COGNITIVE BIASES IN BAYESIAN UPDATING AND OPTIMAL INFORMATION SEQUENCING

Sara Mourad, Ahmed Tewfik

Department of Electrical and Computer Engineering
The University of Texas at Austin

***Abstract*—In this paper, we consider the problem of optimally ordering information to a human subject to maximize detection performance in a binary hypothesis testing problem. We begin by proposing a modification of the traditional Bayesian solution to hypothesis testing problems to incorporate the effect of human cognitive biases. Next, we consider the problem of selecting a subset of information to maximize detection performance in truncated hypothesis testing problems. We then use the solution to that problem to determine the real time ordering of information to enhance human binary hypothesis testing. We verify through simulations that the proposed ordering methods with and without cognitive biases minimize the probability of miss and the probability of false alarm.**

Index — Cognitive biases, Bayesian inference, Likelihood, Ordering, i.i.d. observations

I. INTRODUCTION

Bayesian inference models provide optimal mathematical solution to the problem of inferring the right conclusions based on observing data. They detect the hypothesis with maximal odds of occurrence. However, these optimal solutions do not capture the way human beings integrate the pieces of information and make their decisions. Humans are subject to several cognitive biases, which are heuristics that hinders them from making rational decisions. It is thus important to study modifications of Bayesian decision making under the influence of cognitive biases, and to figure out further the best sequencing of information for humans in order to mitigate the biases.

For hypothesis testing, Neyman-Person test allows the detection of the hypothesis after observing N pieces of independent information. It computes a cumulative log-likelihood and compares it to a threshold which determines the probability of false alarm. This likelihood ratio test is the most powerful test with probability of false alarm less than or equal to a determined value. The sequential probability ratio test (SPRT) updates the cumulative log-likelihood ratio with each new observation and compares it to a lower and upper thresholds which guarantee the probability of false alarm and the probability of miss [11]. For binary hypothesis testing, if the cumulative log-likelihood ratio exceeds the lower or upper threshold, then hypothesis 0 or hypothesis 1 is chosen respectively and the test terminates. Otherwise, if it is between the two thresholds, the test continues and another observation is taken. This test is the optimal test for i.i.d. observations in terms of the expected number of observations needed to make a decision within the required probabilities of miss and of false alarm [11]. The optimality of the SPRT is extended to non-i.i.d. observations. For non i.i.d. observations, it is shown in [10] that time-varying thresholds are optimal over fixed-thresholds. In [8], the thresholds are constructed so that the probabilities of miss and false alarm are upper bounded at each iteration. It is also shown that ordering the observations in the non-decreasing order of the Kullback Leibler (KL) divergence minimizes the average sample number.

However, for these Bayesian inference models to represent how humans integrate information, cognitive biases should be incorporated. Extensive research on human beings and cognitive biases led to the conclusion that human inference is best described in terms of a list of biases and heuristics, regarded as illusions or discrepancies between what people should do and what they actually do [3]. As an example of cognitive biases, we cite the anchoring bias where humans are influenced by starting points or initial beliefs. In [4] and [6] the starting point bias is modeled by the impact of the initial bid value on the willingness to pay. The interviewer introduces successive bid values to a subject, and the initial bid value influences the posterior willingness to pay of the subject. Other biases are the heuristics that humans use in order to optimize the time required for a mental task. In [9], this is modeled by defining a cost function for the estimation error and the number of observations needed to estimate a quantity. The bias is incurred from optimizing the time-accuracy trade-off, and it is thus defined by the time cost and the error cost at the stage of estimating the value. In [5], the bias of overconfidence or underconfidence in current earning forecasts is modeled by introducing a multiplicative constant to the weight of the new observation during the Bayesian updating procedure. In [2], the belief update model is modified to account for the confirmation bias in the context of auditing: the sensitivity of the auditor to confirming and disconfirming information is represented by adequate multiplicative constant for the new observation in the anchoring and adjustment model used. Also in [1], the bias is modeled by modifying the values of the thresholds used in the sequential probability ratio test.

In this paper, we treat the problem of accounting for biases in Bayesian updating and finding out the optimal sequencing of i.i.d. information in the context of cognitive biases and without cognitive biases. In Section II, we review existing models of belief update in the literature with and without biases, and we propose our model: a modified version of the Bayesian updating to account for biases. In Section III, we propose an optimal sequencing of i.i.d. Gaussian observations without biases. We note that this is the first time that the order of i.i.d. observations is shown to have an effect on decreasing the probability of error. In Section IV, we propose an optimal sequencing of i.i.d. Gaussian observations in the context of cognitive biases. Section V concludes the paper.

