Blind Collaborative 20 Questions for Target Localization

Theodoros Tsiligkaridis
EECS Dept., Univ. of Michigan
Ann Arbor, MI, USA
ttsili@umich.edu

Brian M. Sadler
US Army Research Laboratory
Adelphi, MD, USA
brian.m.sadler6.civ@mail.mil

Alfred O. Hero III
EECS Dept. & Dept. Statistics, Univ. of Michigan
Ann Arbor, MI, USA
hero@umich.edu

***Abstract*—We consider the problem of collaborative target localization by several observers, called players, where the reliability of each player is unknown. As in our previous work [1] we formulate this problem as a 20 questions game with noise for collaborative players under a minimum entropy criterion. We extend the setting of [1] to the case where the players' error channels have unknown crossover probabilities. First, we use dynamic programming to characterize the structure of the optimal policy for constructing the sequence of questions. This generalizes the multiplayer policies derived in [1] for the known error channel setting. Second, we prove a separation theorem showing that a sequential bisection scheme achieves the same performance as the optimal joint queries. This generalizes the separation theorem recently derived for the known error channel case in [1]. Third, we derive bounds for the maximum entropy loss per iteration. Finally, we show that even for the one-dimensional case, the optimal query policy for the unknown error channel is not equivalent to a probabilistic bisection policy. This framework provides a methodology for simultaneous sequential estimation of target location and learning the error channels associated with the players.**

I. Introduction

Consider the problem of estimation of an unknown target location by playing 20 questions game with a group of sensors. In this game, sensors are repeatedly queried about target location. The objective is to optimize the sequence of queries when the accuracy of responses of the noisy oracles is unknown, i.e., unknown error channels. This is especially relevant to the case of human-in-the-loop systems where the probability of correct response of the human may be difficult to predict and quantify.

Sequential estimation of target position was studied in [2], for the single player setting, in the context of a noisy 20 questions game, where the objective was to minimize the expected entropy after N questions. In the collaborative case [1], a controller sequentially poses a set of questions about target location to multiple sensors and fuses the sensors' noisy responses to formulate the next questions.

This paper focuses on the unknown error channel case. Our approach is based on jointly estimating the target and the error channels associated with the players. Using dynamic programming, we characterize the optimal policy and provide bounds on the maximum expected entropy loss per iteration. We also derive a separation theorem that shows that a sequential bisection scheme achieves the same expected entropy loss as the jointly optimal scheme.

This research was partially supported by MURI grant W911NF-11-1-0391.

A. Previous Work

The paper by Jedynak et al. [2] formulates the single player 20 questions problem as a controller querying a noisy oracle about whether or not a target X^* lies in a set $A_n \subset \mathbb{R}^d$. Starting with a prior distribution on the target's location $p_0(\cdot)$, the objective is to minimize the expected entropy of the posterior distribution:

$$\inf_{\pi} \mathbb{E}^{\pi}\left[H(p_N)\right] \quad (1)$$

where $\pi = (\pi_0, \pi_1, \dots)$ denotes the controller's query policy and the entropy is the standard differential entropy [3] $H(p) = -\int_{\mathcal{X}} p(x)\log p(x)dx$. The posterior mean or median p_N is used to estimate the target location after N questions. The densities f_0 and f_1 correspond to the noisy channel :

$$\mathbb{P}(Y_{n+1} = y | Z_n = z) = f_0(y) I(z = 0) + f_1(y) I(z = 1)$$

where $Z_n = I(X^* \in A_n) \in \{0, 1\}$ is the channel input. The noisy channel models the conditional probability of the response to each question being correct. For the special case of a binary symmetric channel (BSC), $u^* = 1/2$ and the probabilistic bisection policy [2], [4] becomes an optimal policy. Thm. 2 in [2] shows the bisection policy is optimal under the minimum entropy criterion-i.e., $P_n(A_n) := \int_{A_n} p_n(x)dx = u^* \in \arg\max_{u\in[0,1]} \phi(u)$, where $\phi(u) = H(f_1 u + (1-u) f_0) - uH(f_1) - (1-u)H(f_0)$ is nonnegative.

