decision theory have been developed to study these processes using tools from the theory of stochastic processes. These models have become the dominant approach to modeling decision processes in psychology and cognitive science. The probabilistic framework is theoretically consistent with the idea that the outcome of the choice and the time it takes to make a decision, i.e., the response time or reaction time, are not set or determined at the outset, but are something that emerges out of the deliberation process.

Intuitively a stochastic process is an entity that evolves randomly in time or space. In physics or biology and neuroscience this entity may be particles or neural activation. In psychology and cognitive science, depending on content and context, it is referred to as amount of information, activation, pieces of evidence, preference, and the like.

In psychology, two major classes of stochastic processes have mainly been applied to account for choice frequencies and choice response times. One class of models assumes that evidence for one option is at the same time evidence against the alternative option. They are mostly applied to binary choice situations. Within this class, random walk models accumulate discrete evidence in discrete time, whereas diffusion models accumulate continuous evidence in continuous time. The most commonly used version of the diffusion model is the Wiener diffusion model that linearly accumulates evidence without any decay (Ratcliff, 1978), but other models include the Ornstein–Uhlenbeck model that linearly accumulates evidence with decay (Busemeyer & Townsend, 1993; Diederich, 1995), and the leaky competing accumulator (LCA) model (Usher & McClelland, 2001) that

nonlinearly accumulates evidence with decay. The other class of models consists of accumulator and counter models. They can easily be extended to more than two choice alternatives, because an accumulator/counter is established for each choice alternative separately, and evidence is accumulated in parallel. A decision is made as soon as one counter wins the race to reach one preset criterion. The accumulators/counters may or may not be independent. Poisson-counter models are prominent examples, but random walk and diffusion models, one process for each alternative with a single criterion (absorbing boundary) for each process, can also be employed. Other accumulator models such as LATER (linear approach to threshold with ergodic rate) (Carpenter & Williams, 1995) and LBA (linear ballistic accumulator) (Brown & Heathcote, 2005) assume a deterministic linear increase in evidence for one trial. Randomness in responses occurs by assuming a normal distribution across the linear accumulation rate. These models are not considered further here.

In the following we focus on random walk/diffusion models. Our focus will not be on presenting the state-of-the-art research from the mathematical perspective or the latest debates in psychology and neuroscience. Rather, we we will try to provide the underlying mathematical ideas and tools that have proven to be very successful in the last 60 years in psychological research. Along the way, we will discuss various models that have been proposed and try to lay out the underlying mathematical assumptions that have been placed. In particular, we start with a very simple model, a random walk, and develop more elaborate models from there. The focus will be on Markov chains, which has been referred to as the matrix approach to derive predictions of the models. Numerous examples provide deeper insight to the interplay between psychology and mathematics, i.e., mathematical psychology.

1.2 Probabilistic Modeling in Decision Theory

Sequential sampling models are among some of the most developed decision-making models. These models assume that characteristics of the choice options can be mapped onto a hypothetical numerical dimension representing the instantaneous level of information, activation, evidence, or preference. Further, they assume some random fluctuation of this value over time in the course of the accumulation process. Therefore, sequential sampling models can be built using stochastic processes, that is, a collection of random variables, representing the evolution of some system of random values over time,

$$\{X_\alpha\}_{\alpha \in A}.$$

Note that the *index set* or *parameter set* of the random variable X may be discrete ($\alpha = n \in \mathbb{N}$) or continuous ($\alpha = t \in \mathbb{R}^+$) and relates to the time of a realization of the random variable. We will call this set here *time set* or *time space*. The *state space* S defines the values or *states* that the random variables X_α can take on and may be also discrete ($\mathbb{Z}$) or continuous ($\mathbb{R}$). Let us make a brief remark on the connection between discrete and continuous models. Discrete-time models have

the advantage that their analysis often does not involve hard-core mathematical technology. This simplicity adds to their appeal, especially to the less mathematically sophisticated user. On the other hand, some experience with both of these set-ups is enough to show that the continuous processes are often more amenable to analysis. When at all possible, it is more helpful to derive closed formulas for continuous processes, and even in the absence of such formulas, one can often use various techniques from numerical analysis to return to the discrete set-up and use it as an approximation of the continuous process.

Note that there is often a close connection between discrete and continuous processes. The celebrated Wiener process can be viewed as the limit of suitably scaled random walks. This connection can also sometimes be used to analyze one of these models in terms of the other.

