Computation & Cognition: Assignment 3
Decision Making

Group 6

March 2025

First Name	Last Name	Student Number
Diego	Cañas Jiménez	5621636
Hao	Chen	3990788
Adrien	Im	3984389
Koorosh	Komeilizadeh	3893995
Yihui	Peng	3985571

1 Getting Started

To better understand decision-making, researchers have developed evidence accumulation models, which describe decision making a gradual process in which evidence is gathered and evidence is integrated until a specific threshold is reached. One of the most important model in this framework is the Diffusion Decision Model (DDM).

In this assignment, we explored DDMs mainly through two papers: Bogacz et al. (2006) [1] and Forstmann et al. (2016) [2]. They provide a detailed view of the DDM, from its mathematical definitions, the theoretical background behind them, and also the applications of this model in cognitive science. Bogacz et al. describes that this model is especially useful in Two-Alternative Forced-Choice Task (TAFC), where individuals choose one of two responses based on accumulated evidence [1].

The second paper by Forstmann et al. expands on this by discussing the advantages of sequential sampling models. It explains how the DDM decomposes the decision making process into more manageable components. These components include drift rate, thresholds, and non-decision time [2]. These components, and their effects on decision time will be discussed more in detail in the following sections.

2 Simulating a Random Walk

Before implementing a full evidence accumulation model, we first simulate a random walk. This serves as a baseline for how evidence changes over time when there are no biases. This will allow us to have a foundation for later improved models, where decision thresholds and drift rates will be introduced. A random walk is a stochastic process, where the position of the variable changes over time with the accumulation of small random increments. We implemented a random walk by iteratively updating a variable x with Gaussian noise at each step. This can be characterized mathematically by the following equation:

$$x_{i+1} = x_i + cdW \tag{1}$$

where dW represents a small random increment that is scaled by a factor of $c \cdot dt$.

In our simulation, we initialized $x = 0$, and updated it over 5000 steps. We distinguished between the variables x and xs. x describes the current position in the random walk at each time step, while xs is a list that stores the sequence of values that x takes. The list xs shows us the entire trajectory. Furthermore, to ensure that the results are reproducible, we used the random number generated with a fixed seed: `np.random.seed(5)`. We then verified our implementation by reproducing the expected plot for a random walk.

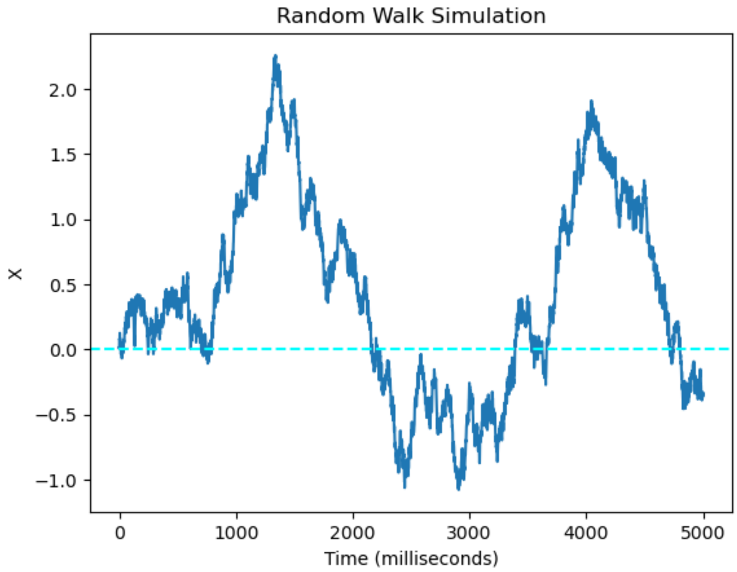

Figure 1: Random Walk Simulation. The plot illustrates the stochastic evolution of x over time. The x-axis represents time in milliseconds, while the y-axis shows the accumulated value of x. The dashed line at $y = 0$ serves as a reference for the mean position.

As shown in Figure 1, we were successful in reproducing the example plot, and that shows that the random walk behaves as we expected. The trajectory fluctuates symmetrically around zero due to the absence of drift and the variance increases. This is consistent with the theory. In conclusion, we have successfully implemented the random walk simulation, which will be extended in the following sections.

3 Adding a Threshold

To actually make a decision, this accumulated value can be used through a threshold. The threshold (a) determines when the decision process stops in the random walk model. Evidence accumulation continues until the x-value reaches either +a or -a, at which point we consider the model to have made a decision. By adjusting a, we observe that a lower threshold allows for faster decisions, while a higher threshold requires more time. This phenomenon can be explained by Figure 1: x undergoes random fluctuations symmetrically around zero. When the threshold is small, x reaches it very fast, leading to shorter decision times.

