Full Bayesian Model and Reduced Baysesian Delta-Rule Model Explaining the Dynamics of Belief Updating in a Changing Environments

Heejin Jeong

Reference Paper: An Approximately Bayesian Delta-Rule Model Explains the Dynamics of Belief Updating in a Changing Environment, *Matthew R. Nassar, Robert C. Wilson, Benjamin Heasly, and Joshua I.Gold*, Journal of Neuroscience 2010

1 Introduction

In order to effectivley make decisions, it is important to maintain and update beliefs appropriately. In unchanging environment, beliefs depends on past experiences. However, in dynamic environment, belief updates may not rely on past experiences too much but instead consider new outcomes more, espically which are unexpected. One common algorithm for such adaptation is the delta rule :

$$B_{t+1} = B_t + \alpha_t \times \delta_t \qquad \text{where } \alpha_t \in [0,1]$$

where B_t is belief at time t, α_t is a learning rate at t, and δ_t is the error made in predicting the most recent outcome. Therefore, when $\alpha_t = 0$, the belief does not take any information from the most recent outcome at time t. On the ohter hand, when $\alpha_t = 1$, the belief ignores the previous belief and fully rely on the most recent outcome.

The prediction errors can be caused by two main sources. One is stochastic fluctuation, or noise, which makes good predicitons poor. Therefore, new outcomes with noise should be considered minimally when updating beliefs. The other source of the prediction errors is a fundamental change point, or volatility, in the action-outcome relationship. Change points indicate that previous outcomes are irrelevant to the new environment after the most recent change point, and thus new outcomes should influence beliefs strongly.

Previous work showed that human subjects elevate learning rates during periods of volatility on probabilisitc decision tasks. Such behavior can be fit by Bayesian model for optimal belief updating and extension of Delta-rule updating (computationally frugal). In the reference paper, they developed a novel task to

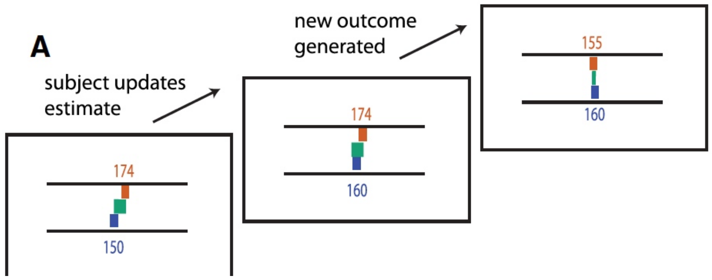

Figure 1: Schematized trial of the estimation task

characterize how well human subjects follow these priniciples under different conditions. One of their main results is that such behaviors were consistent with a modified delta-rule model, derived from a systematic reduction of the Bayesian ideal observer. Therefore, in this project, the Bayesian ideal observer model and the reduced Bayesian model (th modified delta-rule model) were studied and simulated in MATLAB.

2 Behavioral tasks

In this section, the novel task introduced in the reference paper is briefly described.

Estimation task: Subjects were required to predict each subsequent number to be presented in a series of numbers. At each trial t, a number(X_t) is presented that was a rounded pick sampled independently and identically from Gaussian distribution with mean μ_t and standard deviation σ_t (X_t $\mathcal{N}(\mu_t, \sigma_t)$). The mean, μ_t, is changed at unsignaled change points while σ_t was fixed for each of 4 experimental blocks of 200 trials(5,15,25,or 35). Change points in the mean of the generative distribution occurred after at least five trials plus a random pick from an exponential distribution with a mean of 20 trials.

$$\text{rate of change points} = \text{hazard rate}, H = \begin{cases} 0 & \text{for first 5 trials} \\ 0.05 & \text{otherwise} \end{cases}$$

Subjects were told that the numbers were generated from a noisy process and they were instructed to minimize their prediction errors on average across all blocks of the task. Payout depended on how well they achieved this goal. The Figure.?? is an example of display shown to the subjects. The red bar represents the most recent number presented, the blue bar represents the current estimate and the green bar represents the error made in predicting outcomes. The subject updated his/her prediction on each trial to an integer value using video gamepad.