1.2.1 Information Accrual

During the last decades, various probabilistic models have been developed to explain the process of decision-making based on accumulation of evidence for the alternatives at hand. Each model represents the space of available information up to a certain point in time as a process that takes values in a subset of the Euclidean space $\mathbf{R}^d$, where the dimension d is often equal to the number of alternatives.

The process starts at a given point that represents the initial information or bias the decision-maker may have toward one of the choice alternatives. The information accrued up to time T (which may be discrete or continuous) corresponds to a point in this space. Thus, the entire deliberation process can be viewed as driven by a discrete or continuous random process. The discrete models often lead to models based on random walks, whereas continuous models are based on a number of diffusion models. The structure and parameters of this stochastic process highlight the underlying assumptions about the process one wants to investigate, for instance, sensory receptions, storing, memory retrieval, preference, categorizing, and more.

Most of the models that have been developed in these fields are based on Markov processes. The characteristic feature of a Markov process is that the dynamics of the process is determined by the current state of the process.

One of the implications of the Markov property is that the past can only influence the future through the present state of the system. Another way of stating this is that given the present state of the process, its future is independent from its past. This clearly indicates that a Markov process has a weak memory. This drawback aside, the Markov assumption allows one to bring a large body of existing mathematical theory (including tools from linear algebra and analysis) to bear on various applications. There are various generalizations of the Markov property that allow for some amount of memory. These systems that have a bounded amount of memory can also be recast as a Markov model with a different state space (see subsection 1.8). On the other hand, it must be mentioned that non-Markovian models (such as self-avoiding random walks) have proven to be much more difficult to analyze.

Another component of any decision model is the termination criterion. The stochastic process is run until a stopping criterion is satisfied. There are various ways to stop the process. The most obvious one is to have a fixed period for deliberation, after which the process is stopped. This may be related to an experimental set-up, in which the decision-maker is asked to make a response at a predetermined point in time. In other words, it is a time-based criterion. In many models, however, the criterion is satisfied when the process crosses a certain threshold, representing a sufficient amount of evidence in favor of one of the alternatives. Here the decision-maker sets an evidence-based criterion. The former criterion has also been called *fixed stopping rule*, the latter *optional stopping rule* (Busemeyer & Diederich, 2002). In the presence of more than two alternatives, it seems that both the accumulation process and the stopping criterion can be rather involved. In modeling situations in which decision-making is subject to time pressure, the criterion for stopping can also change with time. These will be discussed in the section on stopping times. We will start with a brief and rather informal discussion of the random walk model.

1.2.2 Random Walk Models – An Example

In this subsection, we will give an informal definition of random walks on the line and explain how this model can be used as a prototype of a decision-making system. More details will be given in the subsequent sections. We will start with an example of a *simple random walk with two absorbing boundaries*. A random walk $\{X_n\}$ is a stochastic process with discrete time space, that is, the realization of the random variable occurs at discrete times

$$t = 0, 1, 2, \ldots$$

and discrete state space S. The random process X_n can thus take value in the set $\{0, \pm 1, \pm 2, \ldots\}$. The model as such has few applications in psychology. However, it serves as an intuitive example and introduces basic concepts used throughout the chapter.

Suppose a person has to make a decision between two choice options A and B. At any moment in time, the person may sample information ξ for choosing A or B. For the sake of concreteness, let us assume that the probability of sampling information in favor of option A equals 0.3; the probability of sampling information ξ in favor of option B is 0.7. Amounts of information are coded in units of one. We will represent the total amount of gathered information by a real number. Assume that information in favor of option A amounts to moving from point u to the point $u+1$ and information in favor of B amounts to moving from u to the point $u-1$. Suppose that ξ denotes the unit of gathered information. We will express this by writing

$$\mathbf{Pr}[\xi = +1] = 0.3, \quad \mathbf{Pr}[\xi = -1] = 0.7.$$

The coin-tossing random variables ξ_i are examples of the *signed Bernoulli distribution*. A random variable ξ is said to have a signed Bernoulli distribution if there

is $0 < p < 1$ such that ξ takes values 1 or -1 with probabilities p and $1 - p$. We express this mathematically by writing

$$\mathbf{Pr}[\xi = 1] = p, \qquad \mathbf{Pr}[\xi = -1] = 1 - p.$$

Signed Bernoulli random variables are variations of the ordinary Bernoulli random variables, where, instead of 1 and -1, the random variable takes values of 1 and 0.

