TOPICS
TOPICS IN COGNITIVE SCIENCE

Topics in Cognitive Science 6 (2014) 108–113

ISSN:1756-8757 print / 1756-8765 online
DOI: 10.1111/tops.12070

Quantum Random Walks and Decision Making

Karthik H. Shankar

Center for Memory & Brain, Boston University

Received 25 December 2012; received in revised form 1 March 2013; accepted 28 March 2013

Abstract

How realistic is it to adopt a quantum random walk model to account for decisions involving two choices? Here, we discuss the neural plausibility and the effect of initial state and boundary thresholds on such a model and contrast it with various features of the classical random walk model of decision making.

Keywords: Quantum random walks; Decision making

One-dimensional drift-diffusion models have been very successful in accounting for two crucial features of two-choice decision-making tasks, namely the decision accuracy and the decision latency. At a mechanistic level, these models can be implemented as a set of accumulators working in parallel, each gathering noisy evidence as a function of time. When the sum of evidences in all accumulators reaches a threshold boundary, a decision will be made. There are significant neural signatures in the brain for the process of evidence accumulation and the existence of a decision threshold. For example, when monkeys perform a perceptual decision task to infer the direction of motion of random dots, the LIP (lateral intraparietal area) is known to accumulate the evidence provided by the MT (middle temporal) cortex, and the decision moment synchronizes with the time at which the firing rate of the LIP neurons reach a threshold value (Gold & Shadlen, 2007).

Fuss and Navarro (2013) have proposed an interesting alternative to this classical decision mechanism. Instead of having the accumulators simply interact additively with each other, they consider the CCP (cooperative and competitive decision process) model where the accumulators can either cooperate with or inhibit each other. In summary, if the path of evidence accumulation in an accumulator shifts direction R times, then that accumulator is assigned a complex value of i^R, which determines the extent of cooperation of that accumulator with others. The crucial feature of the CCP model is that the rate of

Correspondence should be sent to Karthik H. Shankar, Center for Memory and Brain, Boston University, Boston, MA, 02215. E-mail: shankark@bu.edu

evidence accumulation toward the boundary is directly proportional to time, as opposed to the classical diffusion model where it is proportional to square root of time. In effect, the CCP model could lead to quicker and more efficient decision making.

1. Neural plausibility

There are, however, some issues in understanding how the brain could implement the CCP model. Consider two accumulators that follow the exact same evidence path up to some time t_0, beyond which they deviate. According to the CCP model, the two accumulators should cooperate until t_0 and then stop cooperating at the next instant. Clearly, this cannot be achieved by simple excitatory and inhibitory connections between groups of neurons representing the two accumulators, because this would require the synaptic weights to drastically (but precisely) vary in time depending on the activity history of the accumulators. It appears that any plausible implementation of the CCP model would require each accumulator to provide a four-way output (corresponding to $+1$, $+i$, -1, $-i$) in which only one of the four can be active at any moment, depending on the number of direction shifts in the activity history of that accumulator. This can potentially be realized as a winner-take-all competition among four output neurons associated with each accumulator. The momentary change in the evidence of each accumulator can then be used to bias the competition among the corresponding four neurons so as to have the appropriate neuron win the competition. Once this is achieved, the activity of each of the four neurons must be separately gathered from all the accumulators, and a global pairwise inhibition needs to be performed—the global outputs corresponding to $+1$ and -1 should inhibit each other and the global outputs corresponding to $+i$ and $-i$ should inhibit each other. Such a network relies on extremely fine-tuned connectivity, raising concerns on its neural plausibility.

2. Initial state of the random walk

Let us now ignore the concern of neural plausibility and simply focus on the utility of quantum random walks in modeling the process of decision making. Fuss and Navarro (2013) point out that the CCP model can be reduced to a problem of quantum random walk. So, for the purposes of computing decision accuracies and decision latencies, it is sufficient to simulate quantum random walks without invoking the parallel array of accumulators. Quantum random walk involves a two-component spinor function $(f_R(x, t), f_L(x, t))$ that deterministically evolves (see eq. 23 and 24 in Fuss & Navarro, 2013) as a function of time over a one-dimensional evidence space x, from an initial state at $t = 0$. At any instant, the spinor function gives the probability amplitude for various evidence values x. The quantum random walk moves the probability density toward the left and right boundaries in a bimodal fashion, which is qualitatively very different from the unimodal evolution of probability density in classical random walks (see Fig. 7 and 8 in Fuss & Navarro, 2013).