II. MODELING COGNITIVE BIASES IN DECISION MAKING

In this Section, we will propose a modification of the Bayesian updating that can model decision making in the presence of biases. We will base our model on other models in the literature that incorporate biases.

Let's first introduce the belief-adjustment model proposed by Hogarth and Einhorn in [7] to represent belief revision in response to new evidence:

$$S_k = S_{k-1} + w_k s(x_k) \tag{1}$$

where :

S_k is the degree of belief in some hypothesis given k pieces of evidence, $0 \leq S_k \leq 1$,
$s(x_k)$ is the subjective evaluation of the kth piece of evidence, $-1 \leq s(x_k) \leq 1$,
w_k is the adjustment weight of the kth piece of evidence,
x_k is the kth piece of evidence.

The adjustment weight w_k for the kth piece of evidence is defined as:

αS_{k-1} for negative evidence (evidence supporting hypothesis $H0$), and
$\beta(1 - S_{k-1})$ for positive evidence (evidence supporting hypothesis $H1$)

where α is the sensitivity towards negative evidence($0 \leq \alpha \leq 1$),and β is the sensitivity towards positive evidence($0 \leq \beta \leq 1$).

The adjustment weight w_k is a function of the prior belief to account for the contrast effect: The higher the anchor, the bigger the adjustment. The importance of this model is that it specifically distinguishes the sensitivity to confirming evidence (β) from the sensitivity to disconfirming evidence (α). Bamber, Ramsay and Tubbs use this model to examine auditors' attitudes to evidence in [2].

The model we propose applies the framework of the belief-adjustment model of Hogarth and Einhorn to the Bayesian updating. Before proposing the model, we first consider a basic problem of testing two hypothesis $H0$ and $H1$, and we consider n observations Y_i, i going from 1 to n, such that:

$$\begin{aligned} H_0 &: Y_n = W_n, \forall n \geq 1 \\ H_1 &: Y_n = m + W_n, \forall n \geq 1 \end{aligned} \tag{2}$$

where W_n $N(0, \sigma^2)$ are i.i.d. and m is the difference in the means under the two hypothesis.

The classical Bayesian updating model between the two hypothesis $H0$ and $H1$ computes the cumulative log-likelihood ratio as follows:

$$L_k = L_{k-1} + l_k \tag{3}$$

where l_k denotes the log-likelihood ratio for observation Y_k, and L_k denotes the cumulative log-likelihood ratio up to observation Y_k. Also,

$$l_k = \log(\frac{f(Y_k|H_1)}{f(Y_k|H_0)}) = \frac{2mY_k - m^2}{2\sigma^2} \tag{4}$$

and

$$L_k = \log(\prod_{i=1}^{k} \frac{f_i(Y_i|H_1)}{f_i(Y_i|H_0)}) = \sum_{i=1}^{k} \log(\frac{f_i(Y_i|H_1)}{f_i(Y_i|H_0)})) = \sum_{i=1}^{k} l_i \tag{5}$$

We propose a model that applies the framework of the belief-adjustment model of Hogarth and Einhorn to the classical Bayesian update model. In other words, we incorporate the anchoring-and-adjustment process to the Bayesian update in order to model the biases. These biases are quantified by the adjustment weight for the new observation. The anchoring bias and the confirmation bias lead the person to treat the new information subjectively, depending on whether it is in accordance to their belief or not. In our new proposed model, the Bayesian update model becomes:

$$L_k = L_{k-1} + p_k l_k \tag{6}$$

where p_k is the adjustment weight that the subject gives to the new observation due to biases. The coefficient p_k is not restricted to a constant, and can be a function of past observations.

In the context of the confirmation bias, the person is sensitive to information that confirm their belief, and tend to neglect disconfirmatory information. So consider the case when the person is biased towards hypothesis $H1$. This can be modeled by making p_k less than one whenever l_k is negative (attenuation of evidence supporting hypothesis $H0$), and by making p_k larger than one whenever l_k is positive (emphasis of evidence supporting hypothesis $H1$). Also, the larger the current cumulative log-likelihood ratio, the more the person is biased towards hypothesis $H1$ and the more p_k is close to zero for negative l_k. So p_k is proportional to L_{k-1}. The Neyman-Pearson test and the SPRT can be applied in the same manner as without biases. However, with the presence of bias, the cumulative log-likelihood ratio is updated while taking into account the effect of biases.