In this experiment, apart from the threshold a, the step size dt also plays a crucial role in determining decision time. In our implementation, we set dt = 0.001, meaning that each iteration of evidence accumulation corresponds to a time increment of 0.001 seconds.

The decision time can thus be calculated as:

$$\text{Decision Time} = N \cdot dt \tag{2}$$

where N represents the number of iterations required for x to reach either +a or -a.

Additionally, we set max_time_steps = 5000 to limit the maximum number of iterations, preventing cases where an excessively large a causes x to never reach a threshold within a reasonable time, thus avoiding infinite loops. From our simulation (code 2.3 of Assignment), we found that the mean decision time across 500 trials was approximately 0.276 seconds.

In each simulation, the final decision outcome is determined by which threshold, +a or -a, is reached first. If x reaches +a first, the trial is classified as a correct decision (upper); if x reaches -a first, it is classified as an incorrect decision (lower). According to the results from Code 2.3, under the conditions of threshold a = 0.5, noise coefficient c = 1, and step size dt = 0.001, the probability that x accumulates to -a (i.e., the error rate of the decision) is approximately 0.474.

From Figure 2, we can observe that under the experimental conditions of threshold a = 0.5, noise coefficient c = 1, and step size dt = 0.001, the majority of decisions occurred within 0.1 to 0.3 seconds, with the highest frequency observed between 0.1 and 0.2 seconds (50 occurrences).

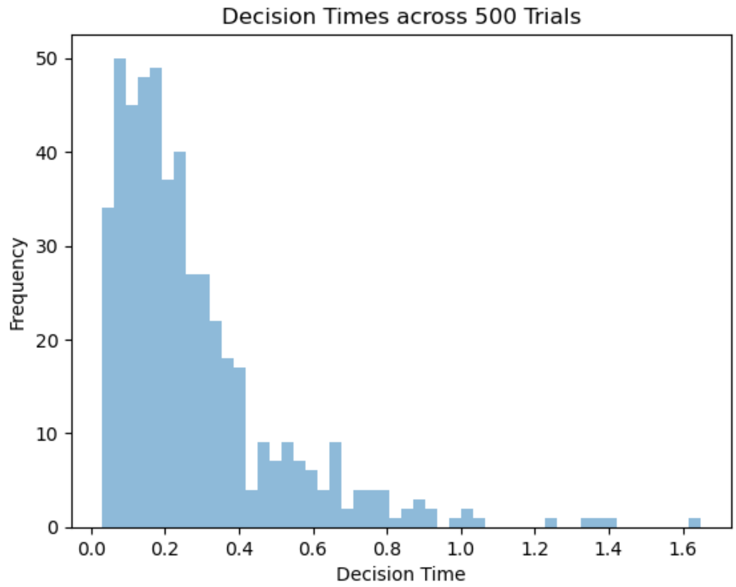

Figure 2: Decision Times across 500 Trials (a=0.5, c=1, dt=0.001). This histogram represents the decision time required for each of the 500 simulations. The x-axis represents the decision time and the y-axis represents the frequency of occurrence of each decision time.

A smaller proportion of trials had decision times exceeding 0.4 seconds, with the longest recorded decision time reaching 1.6 seconds. These longer decision times were rare, with each occurring fewer than 10 times across the 500 trials. The histogram exhibits a right-skewed distribution, indicating that: In most cases, evidence accumulation (x value) reaches the threshold quickly, resulting in shorter decision times. In some cases, the accumulation process takes longer, leading to extended decision times.

4 Adding a Drift

In reality, not all evidence has the same relevance, depending on the context, people are more inclined to one action. This can be referred to as a bias towards one action. To account for this bias, we will introduce drift rate (v) to simulate the potential bias that may exist in real-world decision-making processes. In the previous experiments, the update formula for x was:

$$x = x + \mathcal{N}(0, \sigma) \tag{3}$$

where $\sigma = \sqrt{c \cdot dt}$, since c and dt remain constant in each trial, meaning that $\mathcal{N}(0, \sigma)$ represents the Gaussian noise term, modeling the uncertainty in the evidence accumulation process. In this section, we modify the simulate_single_trial function by introducing drift rate v, leading to the updated equation:

$$x = x + v \cdot dt + \mathcal{N}(0, \sigma) \tag{4}$$

We can observe based on equation 4 when v > 0, evidence x accumulates more quickly towards +a(upper), i.e., the model is more likely to make the right decision, the decision time is shorter, and the error rate decreases. When v < 0, evidence x accumulates more quickly towards -a(lower), i.e. the probability of a wrong decision rises and the decision time may be long. By comparing the results of Code 2.3 (MDT = 0.275, ER = 0.474) and Code 3.2 (MDT = 0.821, ER = 0.126), we can conclude that introducing a drift rate of v = 1 into the random walk model significantly impacts both decision time and error rate:

- Mean Decision Time (MDT) increases substantially from 0.275s to 0.821s.
- Error Rate (ER) decreases significantly from 0.474 to 0.126.