Recently, Tsiligkaridis et al. [1] derived optimality conditions for query strategies in the collaborative multiplayer case. It was shown that even when the collaborative players act independently, jointly optimal policies require overlapping non-identical queries. A sequential bisection policy for which each player responds to a single question was introduced and it was proven that the expected entropy reduction for the jointly optimal scheme is the same as that of the sequential bisection scheme. Thus, while the jointly optimal scheme might be hard to implement as the number of players and dimensions increase, the sequential bisection scheme simplifies the controller design with no performance degradation.

The function $I(A)$ is the indicator function throughout the paper-i.e., $I(A) = 1$ if A is true and zero otherwise.

II. Notation & Assumptions

In this paper, we adopt the setup of [1]. Assume that there is a target with unknown state $X^* \in \mathcal{X} \subset \mathbb{R}^d$. There are M collaborating players that can be asked questions at each time instant. The objective of the players is to come up with the correct answer to a kind of 20 questions game.

In the joint estimation setup, we assume that the controller design queries for M sensors and, after querying, the responses are fused and the next set of questions is formulated (see Fig. 1). Let the mth player's query at time n be "does X^* lie in the region $A_n^{(m)} \subset \mathbb{R}^d$?". We denote this query as $Z_n^{(m)} = I(X^* \in A_n^{(m)}) \in \{0,1\}$ to which the player yields provides a noisy response $Y_{n+1}^{(m)} \in \{0,1\}$. The query region(s) chosen at time n depend on the information available at time n. More formally, $\{Z_n^{(m)} = I(X^* \in A_n^{(m)})\}$ is a predictable stochastic process with respect to the filtration generated by $\{A_n^{(m)}\}$ and $\{Y_{n+1}^{(m)}(A_n^{(m)})\}$, i.e., $\{A_n^{(m)}\}_m \in \mathcal{F}_n := \sigma(\{A_k^{(m)}; Y_{k+1}^{(m)}(A_k^{(m)})\}_m : 0 \le k \le n-1)$ for $n \in \mathbb{N}$. At each iteration a current best target estimate X_n of X^* is produced (which is an $\mathcal{F}_n$-measurable random variable).

The sequential strategy consists of sequentially asking players queries and using the intermediate responses to refine the posterior (see Fig. 2). For each time epoch, indexed by n and called a cycle, the controller formulates and asks the M players questions $A_{n_t} = A_{n,t}$, $t = 0, \ldots, M-1$. Let the mth player's query at time $n_t = (n,t) = n_{m-1}$ be denoted by $Z_{n_t} = I(X^* \in A_{n_t}) \in \{0,1\}$ and its associated noisy response $Y_{n_{t+1}} \in \{0,1\}$. The query region A_{n_t} chosen at time n_t depends on the information available at that time. More formally, define the multi-index (n,t) where $n = 0, 1, \ldots$ indexes over cycles and $t = 0, \ldots, M-1$ indexes within cycles. Define the nested sequence of sigma-algebras $\mathcal{G}_{n,t}$, $\mathcal{G}_{n,t} \subset \mathcal{G}_{n+i,t+j}$, for all $i \ge 0$ and $j \in \{0, \ldots, M-1-t\}$, generated by the sequence of queries and the players' responses. The filtration $\mathcal{G}_{n,t}$ carries all the information accumulated by the controller from time $(0,0)$ to time (n,t). The queries $\{A_{n,t}\}$ formulated by the controller are measurable with respect to this filtration.

Define the random vector $\epsilon = (\epsilon_1, \ldots, \epsilon_M) \in [0,1/2)^M$, the joint posterior distributions $\mathbb{P}(X^* = x, \epsilon^* = \epsilon | \mathcal{F}_n) = p_n(x,\epsilon)$ and $\mathbb{P}(X^* = x, \epsilon^* = \epsilon | \mathcal{G}_{n,t}) = p_{n_t}(x,\epsilon)$.