Confidence task: In this task, subjects are required to not only predict a number but also indicate their confidence. They use a symmetric window around their prediction with 85% confidence that the window would contain the next number. If the next number is generated and it falls within the specified window by a subject, then he or she earn points. Feedback includes a sound to indicate that a subject earned points and a running tally of points earned by the subjects.

3 Models

3.1 Full Bayesian Model

In order to make an optimal decision, we need to infer mean and variance of the probability distribution of the future outcomes given all previous outcomes, $p(X_{t+1}|X_{1:t})$. Bayesian inference can be accomplished by inverting an appropriate generative process. Since a number, X_t, is generated from a Gaussian distribution with mean μ_t, and the mean value depends on whether there was a change point or not, the change points has to be considered first in the generative process. One of ways of representing change points is defining a scalar variable describing the number of outcomes that occurred since the most recent change point ("run length", r). In this full Bayesian model, the generative process depends on run length primarily.

A time interval between two change points is sampled from an exponential distribution with mean of 20:

$$\tau_i = T_i - T_{i-1} \sim \text{exponential}(\lambda)$$

$$\mathbf{E}(\tau) = \frac{1}{\lambda} = 20$$

where T_i is the time of the ith change point. When there is no change point and the current run continues, $r_t = r_{t-1} + 1$ and $\mu_t = \mu_{t-1}$. When a change point occured, r_t is reset to 0 and mu_t is randomly sampled from the uniform distribution $U(0, 300)$. In either case, an outcome(X_t) is generated on each trial from a Gaussian distribution with the mean, μ_t, and the fixed variance (σ^2). This generative process framework is shown in the Figure.??.

With this generative framework, we can express the predictive distribution with possible mean and run

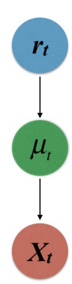

Figure 2: The Generative Framework of the Full Bayesian Model

lengths as :

$$
\begin{aligned}
p(X_{t+1}|X_{1:t}) &= \sum_{\mu_t} p(X_{t+1}, \mu_t|X_{1:t}) \\
&= \sum_{\mu_t} p(X_{t+1}|\mu_t)p(\mu_t|X_{1:t}) \\
&= \sum_{\mu_t} p(X_{t+1}|\mu_t) \sum_{r_t} p(\mu_t, r_t|X_{1:t})
\end{aligned}
$$

The joint distribution of the mean and the run length given previous data can be computed using Bayes' rule:

$$
\begin{aligned}
p(\mu_t, r_t|X_{1:t}) &= \frac{p(X_{1:t}|\mu_t, r_t)p(\mu_t, r_t)}{p(X_{1:t})} \\
&= \frac{p(X_t|\mu_t, r_t) \sum_{\mu_{t-1}} \sum_{r_{t-1}} p(\mu_t|r_t, \mu_{t-1})p(r_t|r_{t-1})p(\mu_{t-1}, r_{t-1}|X_{1:t-1})}{p(X_t|X_{1:t-1})}
\end{aligned}
$$

The prior, $p(\mu_t, r_t)$, can be marginalized across all possible previous mean and run length as in the equation above and $p(\mu_t|r_t, \mu_{t-1})$ and $p(r_t|r_{t-1})$ are simply transition probabilities of mean and run length. These transition probabilities are defined as below:

$$
p(r_t|r_{t-1}) = \begin{cases} 1 - H & \text{if } r_t = r_{t-1} + 1 \\ H & \text{if } r_t = 0 \\ 0 & \text{otherwise} \end{cases}
$$

$$
p(\mu_t|\mu_{t-1}, r_t \neq 0) = \begin{cases} 1 & \text{if } \mu_t = \mu_{t-1} \\ 0 & \text{if } \mu_t \neq \mu_{t-1} \end{cases}
$$

$$
p(\mu_t|\mu_{t-1}, r_t = 0) = p(\mu_t|r_t = 0) = 1/300
$$

Moreover, let's define B_t as a predictive distribution after trial t for the next outcome X_{t+1} and define A_t as a joint distribution of the mean and the run length after trial t given all previous data:

$$
B_t(X_{t+1}) = p(X_{t+1}|X_{1:t})
$$

$$
A_t = A_t(\mu_t, r_t) = p(\mu_t, r_t|X_{1:t})
$$

Then, we can write down a recursive equation as:

$$
A_t(\mu_t, r_t) = \frac{p(X_t|\mu_t, r_t) \sum_{\mu_{t-1}} \sum_{r_{t-1}} p(\mu_t|r_t, \mu_{t-1})p(r_t|r_{t-1})A_{t-1}(\mu_{t-1}, r_{t-1})}{B_{t-1}(X_t)}
$$

