
Distributed Decision Making by Categorically-Thinking Agents

Joong Bum Rhim

Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139
jbrhim@mit.edu

Lav R. Varshney

IBM Thomas J. Watson Research Center
Hawthorne, NY 10532
lrvarshn@us.ibm.com

Vivek K Goyal

Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA 02139
vgoyal@mit.edu

Abstract

This paper considers group decision making by imperfect agents that only know quantized prior probabilities for use in Bayesian likelihood ratio tests. Global decisions are made by information fusion of local decisions, but information sharing among agents before local decision making is forbidden. The quantization scheme of the agents is investigated so as to achieve the minimum mean Bayes risk; optimal quantizers are designed by a novel extension to the Lloyd-Max algorithm. Diversity in the individual agents' quantizers leads to optimal performance.

1 Introduction

Consider a binary decision problem. The Bayes rational decision making strategy is to perform the likelihood ratio test (LRT). Decision makers first compute the likelihood ratio of states of an object based on an observation. Then they make a decision by comparing the ratio to a decision threshold determined by the prior probability of the state and their costs. Not only do LRTs minimize Bayes risk, but also psychology experiments suggest that human decision makers employ them [1, 2].

Optimal LRTs require precise knowledge of the prior probabilities of object states. Much previous research considers the prior probability to be a constant known to decision makers. However, decision makers may face a great variety of objects. For example, soccer referees handle more than twenty-two players in one game and salespeople at stores observe hundreds of customers in one day. This is problematic because players have different prior probabilities of committing fouls and customers have different prior probabilities of making purchases.

Decision makers should use different thresholds uniquely optimized to different objects of decision making, such as players or customers. Computing thresholds and then remembering them is an information processing burden, especially for human decision makers that often resort to categorical and coarse thinking [3, 4]. Human decision makers can afford around seven categories without

getting confused [5]. Thus, we model decision makers as grouping similar objects together and treating them identically by applying a single decision threshold. By classifying all objects into a small number of categories, decision makers can handle infinitely many objects; however decision makers consequently have limited threshold precision, a type of bounded rationality.

In the context of LRTs, categorization of objects is equivalent to quantization of their prior probabilities. With this idea, we move from considering a single object with a constant prior probability to considering an ensemble of objects with performance averaged over the distribution of prior probabilities.

Consider a decision-making group of N agents that chooses between two hypotheses h_0 and h_1 . Agents make local hard decisions without knowing other agents' decisions. Local decisions are combined by a fusion center to produce a global decision. The fusion center has a fixed symmetric fusion rule of the L -out-of- N form whereby the global decision is h_1 when L or more agents choose h_1 . The symmetric fusion rule implies that all agents have an equal voice. Due to information-processing constraints, agents must quantize prior probabilities to one of K values. Our interest here is to design quantizers that lead to the smallest Bayes risk on average.

The study of quantization of prior probabilities originates from [6], which focuses on the minimum mean Bayes risk error (MBRE) quantizer of a single agent. Maximum Bayes risk error is considered in [7, 8]. Recent results and economic implications are reviewed in [9].

We have previously considered a distributed hypothesis testing problem with similar imperfect agents, but where each agent is assumed to know other agents' quantized prior probabilities, whether they have a common interest [8, 10] or whether they have conflicts of interest [11]. The assumption in these prior papers enables agents to optimize decision rules so as to minimize Bayes risk within either the collaboration or the conflict system.

Information about other agents' quantizers should not be taken for granted; it requires a coordination mechanism built on communication channels. Such communication may not be possible in human group decision-making scenarios due to geographic separation, desire to remain clandestine, or if N is too large. In engineering applications, memory or power constraints may prevent detectors from exchanging any information with neighboring detectors. In these scenarios, each agent has to make decisions based on its information—its quantized prior probability and observed signal—only. In this paper, agents do not know how other agents quantize prior probabilities.

Lack of knowledge about others makes it impossible for agents to collaborate by sharing a common goal. Hence, their quantizers need to be cleverly designed so that local decision making becomes harmonious with respect to the global mean Bayes risk (MBR), the distortion measure for quantization. A modified Lloyd-Max algorithm can design MBR-optimal quantizers. It is demonstrated that diversity among agents in quantization of prior probabilities can be helpful to improve the quality of group decision making.