This result confirms our hypothesis that introducing a drift rate v can significantly improve decision accuracy.

In the original random walk model (without drift), the evidence accumulation is purely stochastic, meaning that x has an approximately equal probability of reaching either +a (correct decision) or -a (incorrect decision). As a result, the error rate was close to 50%. However, when v = 1 is introduced, the drift term $v \cdot dt$ systematically biases the accumulation process toward +a, making it more likely that x reaches the upper threshold. As a result, many trajectories that would have previously led to an incorrect decision at -a are now redirected towards +a, reducing the error rate.

Although the introduction of v improves the decision accuracy, it also increases the mean decision time. In the random walk model without drift, some evidence accumulation processes reach -a quickly and have a shorter decision time. When we applied a drift into the model, those trajectories that would have reached -a now require additional time to be "pulled" towards +a leading to an overall increase in decision time. This explains why, despite the reduction in error rate, the MDT increased significantly.

This increase in MDT and decrease in ER is consistent with the expected effects of drift in the diffusion model. To further assess the accuracy of our simulation, we compare our results with the analytical predictions of Bogacz et al. (2006). While the general trend aligns with theoretical expectations, numerical discrepancies exist: the simulated mean decision time (MDT) was higher than the theoretical prediction (0.589s vs 0.231s), and the error rate (ER) was lower than expected (0.126 vs.0.269).

5 The Effect of Threshold and Drift Rates

The drift-diffusion model (DDM) provides a quantitative framework to examine how parameters such as the drift rate(v) and decision threshold (a) influence both the speed and accuracy of decisions. To systematically evaluate these effects, two sets of simulations were conducted: one varying the drift rate (v = 1 and v = 2) while holding the threshold constant (a = 1), and another varying the threshold (a = 1 and a = 1.5) with a fixed drift rate (v = 1).

5.1 Effect of Drift Rate Increase (v=1 to v=2)

Increasing the drift rate from v=1 to v=2 significantly altered both the decision time distribution and error rates. The histogram of decision times (Figure 3) revealed a pronounced leftward shift, with the mean decision time decreasing from 0.82 seconds (v=1) to 0.41 seconds (v=2). This acceleration reflects the stronger directional evidence accumulation induced by higher drift rates, which reduces the time required to reach the threshold[1]. Concurrently, the error rate dropped from 12.4% to 3.6 %, as the increased drift rate diminished the relative impact of noise, making it less probable for the process to cross the incorrect lower boundary. These findings align with the theoretical predictions[1], where higher v reduces both the mean decision time and error rate.

5.2 Effect of Threshold Increase (a=1 to a=1.5)

In contrast, raising the threshold from a=1 to a=1.5 while maintaining v=1 produced opposing effects. The decision time histogram (Figure 4) shifted rightward, with the mean decision time increasing from 0.82 seconds to 1.24 seconds. A higher threshold necessitates more accumulated evidence to trigger a decision, thereby prolonging the deliberation period. However, this cautious strategy also reduced the error rate from 12.4% to 6.8%, as the extended accumulation window allows the drift term to dominate over stochastic noise. This trade-off between speed and accuracy is a hallmark of threshold adjustments in sequential sampling models[2].

5.3 Divergent Roles of v and a

The distinct impacts of v and a highlight their unique roles in decision dynamics. While drift rate (v) primarily governs the efficiency of evidence accumulation—enhancing both speed and accuracy—the threshold (a) modulates the caution of the decision-maker. Higher v optimizes performance by reducing noise susceptibility, whereas higher a prioritizes accuracy at the cost of slower responses. This dichotomy underscores the DDM's ability to disentangle cognitive processes such as stimulus processing (v) and response caution (a), offering a mechanistic explanation for empirical speed-accuracy trade-offs observed in human behavior.