For sets $A \subset \mathbb{R}^d$, define $A^1 = A$ and $A^0 = A^c$. Define the M-tuples $\mathbf{Y}_{n+1} = (Y_{n+1}^{(1)}, \ldots, Y_{n+1}^{(M)})$ and $\mathbf{A}_n = \{A_n^{(1)}, \ldots, A_n^{(M)}\}$. Given the responses $\mathbf{Y}_{n+1}$, the posterior update becomes [1]:

$$p_{n+1}(x,\epsilon) \propto \mathbb{P}(\mathbf{Y}_{n+1} = \mathbf{y}_{n+1} | \mathbf{A}_n, X^* = x, \epsilon^* = \epsilon, \mathcal{F}_n) p_n(x,\epsilon) \tag{2}$$

Assuming that all sensors are queried in sequence starting from $m = 1$ and ending at $m = M$, the posterior updates (after querying the $(t+1)$th player) become:

$$p_{n_{t+1}}(x,\epsilon)$$
$$\propto \mathbb{P}(Y_{n_{t+1}} = y_{n_{t+1}} | A_{n_t}, X^* = x, \epsilon^*_{t+1} = \epsilon_{t+1}, \mathcal{G}_{n_t}) p_{n_t}(x,\epsilon)$$

$$\mathbb{P}(Y_{n_{t+1}} | A_{n_t}, X^*, \epsilon^*_{t+1}, \mathcal{G}_{n_t}) = \begin{cases} f_1^{(t+1)}(Y_{n_{t+1}} | \epsilon^*_{t+1}), & X^* \in A_{n_t} \\ f_0^{(t+1)}(Y_{n_{t+1}} | \epsilon^*_{t+1}), & X^* \notin A_{n_t} \end{cases}$$

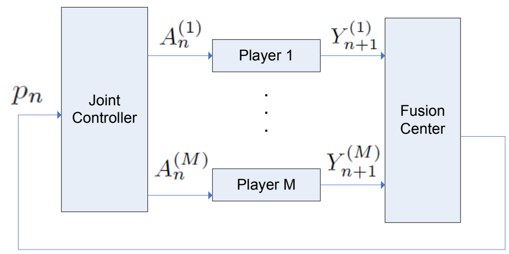

Fig. 1. Joint scheme for M collaborative players responding to binary valued queries about the location X^* of an unknown target.

We make the following assumptions throughout the paper.

Assumption 1. *(Conditional Independence) Assume that the players' responses are conditionally independent:*

$$\mathbb{P}(\mathbf{Y}_{n+1} | \mathbf{A}_n, X^*, \epsilon^*, \mathcal{F}_n) = \prod_{m=1}^{M} \mathbb{P}(Y_{n+1}^{(m)} | A_n^{(m)}, X^*, \epsilon^*) \tag{3}$$

where

$$\mathbb{P}(Y_{n+1}^{(m)} | A_n^{(m)}, X^*, \epsilon^*) = \begin{cases} f_1^{(m)}(Y_{n+1}^{(m)} | \epsilon^*_m, A_n^{(m)}), & X^* \in A_n^{(m)} \\ f_0^{(m)}(Y_{n+1}^{(m)} | \epsilon^*_m, A_n^{(m)}), & X^* \notin A_n^{(m)} \end{cases} \tag{4}$$

Assumption 2. *(Memoryless Binary Symmetric Channels) Players' response channels are independent (memoryless) binary symmetric channels (BSC) [3] with crossover probabilities $\epsilon_m \in [0,1/2)$:*

$$f_j^{(m)}(y^{(m)} | \epsilon_m, A_n^{(m)}) = f_j^{(m)}(y^{(m)}) = \begin{cases} 1 - \epsilon_m, & y^{(m)} = j \\ \epsilon_m, & y^{(m)} \neq j \end{cases}$$

where $m = 1, \ldots, M, j = 0, 1$.