The group decision-making model we consider is described in Section 2. In Section 3, we analyze the mean Bayes risk in terms of endpoints and representation points of quantizers. Then we propose an algorithm to design optimal quantizers. An example of optimal quantizers obtained from our algorithm is presented in Section 4. Section 5 concludes the paper.

2 Distributed Decision-Making Model with Imperfect Agents

We consider a team of N agents and an object in one of two binary states $H \in \{h_0, h_1\}$. The prior probability of the object being in state h_0 , $p_0 = \mathbb{P}\{H = h_0\}$, is a realization of a random variable P_0 drawn from its distribution f_{P_0} . Since the prior probability of being in state h_1 is determined by p_0 through $p_1 = 1 - p_0$, by the term *prior probability* we simply mean p_0 . The prior probability is important for good decision making but Agent i only knows its quantized value of the prior probability, $q_i(p_0)$.

Agent i makes a noisy state measurement Y_i with likelihood functions $f_{Y_i|H}(y_i|h_0)$ and $f_{Y_i|H}(y_i|h_1)$. Agent i then makes a hard decision \hat{H}_i whether the object is in state h_0 or in h_1 based on the quantized prior probability $q_i(p_0)$ and the observation Y_i . Its decision is transferred to a fusion center, which makes a global decision \hat{H} as h_1 if it receives h_1 from L or more agents




Figure 1: A set of quantizers q_1, \dots, q_N categorizes cells $\mathcal{C}_1, \dots, \mathcal{C}_{N_{\text{Cell}}}$.

and as h_0 otherwise. The agents incur cost c_{10} for a false alarm ($\hat{H} = h_1$ when $H = h_0$) and c_{01} for a missed detection ($\hat{H} = h_0$ when $H = h_1$); costs c_{10} and c_{01} are common for all agents. For simplicity, correct global decisions incur zero cost.

Agent i optimizes its decision rule as if it is the only decision maker because it does not have any information about other agents:

$$\frac{f_{Y_i|H}(y_i | h_1)}{f_{Y_i|H}(y_i | h_0)} \underset{\hat{H}_i(y)=h_0}{\overset{\hat{H}_i(y)=h_1}{\geq}} \frac{c_{10}q_i(p_0)}{c_{01}(1 - q_i(p_0))}. \quad (1)$$

This decision rule yields an error with probability $P_{e,i}^I = \mathbb{P}\{\hat{H}_i = h_1 | H = h_0\}$ when $H = h_0$ and with probability $P_{e,i}^{II} = \mathbb{P}\{\hat{H}_i = h_0 | H = h_1\}$ when $H = h_1$.

The agents cannot collaborate to design a decision rule, but they still fuse their decisions to make a global decision. By using L -out-of- N fusion rules, the global decision is wrong if L or more agents send h_1 when $H = h_0$ or if $N - L + 1$ or more agents send h_0 when $H = h_1$. These error probabilities, P_E^I and P_E^{II} , are used in computing the Bayes risk

$$R = c_{10}p_0P_E^I + c_{01}(1 - p_0)P_E^{II}.$$

3 Optimal Quantization of Prior Probabilities

Agent i has quantizer q_i for prior probability p_0 , which has K cells $[0, b_1^{(i)}), [b_1^{(i)}, b_2^{(i)}), \dots, [b_{K-1}^{(i)}, 1]$ with corresponding representation points $a_1^{(i)}, a_2^{(i)}, \dots, a_K^{(i)}$, where $a_k^{(i)} = q_i(p_0)$ for all $p_0 \in [b_{k-1}^{(i)}, b_k^{(i)})$. We define a set of endpoints $\{0, b_1, b_2, \dots, b_{N_{\text{Cell}}-1}, 1\}$, $0 < b_1 < b_2 < \dots < b_{N_{\text{Cell}}-1} < 1$, as the union of endpoints of all quantizers q_1, \dots, q_N and define cells \mathcal{C}_k as the intervals $[b_{k-1}, b_k)$, where N_{Cell} is the number of the cells \mathcal{C}_k . The maximum number of N_{Cell} is $N(K - 1) + 1$. For cell \mathcal{C}_k , we define a vector of representation points $\mathbf{a}_k = (q_1(p_0), q_2(p_0), \dots, q_N(p_0))$, where $p_0 \in \mathcal{C}_k$, see Fig. 1. The necessary conditions of representation points and endpoints for local optimality of the quantizers are now derived.