On the whole, if we apply the transition probability values in this equation, we have:

$$
A_t(\mu_t, r_t = 0) = \frac{N(X_t|\mu_t, \sigma_t^2) \sum_{\mu} \sum_{r} A_{t-1}(\mu, r)H}{B_{t-1}(X_t)}
$$

$$A_t(\mu_t, r_t \neq 0) = \frac{N(X_t|\mu_t, \sigma_t^2) \sum_r A_{t-1}(\mu_t, r)(1-H)}{B_{t-1}(X_t)}$$

Therefore, the solution for the Full Bayesian model is:

$$\hat{X}_{FB,t+1} = \text{argmax}_{i\{\in[0,300]\cap\mathcal{Z}\}} B_t(X_{t+1} = i)$$

3.2 Reduced Bayesian Model

In the reference paper, more computationally tractable and neurally feasible inference algorithm is introduced. This is a systematic reduction of the full Bayesian model and the idea is instead of computing predictive distribution across all possible run lengths, it is updated with respect to a single, expected run length ($\hat{r}$). The probability of a change point (cp) on a given trial is computed by Bayes' rule:

$$\begin{aligned}\Omega_t = p(cp|X_t) &= \frac{p(X_t|cp)p(cp)}{p(X_t)} \\ &= \frac{p(X_t|cp)p(cp)}{p(X_t|cp)p(cp) + p(X_t|\ cp)p(\ cp)} \\ &= \frac{U(X_t|0,300)H}{U(X_t|0,300)H + \mathcal{N}(X_t|\hat{\mu}_t, \hat{\sigma}_t)(1-H)}\end{aligned}$$

where $U(X_t|0, 300)$ is the uniform distribution from which X_t is generated if a change point occurred.

The predictive distribution is determined with a predicted mean and a predicted variance in Gaussian Distribution, $\mathcal{N}(X_{t+1}|\hat{\mu}_{t+1}, \hat{\sigma}_{t+1})$. The expected variance is determined by two main uncertainties. One uncertainty is about the outcome for the given mean, and the other uncertainty is about the actual location of mean:

$$\hat{\sigma}_{t+1}^2 = \sigma_t^2 + \frac{\sigma_t^2}{\hat{r}_t}$$

where σ_t is a standard variation of the generative distribution. The expected mean value of the predictive distribution, μ_{t+1}, is determined considering two possibilities, one that a change point occurred and the other that a change point didn't occur.

$$\begin{aligned}\hat{\mu}_{t+1} &= \mathbf{E}(\mu_{t+1}|X_t) \\ &= \mathbf{E}(\mu_{t+1}|X_t, cp)p(cp|X_t) + \mathbf{E}(\mu_{t+1}|X_t,\ cp)p(\ cp|X_t)\end{aligned}$$

Since when a change point occured, $X_t = \hat{mu}_t$ and $\hat{mu}_{t+1} = \hat{mu}_t$,

$$\mathbf{E}(\mu_{t+1}|X_t, cp) = X_t$$

and when there is no change point, the estimate of the mean keeps being update taking arithmetic mean with new outcome,

$$\mathbf{E}(\mu_{t+1}|X_t,\ cp) = \frac{X_t + \hat{r}_t \times \hat{\mu}_t}{\hat{r}_t + 1}$$

Therefore,

$$\hat{\mu}_{t+1} = X_t\Omega_t + \frac{X_t + \hat{r}_t \times \hat{\mu}_t}{\hat{r}_t + 1}(1 - \Omega_t)$$