III. Noisy 20 Questions with Collaborative Players: Unknown Error Channels

We consider the setting where the error probabilities of the M players are unknown. In this case, the Bayes posterior update (2) is not well-defined, so the probabilistic bisection algorithm cannot be directly used. In the generic setup of unknown $\epsilon^*_m \in [0,1/2)$ with no a priori information, a joint scheme is to estimate the target X^* and the error probabilities $\epsilon^* = (\epsilon^*_1, \ldots, \epsilon^*_M)$. The joint posterior distribution of (X^*, ϵ^*) is considered here because the error probabilities ϵ_m are coupled with the target x through the Bayesian update (e.g. see (4) and (2)).

A. Joint Query Design

We consider the minimum entropy criterion (1). Since the error probabilities of sensors are unknown, the joint policy derived in Thm. 1 in [1] is no longer applicable or valid.

Define the density parameterized by $\epsilon = (\epsilon_1, \ldots, \epsilon_M) \in [0,1/2)^M$ and $\mathbf{i} = (i_1, \ldots, i_M) \in \{0,1\}^M$ as $g(\mathbf{y}|\mathbf{i},\epsilon) = \prod_{m=1}^{M} f_{i_m}^{(m)}(y^{(m)}|\epsilon_m)$. Next, we derive the joint optimality conditions for the case of unknown error probabilities.

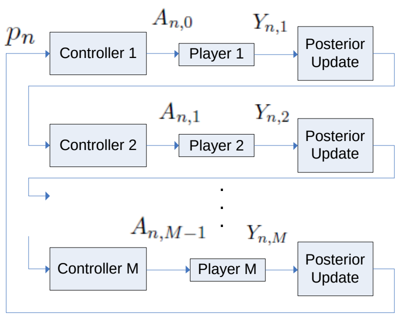

Fig. 2. Sequential scheme for M collaborative players responding to binary valued queries about the location X^* of an unknown target.

Theorem 1. *(Jointly Optimal Policy, Unknown Error Probabilities) Let Assumptions 1 and 2 hold. Consider the problem (1), where the joint policy is made up of the query regions for the M sensors.*

1) Optimal policies $\mathbf{A}_n = (A_n^{(1)}, \ldots, A_n^{(M)})$ at time n satisfy:

$$\underbrace{\sup_{\{A^{(m)}\}_m} H\left(\sum_{\mathbf{i}\in\{0,1\}^M} \int_{\epsilon=0}^{1/2} g(\cdot|\mathbf{i},\epsilon) P_n\left(\bigcap_m (A^{(m)})^{i_m}, \epsilon\right) d\epsilon\right) - \sum_{\mathbf{i}\in\{0,1\}^M} \int_{\epsilon=0}^{1/2} H(g(\cdot|\mathbf{i},\epsilon)) P_n\left(\bigcap_m (A^{(m)})^{i_m}, \epsilon\right) d\epsilon}_{=:G_n^*} \quad (5)$$

2) The maximum information gain at time n is:

$$G_n^* = \sum_{m=1}^{M} \mathbb{E}[C(\epsilon_m)|\mathcal{F}_n] \quad (6)$$

where $\mathbb{E}[C(\epsilon_m)|\mathcal{F}_n] = \int_{\epsilon_m=0}^{1/2} C(\epsilon_m) p_n(\epsilon_m) d\epsilon_m$.