3.1 Representation points

Quantization performance is measured by Bayes risk averaged over P_0 :

$$\mathbb{E}[R] = \int_0^1 (c_{10}p_0P_E^I(q_1(p_0), \dots, q_N(p_0)) + c_{01}(1 - p_0)P_E^{II}(q_1(p_0), \dots, q_N(p_0))) f_{P_0}(p_0) dp_0.$$

Within cell \mathcal{C}_k , since $(q_1(p_0), \dots, q_N(p_0)) = \mathbf{a}_k$ is constant, the mean Bayes risk (MBR) is

$$\mathbb{E}[R]_k = \int_{\mathcal{C}_k} (c_{10}p_0P_E^I(\mathbf{a}_k) + c_{01}(1 - p_0)P_E^{II}(\mathbf{a}_k)) f_{P_0}(p_0) dp_0 = c_{10}\pi_k^I P_E^I(\mathbf{a}_k) + c_{01}\pi_k^{II} P_E^{II}(\mathbf{a}_k),$$

where $\pi_k^I = \int_{\mathcal{C}_k} p_0 f_{P_0}(p_0) dp_0$ and $\pi_k^{II} = \int_{\mathcal{C}_k} (1 - p_0) f_{P_0}(p_0) dp_0$ are constants with respect to \mathbf{a}_k .

Let us fix all representation points except that of q_z, a_{kz} , in \mathcal{C}_k . The mean Bayes risk in \mathcal{C}_k can be written as $\mathbb{E}[R]_k = \alpha_1 P_{e,z}^I + \alpha_2 P_{e,z}^{II} + \alpha_3$, where α_1, α_2 , and α_3 are positive constants. Since $P_{e,z}^{II}$ is strictly convex in the $P_{e,z}^I$ and vice versa [12], $\mathbb{E}[R]_k$ is strictly convex in $P_{e,z}^I(a_{kz})$ and $P_{e,z}^{II}(a_{kz})$.

The convexity is preserved in the entire MBR $\mathbb{E}[R] = \mathbb{E}[R]_1 + \dots + \mathbb{E}[R]_{N_{\text{Cell}}}$ because the MBR in each cell is strictly convex in $P_{e,z}^I(a_{kz})$ and $P_{e,z}^{II}(a_{kz})$ or constant. Therefore, the value of a_{kz} that minimizes the MBR exists uniquely for any $1 \leq k \leq N_{\text{Cell}}$ and $1 \leq z \leq N$. The value of a_{kz} should be the minimum point.

3.2 Endpoints

Let us fix all representation points and endpoints except an endpoint b_j . The endpoint b_j only affects two adjacent cells \mathcal{C}_j and \mathcal{C}_{j+1} , whose boundary is b_j .

$$\begin{aligned} \mathbb{E}[R]_j + \mathbb{E}[R]_{j+1} &= \int_{b_{j-1}}^{b_j} (c_{10} p_0 P_E^I(\mathbf{a}_j) + c_{01} (1 - p_0) P_E^{II}(\mathbf{a}_j)) f_{P_0}(p_0) dp_0 \\ &\quad + \int_{b_j}^{b_{j+1}} (c_{10} p_0 P_E^I(\mathbf{a}_{j+1}) + c_{01} (1 - p_0) P_E^{II}(\mathbf{a}_{j+1})) f_{P_0}(p_0) dp_0 \end{aligned}$$

Taking the derivative of the MBR, we have

$$\begin{aligned} \frac{d}{db_j}(\mathbb{E}[R]) &= \frac{d}{db_j}(\mathbb{E}[R]_j + \mathbb{E}[R]_{j+1}) \\ &= (c_{10} b_j (P_E^I(\mathbf{a}_j) - P_E^I(\mathbf{a}_{j+1})) - c_{01} (1 - b_j) (P_E^{II}(\mathbf{a}_{j+1}) - P_E^{II}(\mathbf{a}_j))) f_{P_0}(b_j) \end{aligned} \quad (2)$$