This update equation can be rearranged as below, and it can be seen as delta-rule update:

$$\hat{\mu}_{t+1} = \hat{\mu}_t + \alpha_t \times \delta_t$$

where

$$\delta_t = X_t - \hat{\mu}_t \quad : \text{Prediction Error}$$

$$\alpha_t = \frac{1 + \Omega_t \hat{r}_t}{\hat{r}_t + 1}$$

Similarly, the expected run length can be computed as:

$$\hat{r}_{t+1} = (\hat{r}_t + 1)(1 - \Omega_t) + 1 \cdot \Omega_t$$

4 Simulation Results

4.1 Initial Condition

The initial conditions for the full Bayesian model is:

$$B_0(X_1) = p(X_1) = 1/300$$

$$B_1(X_2) = p(X_2|X_1) = \sum_{\mu_1} p(X_2|\mu_1)\big(p(\mu_1, r_1 = 0|X_1) + p(\mu_1, r_1 = 1|X_1)\big)$$

where $p(\mu_1, r_1 = 0) = p(\mu_1)p(r_1 = 0) = H/300$ and $p(\mu_1, r_1 = 1) = p(\mu_1)p(r_1 = 1) = (1 - H)/300$. For the reduced Bayesian model,

$$\Omega_1 = H$$

$$r_1 = 1$$

4.2 Prediction of FB and RB models

As same as the task, time intervals between two change points were generated by exponential distribution with mean 20. Since this is not compared to subject experiments, we ignored the first five training trials in the estimation task. The true means were generated from the uniform distribution at those change points, and thus a true number at each trial t was generated from Gaussian distribution with the mean at t and $\sigma = 5$. Total ten sets were simulated and in each set there were 200 trials.

The Figure.**??** and Figure.**??** show simulation results of two of the ten sets. Red line represents true numbers, blue line represents predicted numbers from the full Bayesian model (FB), and the green line

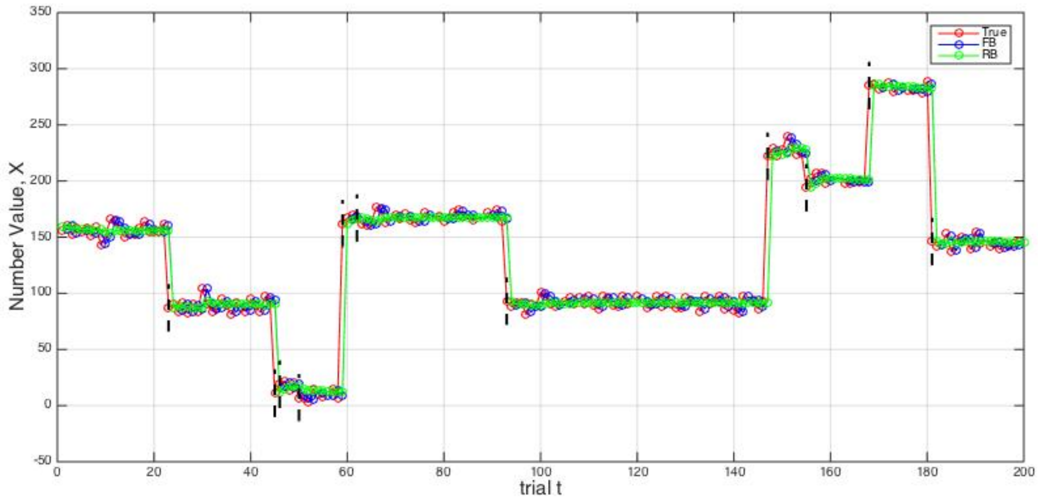

Figure 3: Simulation Result, Set 1. True numbers(red), Predicted values from the full Bayesian model(blue), Predicted values from the reduced Bayesian model(green)

represents predicted numbers from the reduced Bayesian model(RB). The black dot line indicates where change points occurred. Both FB and RB follows well the abrupt changes at the change points. However, during the non-change point periods, FB seems more sensitive to small changes than RB. The predicted values by RB were gradually changed. The Figure.**??** shows learning rate at each trial computed in RB. As expected, the learning rate goes to 1 after the change points occurred, which means the expected mean is updated only by the most recent outcomes. In order to compare the performance of the models, Mean Absolute Error (MAE) was computed. For T trials, MAE is:

$$\text{MAE} = \frac{1}{T}\sum_{t=1}^{T}|X_t - \hat{X}_t|$$

where X_t is a true value at t and $\hat{X}_t$ is a predicted value for t. The Figure.6 shows MAE of FB(1) and RB(2) for all sets.