Proof: 1) Optimality conditions
The solution of (1) yields the Bellman recursion:

$$V_n(p_n) = \inf_{\mathbf{A}} \mathbb{E}\left[V_{n+1}(p_{n+1})|\mathbf{A}_n = \mathbf{A}, \mathcal{F}_n\right]$$

Using a similar argument as in Thm. 2 in [2], the optimal solution at time n is given by maximizing the entropy loss:

$$\begin{aligned} G_n^* &= \sup_{\mathbf{A}} I((X^*, \epsilon^*); \mathbf{Y}_{n+1}|\mathbf{A}_n = \mathbf{A}, \mathcal{F}_n) \\ &= \sup_{\mathbf{A}} H(p_n) - \mathbb{E}\left[H(p_{n+1})|\mathbf{A}_n = \mathbf{A}, \mathcal{F}_n\right] \\ &= \sup_{\mathbf{A}_n} H(\mathbf{Y}_{n+1}|\mathbf{A}_n, \mathcal{F}_n) - \mathbb{E}\left[H(\mathbf{Y}_{n+1})|X^*, \epsilon^*, \mathbf{A}_n, \mathcal{F}_n\right] \end{aligned} \quad (7)$$

and the value function is given by $V_n(p_n) = H(p_n) - \sum_{k=n}^{N-1} G_k^*$ for $n < N$ and $V_N(p_N) = H(p_N)$. Rewriting (7) inside the supremum, we obtain (5) [5].
2) Bounds on Maximum entropy loss
Note that the second term in (5) is independent of the queries, so the supremum can be restricted to only the first term without loss of generality. This follows from the additivity of the entropy of a product distribution-i.e., $H(g(\cdot|\mathbf{i},\epsilon)) = \sum_{m=1}^{M} H(f_{i_m}^{(m)}(\cdot|\epsilon_m)) = \sum_{m=1}^{M} h_b(\epsilon_m)$. Using part 1), the capacity formula of BSC ($C(\epsilon_m) = 1 - h_b(\epsilon_m)$), and the fact that the uniform distribution maximizes the entropy (see Ch.2 in [3]), the maximum entropy loss can be bounded as $G_n^* \leq \mathbb{E}\left[\sum_m C(\epsilon_m)\Big|\mathcal{F}_n\right]$ [5]. Using the concavity of $H(\cdot)$ along with Thm. 1 from [1], it can be shown that $G_n^* \geq \mathbb{E}\left[\sum_{m=1}^{M} C(\epsilon_m)\Big|\mathcal{F}_n\right]$ [5]. ■

1) Lower Bound on MSE Performance: The maximum entropy loss derived in Thm. 1 is used next to provide a lower bound on the MSE of the joint sequential estimator.

Theorem 2. *(Lower bound on Joint MSE) Assume $H(p_0)$ is finite. Then, the joint MSE of the joint query policy in Thm. 1 satisfies:*

$$\frac{K}{2\pi e} d \exp\left(-\frac{2n\bar{C}_n}{d}\right) \leq \mathbb{E}[\|X_n - X^*\|_2^2] + \mathbb{E}[\|\epsilon_n - \epsilon^*\|_2^2] \quad (8)$$

where $K = \exp(2H(p_0))$ is a constant and $X_n = \mathbb{E}[X^*|\mathcal{F}_n], \epsilon_n = \mathbb{E}[\epsilon^*|\mathcal{F}_n]$. The average entropy loss after n questions is $\bar{C}_n = \frac{1}{n}\sum_{k=0}^{n-1} G_k^*$.

Proof: The proof is similar to the proof of Thm. 3 in [1] and is included in [5]. ■

2) Discussion: The jointly optimal policy derived in Thm. 1 bears some similarity with the jointly optimal policy of Thm. 1 in [1] that does not apply to the case of unknown channels. We remark that in the unknown channel setting, the maximum entropy loss G_n^* given in (5) is not time-invariant, unlike in the case of known error channels, in which the maximum entropy loss was the sum of the capacities of the players' channels $G^* = \sum_m C(\epsilon_m)$. This observation motivates an adaptive sensor selection; given the constraint that only one sensor may be queried at each time instant, then, unlike in the known channel case, the maximal information gain may be obtained by querying different sensors at different time instants based on the collected information.