If we compare each entry of two vectors \mathbf{a}_j and \mathbf{a}_{j+1} , at least one entry has different values. For any entry that has a different value, \mathbf{a}_{j+1} has a greater value than \mathbf{a}_j does because the former represents larger P_0 . A bigger representation point leads to a smaller local false alarm probability. Thus, $P_E^I(\mathbf{a}_{j+1}) < P_E^I(\mathbf{a}_j)$. On the contrary, $P_E^{II}(\mathbf{a}_{j+1}) > P_E^{II}(\mathbf{a}_j)$. Let $\beta_1 = P_E^I(\mathbf{a}_j) - P_E^I(\mathbf{a}_{j+1}) > 0$ and $\beta_2 = P_E^{II}(\mathbf{a}_{j+1}) - P_E^{II}(\mathbf{a}_j) > 0$.

$$\frac{d}{db_j}(\mathbb{E}[R]) = ((c_{10} \beta_1 + c_{01} \beta_2) b_j - c_{01} \beta_2) f_{P_0}(b_j). \quad (3)$$

This first derivative is zero at only one or no point if $f_{P_0}(p_0) > 0, \forall p_0 \in [0, 1]$. This means that $\mathbb{E}[R]$ has only one or zero extreme point for $b_j \in (b_{j-1}, b_{j+1})$: if it has one extreme point, then it is the minimum point. Otherwise, either b_{j-1} or b_{j+1} is the minimum point. The value of b_j should be the minimum point.

3.3 Algorithm

The iterative Lloyd-Max algorithm is applied to find an optimal quantizer in a single-agent decision-making model [6]. In this problem, however, the algorithm needs to be modified so as to optimize N different quantizers together. The key to the Lloyd-Max algorithm is alternating iterations of finding optimal endpoints while fixing representation points and finding optimal representation points while fixing endpoints.

In our group decision-making model, optimization steps are complicated because of dependency among variables; a change of one representation point also changes optimal values of other representation points. Hence, representation points are repeatedly adjusted until every representation point is optimal for the other representation points and given endpoints. Likewise for optimization of endpoints.

We use the following alternating nested-iteration optimization algorithm:

1. Assign initial values to endpoints and representation points.
2. (E-Step) Optimize endpoints with representation points fixed.
 - (a) From the first endpoint variable $b_1^{(1)}$ to the last one $b_{K-1}^{(N)}$, successively optimize each variable.




Figure 2: Bayes risk for uniformly distributed P_0 and $K = 2$. (a) An optimal set of quantizers (cell endpoints as '+'s and representation points as 'o's) and the resulting Bayes risk. (b) The performance loss in terms of Bayes risk due to the quantization of prior probabilities.

- (b) Repeat step (a) until all endpoints become stable, i.e., a new iteration does not change any endpoints.
3. (R-Step) Optimize representation points with endpoints fixed.
 - (a) From the first representation point variable $a_1^{(1)}$ to the last one $a_{K-1}^{(N)}$, successively optimize each variable.
 - (b) Repeat step (a) until all representation points become stable.
4. Iterate E-Step and R-Step until all endpoints and representation points become stable.

4 Example

As an example, let us consider the following measurement model for $N = 3$ agents:

$$Y_i = s_m + W_i, \quad i = 1, \dots, N, \quad m \in \{0, 1\}, \quad (4)$$

where $s_0 = 0$, $s_1 = 1$, and W_i is a zero-mean Gaussian random variable with variance $\sigma^2 = 1$. The Bayes costs are $c_{10} = c_{01} = 1$. The local decisions are fused by MAJORITY rule (2-out-of-3 rule).

Fig. 2a shows Bayes risk when the agents can distinguish 2 categories, i.e., they use 2-level quantizers. The Bayes risk (solid piecewise line) is compared to the Bayes risk when the agents can distinguish any prior probability exactly and collaborate with others (dashed curve) like in [10], which is the best performance that the agents can achieve. The excess Bayes risk, the difference between the Bayes risks with and without quantization, is depicted in Fig. 2b. It shows the performance loss due to quantized prior probabilities compared to the best performance. For comparison, Fig. 2b also shows the performance loss when the agents are forced to use identical quantizers (gray dashed curve) and the performance loss when the agents use diverse quantizers and can collaborate by sharing their quantized values (dash-dot curve). The latter is the best performance that the agents can achieve with quantized prior probabilities [8, 10].