The only degree of freedom available in modeling the decision process as a quantum random walk is the choice of the initial state ($f_R(x, 0)$, $f_L(x, 0)$). Fuss and Navarro (2013) take the initial state to be ($\Phi(x)$, $e^{-i\theta}\Phi(x)$), where θ is a free parameter ranging between 0 and 2π, and $\sqrt{2}\Phi(x)$ is the initial probability density over evidence space. We shall refer to this as model 1. When $\theta = \pi/2$ or $3\pi/2$, the probability density symmetrically moves toward the two boundaries in an unbiased fashion, and when $\theta = 0$ or π, the probability density preferentially moves toward one of the boundaries with maximum bias (Kempe, 2003). The top row of Fig. 1 shows the probability density over evidence space after 100 steps of quantum walk starting from an initial Gaussian probability distribution (shown in gray) with a mean evidence at $x = 0$ and a variance of 4. For any value of θ, we see two probability peaks—one in the left and one in the right of the initial probability peak. For all practical purposes, the entire probability lies within these two peaks with apparently negligible probability in between them. For illustration, let us take the right boundary at $x = 100$ as the correct decision boundary and the left boundary at $x = -100$ as the error decision boundary. The ratio of the right peak to the sum of right and left peaks gives an estimate of the decision accuracy. In the unbiased case ($\theta = \pi/2$), this ratio is 0.5, and hence the decision accuracy is 50%. Interestingly, in the case of maximum bias ($\theta = 0$), there is a significant nonzero probability under the left peak leading to an accuracy of only 85%—the maximum accuracy that can be achieved in this model. In the simple classical case of Bernoulli random walk, the parameter q (that denotes the probability of moving in the right) can be varied between 0 and 1 to achieve any accuracy between 0% and 100%. But in model 1, there is a constraint on the maximum accuracy (85%) that can be achieved. This seems like an unnecessary constraint that can be eliminated by considering a different parameterization for the initial state.

In model 1, the initial states of both f_R and f_L have equal probability amplitudes of $1/\sqrt{2}$. A more general version of this model that allows for different initial probability amplitudes could yield different asymptotic distributions. Let us consider the initial state to be $(\frac{r}{\sqrt{l^2+r^2}}\Phi(x), \frac{l}{\sqrt{l^2+r^2}}\Phi(x))$, where $\Phi(x)$ is the initial probability density over evidence space, and l and r are real numbers. Assuming that the right boundary is the correct boundary, let us fix $r = 1$ and treat l as a free parameter ranging between $+1$ and -1. We shall refer to this as model 2. The bottom row of Fig. 1 shows the probability density over evidence states for model 2 after 100 steps of quantum walk starting with the same initial probability state as before. Note that for $l = 0.4$, the accuracy is practically 100% and for $l = -0.4$, the accuracy is about 50%. Thus, model 2 does not face the constraint faced by model 1.

3. Choice of boundary

Let us now focus on the choice of the decision boundary and demonstrate qualitatively distinct behaviors with small and large boundaries. We previously observed that most of the probability lies under the two peaks with practically no probability in between them. This is, however, true only when the number of time steps is small. After a large number

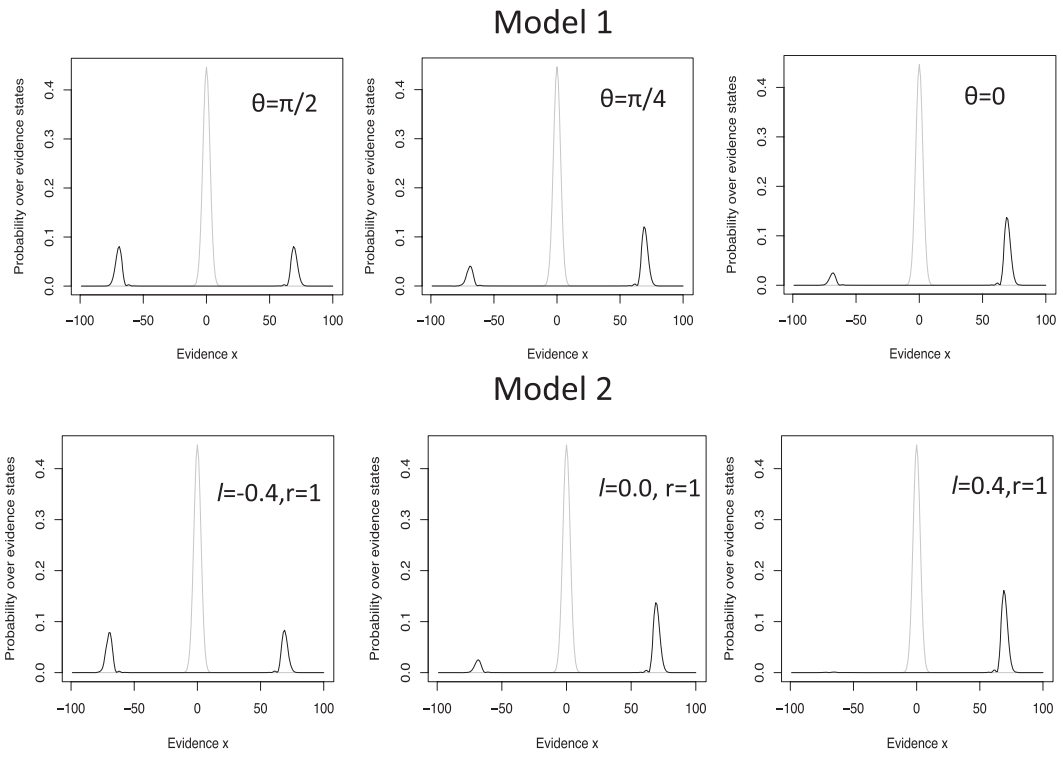

Fig. 1. The probability distribution after 100 time steps. The initial state of probability distribution is shown in gray. The top row shows model 1 for various values of θ, and the bottom row shows model 2 for various values of *l*.