$$MAE_{FB} = 9.2395 \quad MAE_{RB} = 8.5508$$

The Figure.7 and Figure.8 show the message passing algorithm graph of FB when no change point ocurred and when a change point occurred, respectively. The y-axis is the run length and the x-axis is trial. As the graph shows, in this model, distribution at each run length has to be computed at every trial. On the other hand, in RB, only expected run length is considered, thus much small number of run lengths are considered. The Figure.9 shows an example of message passing algorithm of RB drawn from one of the simulated sets.

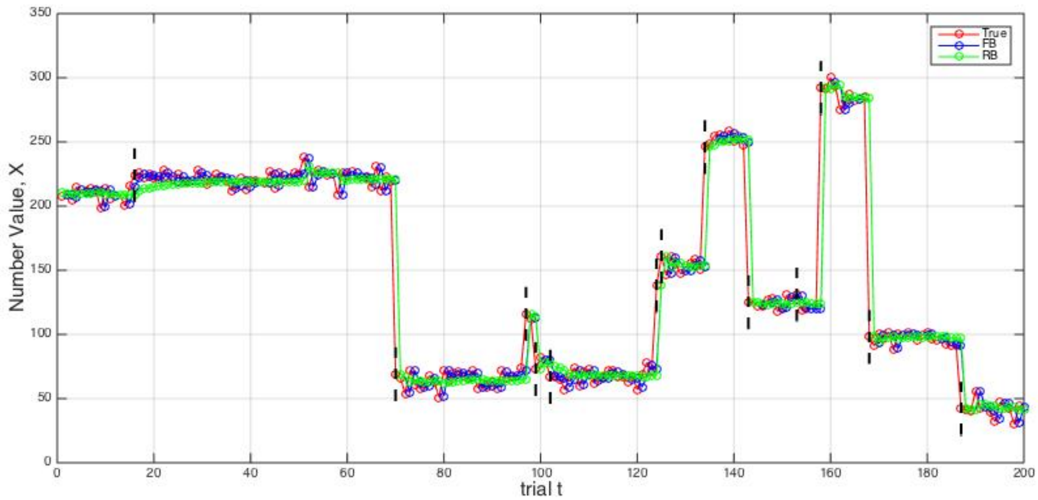

Figure 4: Simulation Result, Set 2. True numbers(red), Predicted values from the full Bayesian model(blue), Predicted values from the reduced Bayesian model(green)

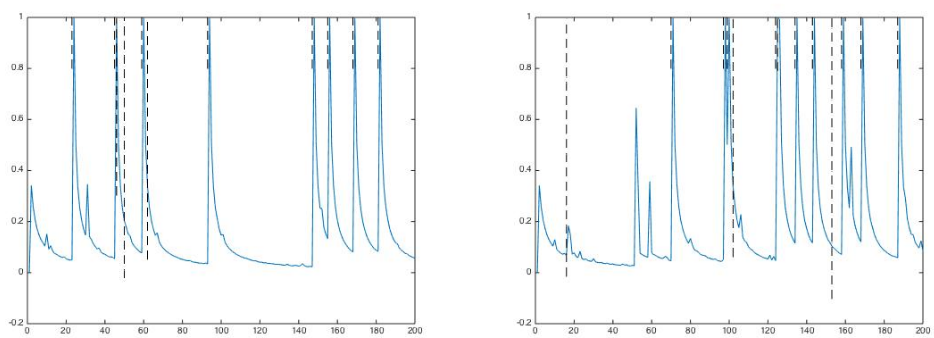

(a) Set 1

(b) Set 2

Figure 5: Learning rates computed in the reduced Bayesian model

5 Discussion

In this projects, we studied the full Bayesian model and the Reduced Bayesian model introuduced in the reference paper. In the reference paper, the trial subscripts of the expected mean and standard variation in the reduced Bayesian model were denoted as one step behind. Thus, here we used a corrected version of the model. As mentioned above, predicted values in RB were gradually changed and got closer to the true mean value, while predicted values in FB were relatively sensitive to the changes in true numbers. As a result, RB performed slightly better than FB unlike the result in the reference paper where FB performs less than 0.5