III. ORDERING INFORMATION WITHOUT COGNITIVE BIASES

The main goal of this paper is to study the problem of ordering in real time observations presented to a human subject to optimize the human decision-making. Before solving this problem, we consider the problem of optimally ordering in real time observations in the absence of cognitive biases to maximize performance when a hypothesis test terminates early. Specifically, we reconsider the detection problem mentioned in the previous section. The number of i.i.d. Gaussian observations is N, and we are interested in the problem of finding out which N' observations to show to the subject in order to minimize the probability of error due to the loss of $N - N'$ observations. The solution we propose is to order the log-likelihood ratios of the N observations from the highest absolute value to the smallest, and then pick the N' observations corresponding to the highest N' values of the ordered log-likelihood ratios.

The intuition is that, even though the observations are i.i.d., large absolute value of the log-likelihood ratio has high quality of information.

Moreover, another intuition is the following: the cumulative log-likelihood ratio is $L_N = \sum_{i=1}^{N}(l_i)$. We need to keep the most relevant N' observations with $L_{N'}$=$\sum_{k=1}^{N'}(l_{n_k})$, where k goes from 1 to N' and n_k takes different values from 1 to N. In other words, we denote the error by $|L_N - L_{N'}|$ = $\left|\sum_{i=1}^{N}(l_i) - \sum_{k=1}^{N'}(l_{n_k})\right|$ = $\left|\sum_{k=N'+1}^{N}(l_{n_k})\right|$; we need the error to be as small as possible. The problem of minimizing $\left|\sum_{k=N'+1}^{N}(l_{n_k})\right|$ over all possible combinations of N' out of the N observations has a polynomial time solution

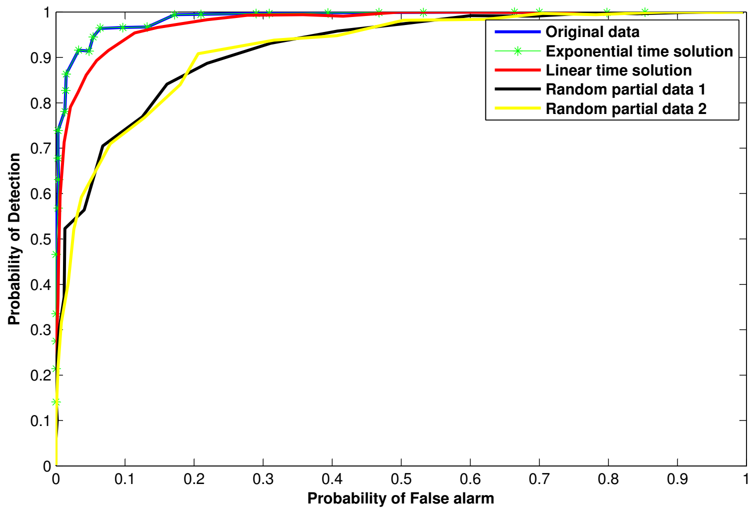

Fig. 1: ROC curves under different orderings of Gaussian i.i.d. observations without bias

$O(N^{N'})$ if we consider N' a constant, and a solution with exponential complexity if we let N' scale with N. Now since $\left|\sum_{k=N'+1}^{N}(l_{n_k})\right| \leq \sum_{k=N'+1}^{N}(|l_{n_k}|)$, then minimizing the RHS would guarantee an "okay" solution, meaning it would guarantee a small error, but not the smallest. Also, since the infimum of the error is zero, a good approach to minimize the error would be to first pick the log-likelihood ratios between $N'+1$ and N as opposite numbers that sum to zero, and when nearly all the log-likelihood ratios left have the same sign, pick the smallest absolute values of the log-likelihood ratios.

In our case, we first note that $l_k = \frac{2mY_k - m^2}{2\sigma^2}$ so l_k follows a Gaussian distribution with variance $Var[l_k|H_0] = Var[l_k|H_1] = m^2/(\sigma^2)$ and mean $E[l_k|H_0] = -m^2/(2\sigma^2)$ and $E[l_k|H_1] = m^2/(2\sigma^2)$. The distribution of l_k is either a Gaussian shifted towards the left of the y-axis or towards the right. Therefore, the opposite values with the highest probability of occurrence lie close to zero, and thus picking the log-likelihoods with smallest absolute values highly guarantees opposite values than sum to zero, and the values left would have a small absolute value as well.