B. Sensor Selection Scheme

The control $u_n = u$ denotes that the uth sensor is queried at time n and $A_n^{(u)} = A$ is the associated query region.

Theorem 3. *(Sensor Selection Policy, Unknown Error Probabilities) Consider the problem (1), where the policy is made up of which sensor to choose and the associated query region.*

1) At time n, optimal query policies satisfy:

$$\max_{1\leq u\leq M} G_n^*(u) = \sup_{A} H\left(\int_{\epsilon_u=0}^{1/2} \sum_{i=0}^{1} f_i(\cdot|\epsilon_u) P_n^{(u)}(A^i, \epsilon_u) d\epsilon_u\right) - \int_{\epsilon_u=0}^{1/2} \sum_{i=0}^{1} H\left(f_i(\cdot|\epsilon_u)\right) P_n^{(u)}(A^i, \epsilon_u) d\epsilon_u \quad (9)$$

2) At time n, the maximum entropy loss is:

$$G_n^* = \max_u G_n^*(u) = \max_u \mathbb{E}[C(\epsilon_u)|\mathcal{F}_n]$$

Proof: The proof follows using techniques similar to Thm. 1 and is included in [5]. ■

The optimal policy for the minimum expected entropy criterion (1) shown in Thm. 3 prescribes to use the sensor u with the maximum information gain (measured through the uth sub-marginal distribution $p_n^{(u)}(x, \epsilon_u)$). While the form (9) bears some similarity to the optimality conditions of the known error models (see Thm. 1 in [1]), the bisection policy is no longer optimal.

1) One-dimensional Case: The next corollary specifies the form of the optimal policy derived in Thm. 3 for one-dimensional targets. For simplicity, consider the unit interval $\mathcal{X} = [0, 1]$ as the target domain.

Corollary 1. *(Sensor Selection Policy, Unknown Error Probabilities, One-dimensional Target) Consider the problem (1) for the optimal sensor and query selection policy. Consider the query regions $A_n = [0, x_n]$. The optimal sensor u and associated query region $A = [0, x]$ at time n is given by:*

$$\max_u \left\{ \max_{x \in [0,1]} h_B(g_{1,n}^{(u)}(x)) - c_n^{(u)} \right\} \tag{10}$$

where $h_B(\cdot)$ is the binary entropy function [3] and

$$c_n^{(u)} = \int_{\epsilon_u=0}^{1/2} h_B(\epsilon_u) p_n^{(u)}(\epsilon_u) d\epsilon_u$$
$$g_{1,n}^{(u)}(x) = \int_0^x \mu_n^{(u)}(t)dt + \int_x^1 (p_n(t) - \mu_n^{(u)}(t))dt$$
$$\mu_n^{(u)}(t) = \int_{\epsilon_u=0}^{1/2} \epsilon_u p_n^{(u)}(t, \epsilon_u) d\epsilon_u$$
$$p_n(t) = \int_{\epsilon_1=0}^{1/2} \cdots \int_{\epsilon_M=0}^{1/2} p_n(t, \epsilon_1, \ldots, \epsilon_M) d\epsilon_1 \cdots d\epsilon_M$$

Proof: The proof follows from Thm. 3 [5]. ■

We note that the optimal policy derived for the case of unknown probability in (10) is *not* equivalent to the probabilistic bisection policy-i.e., obtaining $P_n^{(u)}([0, x_n^{(u)}]) = 1/2$ for each sensor u and then evaluating the information gain and choosing the sensor with the maximum information gain. This heuristic scheme would yield a suboptimal information gain as compared to the maximal information gain given by (10). Thus, in the unknown probability setting, the optimal control law is no longer equivalent to the known probability setting (after marginalizing out the noise parameters $\epsilon_1, \ldots, \epsilon_M$). This result shows that the two settings are quite different. We empirically observed that there is a unique query point $x = x_n^* = x_n^{(u^*)}$ that maximizes the function (10). This is similar to the one-dimensional case for the known channel setting when the query region is of the form $A = [0, x]$; i.e., the optimal point is the (unique) median.