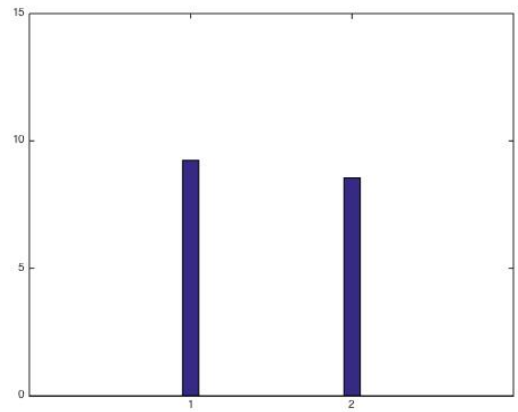

Figure 6: Mean Absolute Errors of FB and RB

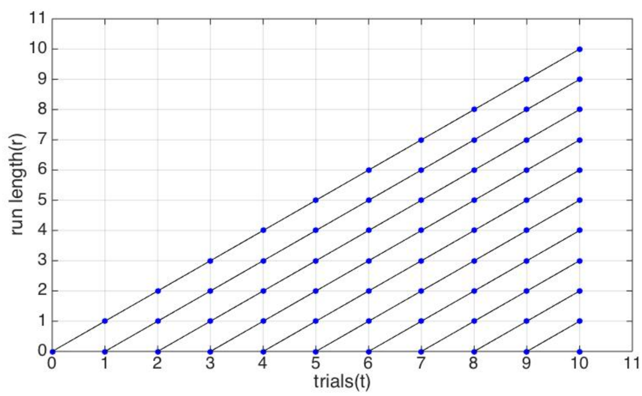

Figure 7: Message-passing Algorithm - FB

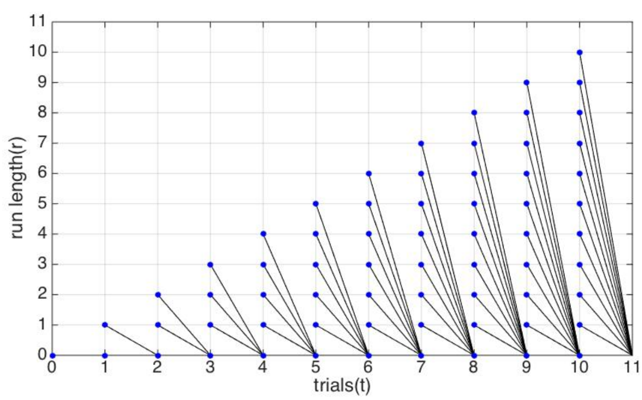

Figure 8: Message-passing Algorithm - FB

better than RB in MAE evaluations. However, as the paper suggested, the reduced Bayesian model took less time and require less computational capacity than the full Bayesian model.

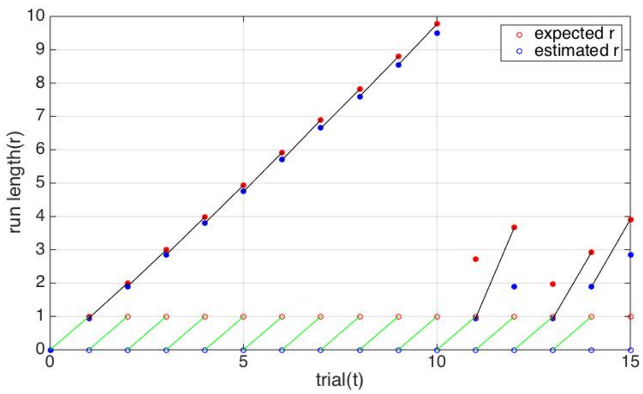

Figure 9: Message-passing Algorithm - RB, Black solid line : When no change point occurred, Green solid line : When a change point occurred