We note that even though this approach is not guaranteed to necessarily give the ordering with the smallest feasible error, its simplicity and good performance make it a good solution. To verify our claim on the proposed ordering of the observations, we consider the total number of i.i.d. Gaussian observations N to be 10, and the number of observations to keep $N' = 4$. We simulate the corresponding ROC curves under 4 cases: the original case where the 10 observations are considered, the second case where we order the observations from their highest absolute value of the log-likelihood ratio to the smallest value and pick the observations corresponding to the largest $N' = 4$ values, the third case corresponding to the exponential time solution, and two other cases corresponding to random orderings of 4 observations from the 10 observations.

The ROC curve of a test plots the probability of detection (Pd) versus the probability of false alarm (Pf) of a test. Only the pairs (Pd,Pf) below the curve are achievable. A test is therefore considered better than another if its ROC curve is higher, meaning the area under its ROC curve is larger.

In Figure 1, the highest ROC curve corresponds to the original case where all observations are considered. As expected, the ROC curve obtained using the exponential time solution is nearly overlapping with the original ROC curve. Also as expected, the next best ROC curve is for the case when we order the observations from the highest absolute value of the log-likelihood ratio to the smallest. It is clear that the ROC curves for random case are lower than the ROC curve with our proposed ordering. Thus, the proposed ordering is a simple linear solution which guarantees low probability of error.

IV. ORDERING INFORMATION WITH COGNITIVE BIASES

In the context of cognitive biases, the Bayesian update is no more the classical one, as shown in II. The cumulative log-likelihood ratio up to the N'th observation $L^b_{N'}$ becomes: $L^b_{N'}=\sum_{k=1}^{N'}(p_{n_k} l_{n_k})$. The error becomes: $\left|\sum_{k=N'+1}^{N}(l_{n_k}) + \sum_{k=1}^{N'}(l_{n_k}(1-p_{n_k}))\right|$. We note that minimizing this error has a polynomial time solution $O(N^{N'})$ if we consider N' a constant, and a solution with exponential complexity if we let N' scale with N. It is a harder problem with the bias model than without considering the bias since in the bias model, different orderings of a specific combination of N' observations lead to different cumulative log-likelihood ratios given the N' observations. This is because for each new observation, the adjustment weight depends on the cumulative log-likelihood ratio at the time of the acquisition of the new information as shown in Section II. Consider the setting where the information is attenuated if it is not supporting the current belief of the subject, i.e. the current cumulative log-likelihood ratio L_{k-1}, and emphasized when it is supporting the current belief. Also, as pointed out in Section II, the adjustment weight p_k is a function of L_{k-1}. We consider an adjustment weight p less than 1 as attenuating an observation, and an adjustment weight p larger than 1 as emphasizing an observation. The

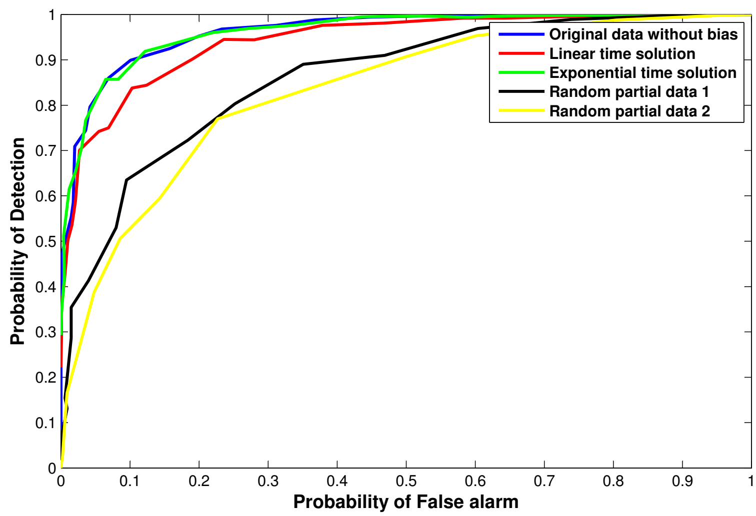

Fig. 2: ROC curves under different orderings of Gaussian i.i.d. observations with bias