of time steps, both the peaks shrink in height and spread in width. In fact, it can be analytically shown that after a large number of time steps N, the probability distribution is essentially a uniform distribution in the interval between $x = -N/\sqrt{2}$ and $x = +N/\sqrt{2}$ (Ambainis et al., 2001). Hence, it is incorrect to view the probability distribution as two separate peaks after a large number of time steps. As the two peaks move in opposite directions, it is easy to see that the overall variance of the probability distribution increases proportional to N^2. But it is rather subtle to realize that the variance of each peak separately increases with N.

To quantify this, let us consider the case of $\theta = \pi/2$ in model 1 (top-left panel of Fig. 1). Fig. 2 plots the variance of the probability distribution restricted to the positive x-axis as a function of the number of time steps N, so that only the right peak is taken into consideration. It is clear that the variance of the right peak increases linearly with N for large N. This is a clear indication that for large N, the probability distribution indeed becomes uniform. So, if we choose the decision boundary a to be very large, then the probability of finding the evidence state at a after N time steps, $p_a(N)$, increases from 0 to a small value when $N = a\sqrt{2}$ and then gradually decays off inversely proportional to $\mathcal{N}$. On the other hand, if we choose a to be small, then $p_a(= N)$ has a very narrow

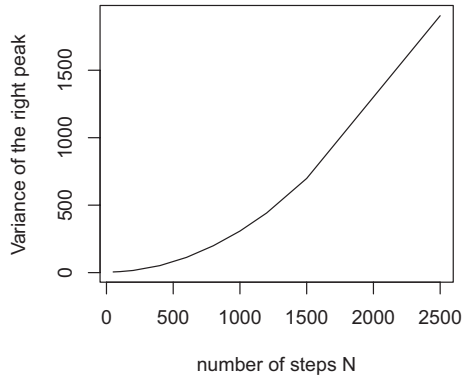

Fig. 2. Variance of the probability distribution restricted to the positive evidence states for $\theta = \pi/2$ in model 1 is plotted as a function of the number of time steps.

support near $\mathcal{N} = a\sqrt{2}$. To give a quantitative perspective on the size of the boundary, for the illustrations shown in Fig. 1 with initial Gaussian probability state of variance = 4, 50 is a small value for the boundary while 5000 is large. The choice of the boundary thus plays a qualitatively more interesting role in quantum random walk models than in the classical drift diffusion models.

4. Wave function collapse–decision mechanism

Finally, to fully validate a decision-making model, it is not sufficient to just account for the probability distribution of decision latency and accuracy; we need a mechanistic account of how the choice is made in each trial. In classical drift diffusion models, the evidence moves toward only one of the boundaries in each trial and hence the choice in each trial is straightforward. But in the CCP model and its simplified equivalent of quantum walk model, the single-trial description is ignored and only the probability density over evidence states is described. This could potentially be augmented with a separate collapse mechanism which at each moment reads out the probability of the evidence state being at either boundaries, and probabilistically generate one or none of the choices. Only with such an augmentation can the CCP model be considered as a mechanistic account of decision making.

Quantum physics faces a similar situation. The behavior of an electron in a given trial can only be described probabilistically because its wave function as a whole is not measurable. Any collapse mechanism that collapses the wave function in each trial is believed to act instantaneously over all space leading to a tussle with Einstein's Special Relativity. This makes it impossible to provide a satisfactory collapse mechanism in quantum physics. But the problem is much simpler for cognitive decision making. A spatially local collapse mechanism can certainly be prescribed to augment the CCP model. However, in light of existing neural signatures of evidence accumulation to a threshold

boundary (Gold & Shadlen, 2007), which supports classical drift diffusion models, the plausible neural signatures of a collapse mechanism and the CCP model needs a lot of scrutiny.

Acknowledgments

The author acknowledges support from National Science Foundation grant, NSF BCS-1058937, and Air Force Office of Scientific Research grant AFOSR FA9550-12-1-0369. The author also thanks the Aspen Center for Physics and the NSF Grant - 1066293 for the hospitality while this work was done.

References

Ambainis, A., Bach, E., Nayak, A., Vishwanath, A., & Watrous, J. (2001), One-dimensional quantum walks, *Proceedings of the thirty-third annual ACM symposium on theory of computing*, 37–49.

Fuss, I. G., & Navarro, D. J. (2013). Open parallel cooperative and competitive decision process: A potential provenance for quantum probability decision models, *Topics in Cognitive Science*, *5*(4), 818–843.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574.

Kempe, J. (2003). Quantum random walks: and introductory overview. *Contemporary Physics*, *44*, 307–327.