C. Sequential Query Design

In this section, we show a version of the separation theorem (Thm. 2 in [1]) for the unknown error channel case.

Theorem 4. *(Separation, Unknown Error Probabilities) Consider the sequential and joint schemes. Then, it follows that $G_{seq,n}^* = \mathbb{E}[\sum_m C(\epsilon_m)|\mathcal{G}_n]$ and $G_n^* = \mathbb{E}[\sum_m C(\epsilon_m)|\mathcal{F}_n]$ for all n.*

Proof: After querying all M players in sequence, using the tower property of expectation and Thm. 1 with $M = 1$ for each sub-instant n_t, the maximal entropy loss can be shown to be $G_{seq,n}^* = \sup_{\{A_{n_t}\}_{t=0}^{M-1}} \mathbb{E}[H(p_n) - H(p_{n+1})|\mathcal{G}_n] = \mathbb{E}[\sum_{m=1}^M C(\epsilon_m)|\mathcal{G}_n]$ [5]. The second part follows from Thm. 1 part 2). ■

IV. Simulation

Fig. 3 shows a simulation result of the MSE performance for $M = 1$ sensor with unknown error probability. This simulation implies that the binary responses obtained from one player carry enough information to accurately estimate both the target and its error probability.

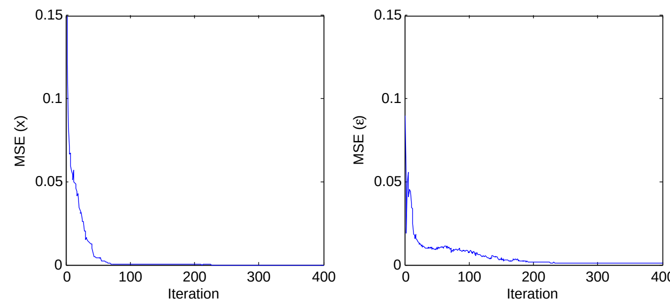

Fig. 3. Monte Carlo simulation for MSE performance of the joint sequential estimator (of the target X^* and the error probability ϵ^*). The MSE for X is shown on the left and MSE for ϵ on the right, as a function of iteration. 100 Monte Carlo trials were used. The true error probability was set to $\epsilon^* = 0.3$ and the true target location was $X^* = 0.75$. The initial distribution was a product of uniform distributions $p_0(x) = I(x \in [0, 1])$ and $p_0(\epsilon) = I(\epsilon \in [0, 1/2))$-i.e., $p_0(x, \epsilon) = p_0(x)p_0(\epsilon)$.

V. Conclusion

We studied the problem of collaborative 20 questions with noise for the multiplayer case under unknown error channels. In this setting, we characterized jointly optimal policies and derived a separation theorem that shows the jointly optimal design is equivalent to a sequential bisection design that can be more easily implemented. Simulations were provided to numerically evaluate the performance of the proposed sequential estimator. Future work may include cost constraints associated with the use of sensors.

References

[1] T. Tsiligkaridis, B. M. Sadler, and A. O. Hero III, "A collaborative 20 questions model for target search with human-machine interaction," in *Proceedings of ICASSP, Vancouver, CA*, May 2013.

[2] B. Jedynak, P. I. Frazier, and R. Sznitman, "Twenty questions with noise: Bayes optimal policies for entropy loss," *Journal of Applied Probability*, vol. 49, pp. 114–136, 2012.

[3] T. D. Cover and J. A. Thomas, *Elements of Information Theory*, Wiley, 2006.

[4] R. Castro and R. Nowak, "Active learning and sampling," in *Foundations and Applications of Sensor Management*. Springer, 2007.

[5] T. Tsiligkaridis, B. M. Sadler, and A. O. Hero III, "Collaborative 20 Questions for Target Localization," *arXiv: 1306.1922*, June 2013.