Two important assumptions have been made here: identical distribution of steps, and their independence.

Assumption 1.1 The random variables ξ_i are *identically distributed*. In other words, the probability that the information accrued at time i is ± 1 is the same as the probability that the information at time j be equal to ± 1:

$$\mathbf{Pr}[\xi_i = 1] = \mathbf{Pr}\left[\xi_j = 1\right] = \frac{1}{2}, \quad \mathbf{Pr}[\xi_i = -1] = \mathbf{Pr}\left[\xi_j = -1\right] = \frac{1}{2}.$$

This assumption essentially states that the information source has reached a stationary distribution. Such assumptions are rather common in modeling problems.

Assumption 1.2 The random variables ξ_i are *independent*. In other words, the information accrued at time i will not influence any ξ_j for $j \neq i$. This can be mathematically expressed as follows. For any values of $\epsilon_1, \ldots, \epsilon_n = \pm 1$, we have

$$\mathbf{Pr}[\xi_1 = \epsilon_1, \ldots, \xi_n = \epsilon_n] = \prod_{1 \leq j \leq n} \mathbf{Pr}\left[\xi_j = \epsilon_j\right].$$

Let X_n denote the amount of information accumulated up to time unit n, $n \in \mathbb{N}$, i.e.

$$X_n = \sum_{i=1}^{n} \xi_i. \tag{1.1}$$

Continuing with our example, assume that at the beginning of the trial no amount of information has been sampled yet, $X_0 = 0$, and that the sampling process stops as soon as a critical value is reached for either initiating a response in favor of option A or in favor of option B. Let us assume that the critical value for option A is $\theta_A = 4$ and for option B is $\theta_B = -4$. The state space for this example is therefore given by

$$S = \{0, \pm 1, \pm 2, \pm 3, \pm 4\}.$$

The critical values $\{-4, 4\}$, a subset of the state space, are called *absorbing states*; the remaining intermediate states $S^* = \{-0, \pm 1, \pm 2, \pm 3\} \subset S$ are called *transient states*. Eventually the accumulated information leaves the transient states and is captured by one of the absorbing states. That is, the probability of absorption is 1.

The states and the probabilities for moving up or down, that is, the *transition probabilities* (defined properly later), can be conveniently presented in matrix form (Equation (1.2)). The rows of the matrix display all states, in our example $-4, -3, \ldots +3, +4$, and are represented by rows $1, 2, \ldots, 9$, respectively; and

similarly for the columns. This matrix (the entries within the brackets) is called the *transition probability matrix* and is denoted by **P**.

$$\mathbf{P} = \begin{array}{c|c|ccccccccc} \textit{index} & & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ & \textit{state} & -4 & -3 & -2 & -1 & 0 & +1 & +2 & +3 & +4 \\ \hline 1 & -4 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 2 & -3 & .7 & 0 & .3 & 0 & 0 & 0 & 0 & 0 & 0 \\ 3 & -2 & 0 & .7 & 0 & .3 & 0 & 0 & 0 & 0 & 0 \\ 4 & -1 & 0 & 0 & .7 & 0 & .3 & 0 & 0 & 0 & 0 \\ 5 & 0 & 0 & 0 & 0 & .7 & 0 & .3 & 0 & 0 & 0 \\ 6 & +1 & 0 & 0 & 0 & 0 & .7 & 0 & .3 & 0 & 0 \\ 7 & +2 & 0 & 0 & 0 & 0 & 0 & .7 & 0 & .3 & 0 \\ 8 & +3 & 0 & 0 & 0 & 0 & 0 & 0 & .7 & 0 & .3 \\ 9 & +4 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{array}. \tag{1.2}$$

Each cell, p_{ij}, of this matrix represents the transition probability of going from state s_i to state s_j. For example, row 5 reflects the transition probabilities from the neutral state (state 0) to either the state one step up (to state $+1$ from column 5 to column 6 with probability .3); or remaining in neutral (with probability 0); or to the state one step down (to state -1 from column 5 to column 4 with probability .7). The remaining rows are defined in a similar manner.