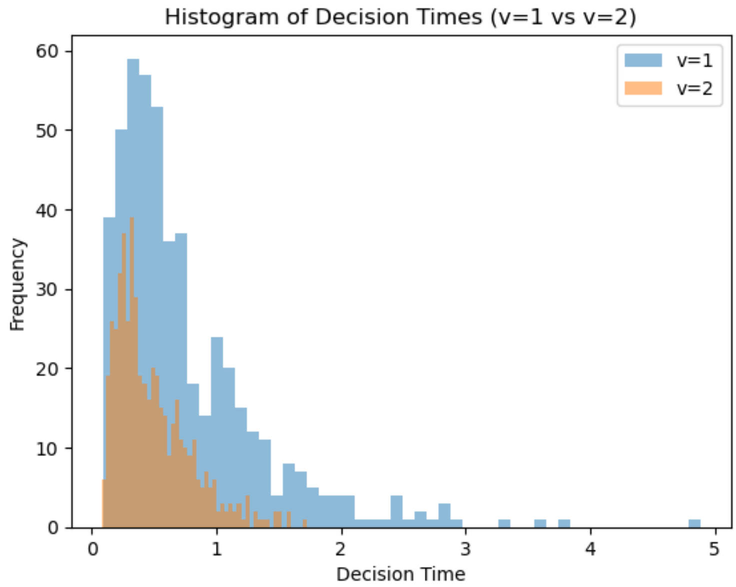

Figure 3: Comparison of decision times across 500 trials for different drift rates v=1 in blue and v=2 in orange, and a fixed threshold of a=1. The x-axis represents the decision time and the y-axis represents the frequency of occurrence of each decision time. Represents the reduction of relevance of noise with the increase of drift rate, leading to faster and more precise action selections.

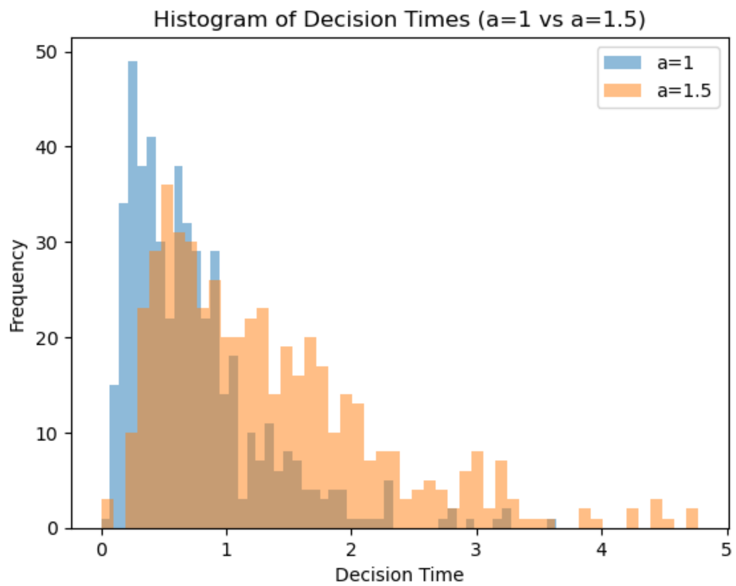

Figure 4: Comparison of decision times across 500 trials for different threshold values a=1 in blue and a=1.5 in orange, and fixed drift rate of v=1. The x-axis represents the decision time and the y-axis represents the frequency of occurrence of each decision time. The graph reflects the tradeoff between precision and speed of crossing the threshold.

6 Evaluation

6.1 Utility of Decision Time Histograms

All in all, the analysis of decision time histograms serves as a critical validation tool for the DDM. Human decision-making experiments consistently report right-skewed reaction time distributions with long tails, a pattern replicated in the simulated histograms. Such distributions arise naturally from the stochastic accumulation process, where noise introduces variability in the time required to reach a threshold. By comparing simulated histograms to empirical data, researchers can assess whether the DDM captures key features of human behavior, such as the distribution shape and parameter-dependent shifts. Furthermore, histograms reveal subtle effects of parameter changes—for instance, the narrowing of the distribution under higher v reflects reduced variability in high-confidence decisions, a phenomenon observed in perceptual tasks [2].

6.2 Challenges and Limitations

Some challenges and limitations we faced include implementing the DDM presented several challenges. First, translating the continuous stochastic differential equation ($dx/dt = vdt + cdW$) into discrete simulations required precise handling of noise scaling $\left(\sqrt{c \cdot dt}\right)$ to ensure numerical stability. Second, edge cases such as v=0 necessitated ad-hoc solutions to avoid division-by-zero errors in analytical calculations. Third, the simulations implicitly excluded trials where thresholds were not reached within the 5-second limit, potentially biasing estimates of mean decision times. A more robust approach would explicitly account for such trials, either by extending the time window or incorporating a "non-decision" component.

6.3 Conclusion

In conclusion, the DDM's capacity to generate human-like decision time distributions and parameter-dependent effects underscores its value as a cognitive model. By systematically manipulating v and a, this study demonstrated how drift rates and thresholds dissociate speed and accuracy—a cornerstone of evidence accumulation theory. Future work could extend these simulations to incorporate trial-to-trial variability in v or a, better mirroring the dynamic nature of real-world decision tasks.

References

[1] R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J. Cohen, "The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks," *Psychological review*, vol. 113, pp. 700–65, 10 2006.

[2] B. Forstmann, R. Ratcliff, and E.-J. Wagenmakers, "Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions," *Annual Review of Psychology*, vol. 67, no. Volume 67, 2016, pp. 641–666, 2016. [Online]. Available: https://www.annualreviews.org/content/journals/10.1146/annurev-psych-122414-033645