The Bayes risks when the agents use 4-level quantizers are depicted in Fig. 3.

5 Conclusion

We have discussed decision making by multiple agents that have imperfect perception ability. There are two factors that worsen the quality of global decisions. First, they perform local testing based on quantized prior probabilities. Second, they do not know how other agents quantize prior probabilities. We have determined the effect of these factors on Bayes risk in decision making.

To minimize the negative influence from these factors, we have defined mean Bayes risk as the optimization criterion for prior-probability quantizers. The Lloyd-Max algorithm is modified to an algorithm with double-iteration structure to design optimal quantizers. Using the algorithm, we




Figure 3: Bayes risk for uniformly distributed P_0 and $K = 4$. (a) An optimal set of quantizers (cell endpoints as '+'s and representation points as 'o's) and the resulting Bayes risk. (b) The performance loss in terms of Bayes risk due to the quantization of prior probabilities.

have provided an example of additive white Gaussian noise model. The result shows that the MBR when the agents use diverse quantizers is lower than the MBR when they use identical quantizers. It is reasonable because $N_{\text{Cell}} (= N(K - 1) + 1)$ when they use diverse quantizers is greater than $N_{\text{Cell}} (= K)$ when they use identical quantizers. Therefore, we can conclude that the diversity among agents is still helpful even though they cannot fully utilize the diversity because of lack of knowledge about other agents.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grants 0729069 and 1101147 and the Korea Foundation for Advanced Studies (KFAS) Fellowship.

References

- [1] J. A. Swets, T. Wilson P, Jr., and T. G. Birdsall, "Decision process in perception," *Psychological Review*, vol. 68, no. 5, pp. 301–340, Sep. 1961.
- [2] M. Glanzer, A. Hilford, and L. T. Maloney, "Likelihood ratio decisions in memory: Three implied regularities," *Psychonomic Bulletin & Review*, vol. 16, no. 3, pp. 431–455, Jun. 2009.
- [3] G. A. Miller, "Human memory and the storage of information," *IRE Trans. Inf. Theory*, vol. 2, no. 3, pp. 129–137, Sep. 1956.
- [4] R. Fryer and M. O. Jackson, "A categorical model of cognition and biased decision making," *The B.E. Journal of Theoretical Economics*, vol. 8, no. 1, Feb. 2008.
- [5] G. A. Miller, "The magical number seven, plus or minus two: Some limits on our capacity for processing information," *Psychological Review*, vol. 63, no. 2, pp. 81–97, Mar. 1956.
- [6] K. R. Varshney and L. R. Varshney, "Quantization of prior probabilities for hypothesis testing," *IEEE Trans. Signal Process.*, vol. 56, no. 10, pp. 4553–4562, Oct. 2008.
- [7] —, "Multilevel minimax hypothesis testing," in *Proc. IEEE/SP Workshop Stat. Signal Process.*, Nice, France, Jun. 2011, pp. 109–112.
- [8] J. B. Rhim, L. R. Varshney, and V. K. Goyal, "Quantization of prior probabilities for collaborative distributed hypothesis testing," *arXiv:1109.2567*, 2011.
- [9] L. R. Varshney, J. B. Rhim, K. R. Varshney, and V. K. Goyal, "Categorical decision making by people, committees, and crowds," in *Proc. Information Theory and Applications Workshop*, La Jolla, CA, Feb. 2011.
- [10] J. B. Rhim, L. R. Varshney, and V. K. Goyal, "Collaboration in distributed hypothesis testing with quantized prior probabilities," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 2011, pp. 303–312.
- [11] —, "Conflict in distributed hypothesis testing with quantized prior probabilities," in *Proc. IEEE Data Compression Conf.*, Snowbird, UT, Mar. 2011, pp. 313–322.
- [12] H. L. van Trees, *Detection, Estimation, and Modulation Theory, Part I*. New York: John Wiley & Sons, 1968.