Three possible *realizations* of the information accumulation process X, also called *sample paths* or *trajectories*, are

trial 1 $(0, +1, 0, -1, -2, -3, -4, -3, -4)$
trial 2 $(0, +1, +2, +3, +2, +3, +4)$
trial 3 $(0, -1, 0, -1, -2, -1, -2, -3, -4, -3, -4)$

In the first trial, the first amount of information gives evidence for choosing option A and the state changes from neutral to a value that moderately favors A $(+1)$. However, the evidence is not sufficiently strong to decide. The next amount of information favors B, producing an evidence step back to neutral. The process continues until one of the critical values is reached, at which point the respective response is chosen, i.e., as soon as $X_n = \theta_A = +4$ or $X_n = \theta_B = -4$. In the case of trial 1, the evidence is sufficiently strong to make a decision in favor of B after the eighth sample. In the second trial there was enough information accumulated in favor of B after the sixth sample; in the last trial there was sufficiently strong evidence to make a decision in favor of B after the tenth sample. For the last example, the cumulative amount of information at each time is $X_0 = 0; X_1 = -1; X_2 = 0; X_3 = -1; X_4 = -2; X_5 = -1; X_6 = -2; X_7 = -3; X_8 = -4; X_9 = -3; X_{10} = -4$.

In the subsequent sections we will provide many relevant mathematical concepts that are related to the probabilistic modeling of decision-making. In any modeling problem, it is of paramount importance to underline the mathematical underpinnings of the model. This is necessary to see exactly how our assumptions about the decision-making processes are interpreted in mathematical formalism.

This will also help to find the pitfalls of such models. Apart from a rigorous mathematical theory, we will also discuss tools that can be helpful for practitioners. Most of these tools are standard to mathematicians, but may not be as well known to math users, hence their inclusion in this chapter. In the next section we will discuss the notion of Markov chain, which is the set-up for most discrete decision-making models.

1.3 Markov Chains

In this section we will formalize and extend the ideas discussed in subsection 1.2.2. As indicated above, in many modeling problems, one deals with sequences of random quantities

$$X_1, X_2, \ldots, X_n, X_{n+1}, \ldots$$

whose values are revealed with the time progress. The most basic situation is about a sequence of *independent* quantities. This means that a full knowledge of $X_1, \ldots, X_n$ does not influence the distribution of the variable X_n, or more precisely,

$$\mathbf{Pr}\left[X_n = s_n | X_{n-1} = s_{n-1}, \ldots, X_1 = s_1\right] = \mathbf{Pr}\left[X_n = s_n\right]. \tag{1.3}$$

Here we use the notation $\mathbf{Pr}[A|B]$ to denote the conditional probability[1]. Note that the equality (1.3) indicates that the extra information provided by $X_{n-1} = s_{n-1}, \ldots, X_1 = s_1$ has no bearing on predicting the value of X_n. The proverbial tossing of coins is perhaps the most well-known example of such processes. The sequence $\xi_1, \xi_2, \ldots$ discussed in the previous section is an example of such a sequence.

The theoretical simplicity is one of the appeals of such processes. However, in many real-world situations, one needs to deal with processes in which the distribution of the quantity X_n at time n somehow depends on that of the history of the process. As a motivating example, consider the following situation. Let P_d denote the price of a unit of the stock of company C on the dth day of the year. It is intuitively clear that the sequence $P_1, P_2, \ldots$ does not behave as a sequence of independent quantities. In fact, one expects the value P_{n+1} to be highly correlated with the values of P_m, at least for values of m close to n, even though this dependence is not deterministic. In other words, the knowledge of the value of the stock on days $d = 1, \ldots, n$ is relevant for predicting its value on day $d = n + 1$. Here is one way of formulating this property. We assume that the value X_n of the sequence at time n is given by

$$X_{n+1} = F(X_n, \ldots, X_{n-k}, S_n),$$

where S_n is a *random seed*, that is, a source of randomness independent of the sequence X_m. For instance, in the example of the random walk described in the previous section, we have

1 Recall that the conditional probability of A given B is defined by the ratio $\mathbf{Pr}[A|B] = \mathbf{Pr}[A \cap B]/\mathbf{Pr}[B]$ for $\mathbf{Pr}[B] > 0$.

$$X_{n+1} = X_n + \xi_{n+1}.$$

In this example, the step taken at time $n+1$ is the random seed involved in determining the position of the chain at time $n+1$.

Note that we have postulated that the values of $X_1, \ldots, X_{n-k-1}$ do not directly affect the value of X_{n+1}. In other words, the system has a *limited memory*.