solution we propose is to keep observations with the highest absolute values of the log-likelihood ratio as shown in Section III, and we show that this method guarantees low probability of error while reducing the exponential complexity of minimizing $\left|\sum_{k=N'+1}^{N}(l_{n_k}) + \sum_{k=1}^{N'}(l_{n_k}(1-p_{n_k}))\right|$ over all the possible permutations of N' out of N observations. Another advantage for the method proposed is that it doesn't require knowing the adjustment weight neither the initial log-likelihood ratio that a human has before any observation. We simulated the ROC for the same problem as above (with $N' = 4$, $N = 10$). An adjustment weight p_k less than one discriminates disconfirming observations and an adjustment weight p_{n_k} higher than one keeps the confirming observations. We note that p_{n_k} is proportional to the value of the prior cumulative log-likelihood ratio L_{k-1} before the acquisition of the new observation as suggested in section II.

In Figure 2, the highest ROC curve corresponds also to the original case without bias where all the N observations are considered. The ROC curve of the exponential time solution for the problem of minimizing the error: $\left|\sum_{k=N'+1}^{N}(l_{n_k}) + \sum_{k=1}^{N'}(l_{n_k}(1-p_{n_k}))\right|$ is very close to the original ROC curve. As expected, the next lower ROC curve (in red) corresponds to the case where ordering of the observations is done from the highest absolute value of the log-likelihood ratio to the smallest (as in section III). The lowest ROC curves are for random orderings. The bias introduces additional errors due to over or under estimating the log-likelihood ratio of the observations. Ordering observations from highest to lowest absolute value of log-likelihood ratio gives good results and low probability of errors since it keeps the most relevant observations with the highest absolute value of log-likelihood ratio. It has linear complexity, and doesn't require knowing neither the initial log-likelihood ratio nor the parameters of the bias.

V. CONCLUSION

In this paper, we modified the Bayesian updating model to account for cognitive biases. We then proposed ordering i.i.d. Gaussian observations from highest to lowest absolute value of log-likelihood ratio to guarantee low probability of error. In the context of cognitive biases, we proposed ordering observations while also keeping observations with the highest absolute value of the log-likelihood ratio. This method has low complexity, and doesn't require knowing neither the parameters of the bias nor the initial log-likelihood ratio when accounting for the biases. We showed with simulations that both orderings actually guarantee lower probability of errors and larger area under the ROC curves than random orderings.

REFERENCES

[1] N. Akl and A. Tewfik. Optimal information ordering in sequential detection problems with cognitive biases. In *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, pages 1876–1880, May 2014.

[2] E. Bamber, R. J. Ramsay, and R. M. Tubbs. An examination of the descriptive validity of the belief-adjustment model and alternative attitudes to evidence in auditing. *Accounting, Organizations and Society*, 22(34):249 – 268, 1997.

[3] M. H. Birnbaum and B. A. Mellers. Bayesian inference: Combining base rates with opinions of sources who vary in credibility. *Journal of Personality and Social Psychology*, 45(4):792, 1983.

[4] Y.-L. Chien, C. J. Huang, and D. Shaw. A general model of starting point bias in double-bounded dichotomous contingent valuation surveys. *Journal of Environmental Economics and Management*, 50(2):362–377, 2005.

[5] G. Friesen and P. A. Weller. Quantifying cognitive biases in analyst earnings forecasts. *Journal of Financial Markets*, 9(4):333 – 365, 2006.

[6] J. A. Herriges and J. F. Shogren. Starting point bias in dichotomous choice valuation with follow-up questioning. *Journal of Environmental Economics and Management*, 30(1):112 – 131, 1996.

[7] R. M. Hogarth and H. J. Einhorn. Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24(1):1 – 55, 1992.

[8] R. Iyer and A. Tewfik. Optimal ordering of observations for fast sequential detection. In *Signal Processing Conference (EUSIPCO), 2012 Proceedings of the 20th European*, pages 126–130, 2012.

[9] F. Lieder, T. L. Griffiths, and N. D. Goodman. Burn-in, bias, and the rationality of anchoring. In *NIPS*, pages 2699–2707, 2012.
[10] Y. Liu and S. Blostein. Optimality of the sequential probability ratio test for nonstationary observations. *Information Theory, IEEE Transactions on*, 38(1):177–182, 1992.
[11] A. Wald. Sequential tests of statistical hypotheses. *The Annals of Mathematical Statistics*, 16(2):117–186, 06 1945.