A particularly important case is when $k = 1$. In this case, X_n depends on X_{n-1} and the random seed S_n, but does not depend on the "older history" consisting of the values $X_{n-1}, \ldots, X_1$. A succinct (although not precise) description is that the only part of history relevant to predicting future is the current state. In the next section we will make this notion more precise.

1.4 Markov Property

We will begin with the formal definition of Markov chains. We will then proceed to some examples.

Definition 1.1 Let S be a finite set. A sequence $X_1, X_2, \ldots$ of random variables which take values in S is called a Markov chain if for all $n \geq 1$ and all $s_1, \ldots, s_n \in S$ we have

$$\mathbf{Pr}\left[X_n = s_n | X_{n-1} = s_{n-1}, \ldots, X_1 = s_1\right] = \mathbf{Pr}\left[X_n = s_n | X_{n-1} = s_{n-1}\right]. \tag{1.4}$$

This condition is often referred to as the *Markov property*. Let us compare (1.4) to (1.3). Suppose that the elements of S correspond to the states of a system that evolves with time, in that we think of X_n as the *state of the system at time n*. Note that we have implicitly assumed that the time is discrete. Later we will consider analogous situations in which the time is assumed to be continuous. In (1.3), the entire history of the past is considered to be irrelevant for the prediction of the future. In (1.4), however, we are assuming that the value of X_n may depend on the past through its current value. Note that the information given by

$$X_{n-1} = s_{n-1}, \ldots, X_1 = s_1$$

describes the full history of the system up to time $n-1$. The Markov property simply states that the distribution of X_n only depends on its past inasmuch as it determines the value of X_{n-1}. In other words, the history of the system prior to time $n-1$ is only relevant through the value of X_{n-1}. As the system evolves with time, all the history except for the current state is erased. To study a Markov chain we will need to work with the transition probabilities defined next:

Definition 1.2 The *transition probabilities* of the Markov chain (X_n) are defined by

$$p_{ij}^{[n]} = \mathbf{Pr}\left[X_n = j | X_{n-1} = i\right].$$

In most of the applications in decision theory, we work with Markov chains that are *time-homogeneous*. This condition means that the transition probability $p_{ij}^{[n]}$ does not depend on the time n in which the transition is happening. In such cases, the superscript n is suppressed and the shorthand notation p_{ij} will be used.

Given a time-homogeneous Markov chain, a systematic way of keeping track of the transition probabilities is to put them in an $N \times N$ matrix, where N denotes the number of states. From now on, we will always describe a time-homogeneous Markov chain with N states $s_1, \ldots, s_N$ by an $N \times N$ matrix

$$\mathbf{P} = (p_{ij})_{1 \leq i,j \leq N}$$

where p_{ij} denotes the transition probability of the chain from state s_i to state s_j. The matrix $\mathbf{P}$ defined above is called the *transition matrix* of the Markov chain. For instance, p_{ii} is the probability that the Markov chain stays put, given that it is at state s_i. In the following theorem, we sum up two of the fundamental properties of the transition matrix.

Theorem 1.3 *The transition matrix* $\mathbf{P}$ *of any Markov chain satisfies the following two properties:*

1. *The transition probabilities are non-negative, that is, for all* $1 \leq i, j \leq N$,

$$p_{ij} \geq 0.$$

2. *Each row adds up to* 1, *that is, for* $1 \leq i \leq N$, *we have*

$$\sum_{j=1}^{N} p_{ij} = 1.$$

The first property is obvious. Let us explain the second property. Note that the entries in the ith row of $\mathbf{P}$ indicate the probabilities of making a transition from a fixed state s_i to one of the other states s_j, where j could vary. Because exactly one of these transitions will happen, the law of total probability implies that the sum of respected probabilities is 1, that is

$$\sum_{j=1}^{N} p_{ij} = \sum_{j=1}^{N} \mathbf{Pr}\left[X_n = j | X_{n-1} = i\right] = 1.$$

A matrix with these two properties is called a *stochastic matrix*. So, the transition matrix of any Markov chain is a stochastic matrix. Conversely, it is easy to see that if $\mathbf{P}$ is a stochastic matrix, one can define a Markov chain with transition matrix $\mathbf{P}$. Note that the same cannot be said about the columns. In some of the examples below we will encounter Markov chains where the sum of entries of a column is not equal to 1.