2015 American Control Conference
Palmer House Hilton
July 1-3, 2015. Chicago, IL, USA

To Ask or Not to Ask: A Foundation for the Optimization of Human-Robot Collaborations

Hong Cai and Yasamin Mostofi

Abstract— **In this paper, we propose a new paradigm for human-robot collaboration. In this paradigm, the collaboration properly takes advantage of the superior visual performance of the humans and the field exploration capabilities of robots, allowing the robot to only ask humans for help when needed. More specifically, we consider a robotic field exploration and classification task with limited communications with a human operator and under a given energy budget. By learning the visual performance of humans probabilistically, we show how the robot can optimize its path planning, sensing, and communication with humans. More specifically, we show when the robot should ask humans for help, when it should rely on its own judgment and when it should gather more information from the field. In order to show the performance of our framework, we then collect several human data using Amazon Mechanical Turk. Our simulation results with real data then confirm that our approach can save the resources considerably. They further reveal interesting behaviors in terms of when to ask humans for help, which we also mathematically characterize.**

I. Introduction

Recent years have seen great developments in robotics, in areas such as navigation, information gathering and group operation. An unmanned vehicle or a network of them can have a tremendous impact in many different areas such as national security, surveillance and battlefield operation. Although robots are becoming more capable of various kinds of missions, there still exist a great number of tasks that robots simply cannot perform to a satisfactory level, when compared to humans. A complex visual task, such as recognition and classification in the presence of uncertainty, is one example of such tasks [1]. Thus, a proper collaboration of humans and robots can be very beneficial to the mission.

More recently, the research community has started to look into the role of humans and different aspects of human-robot collaboration. In control and robotics, for instance, Drift Diffusion Model (DDM) from cognitive psychology [2]–[4] has been heavily utilized in modeling human decision making and the overall collaboration. Various experimental studies have also been conducted on how humans and robots interact and cooperate in simulated scenarios, such as urban search and rescue operations [5], [6]. Chipalkatty [7] shows how to incorporate human factors into a Model Predictive Control (MPC) framework, in which human commands are predicted ahead of time. Srivastava [8] has carried out a comprehensive study on designing a Decision Support System (DSS) to interface field robots and human operators. In computer vision, there are also works on human-machine interface, but without robotic navigation and path planning elements. For instance, Branson et al. [1] proposes a collaboration that resembles the 20-question game between a machine and a human for bird classification.

In this paper, we are interested in the optimization of the human-robot collaboration such that the strengths of both are properly taken into account in task planning and execution. We know humans can do complex visual tasks, such as recognition and classification, in the presence of a high level of uncertainty, while robots can explore terrains harsh for humans. We then propose a new paradigm on an important aspect of the collaboration, *when to ask for human's help*, which has received little attention in the literature to the best of our knowledge. In this approach, the collaboration properly takes advantage of human's superior visual performance and robot's exploration capability, allowing the robot to only ask for human's help when needed. More specifically, consider a robotic field exploration and target classification task, with a given budget for communication with a human operator and a limited motion energy. We show when the robot should ask humans for help, when it should rely on its own classification, and when it should gather more information from the field. This is an important problem since the robot may have a limited chance of communication with the human (due to limited bandwidth for instance) as well as limited energy/time budget for field exploration. Thus, it cannot bug humans all the time for help with classification (which means sending raw sensing data to humans). On the other hand, it may not have enough resources to explore the field (and reduce the sensing uncertainty) to the level that its own classification over the whole field becomes acceptable. We then show how the robot can optimally take advantage of its collaboration with the human via a co-optimization of its navigation, sensing and query to humans.

In order to do so, the robot only needs to understand the extent of human's visual capabilities, as compared to its own performance. We therefore first start with characterizing human's visual performance, i.e. when/under what conditions humans can make a correct visual decision in the presence of uncertainty. For instance, a robot may collect data with a high level of uncertainty. Yet, the human may be able to make sense out of this data and do an accurate classification of the target of interest. If the robot can properly understand this, it can then judge if it should stop sensing and present the data to the human, or if it should sense and gather more data. By using input from several human subjects via Amazon Mechanical Turk, we first show how

Hong Cai and Yasamin Mostofi are with the Department of Electrical and Computer Engineering, University of California, Santa Barbara, USA (email: {hcai, ymostofi}@ece.ucsb.edu).

This work is supported in part by NSF NeTS award number 1321171.

to probabilistically model human's and robot's performances in Section II. In Sections III and IV, we then propose our approach for the co-optimization of human collaboration and field exploration. We further discuss and mathematically characterize important properties of "optimal bugging" that emerge in our framework. Our simulation results with real data then confirm that our approach can save the resources considerably. They further reveal interesting behaviors in terms of when to ask humans for help.

II. Human and Robot's Performance in Target Classification

As we discussed in the introduction, the proposed underlying methodology of this paper is that the robot does not need to understand how human makes superior visual decisions but rather needs to understand the extent of human's visual capability, as compared to its own, in order to best take advantage of its collaboration. In this section, we focus on this observation and analyze human and robot's performances in target classification in the presence of noise/uncertainty. The understanding we gain in this section will then be utilized for the optimization of the collaboration and robotic exploration in the next sections. We note that our underlying methodology is applicable to any collaborative task and not just target classification.

Consider the case where the robot has sensed a target, for instance via taking a picture, and needs to classify it based on an a priori known set of target possibilities. For instance, Fig. 1 shows 4 possible images that are presented to the robot before it starts its task. The robot then needs to classify an image it takes during its mission to one of these four images. The sensing of the robot in the field, however, is prone to uncertainty. For instance, its sensing could be corrupted by noise, missing parts, and low resolution, which will impact its classification capability. If the robot could model the exact way that different sources of uncertainties have impacted its sensing and derive the best detector (the one that maximizes the posterior probability), then its performance would actually be better than that of human's. For instance, if each image was corrupted by an Additive White Gaussian Noise (AWGN) with known statistics, then the robot will outperform the human. This is due to the fact that in this case, the robot can derive the exact form of the optimum detector to use, while human's visual performance is somewhat worse than the optimum detector. In fact, in cognitive psychology, some imperfections are added to the optimum detector to model human's visual performance in this simple case [9]. In most realistic tasks, however, it is not possible for the robot to know/properly model the way all the different sources of uncertainties have impacted its sensing. Furthermore, deriving the optimum detector, with the current state-of-the-art in vision and image processing, may not be possible due to the complexity of the visual task. This is why the robot can benefit from the collaboration with the human tremendously.

We therefore consider the following setting in this paper to capture an example of the case where human's visual

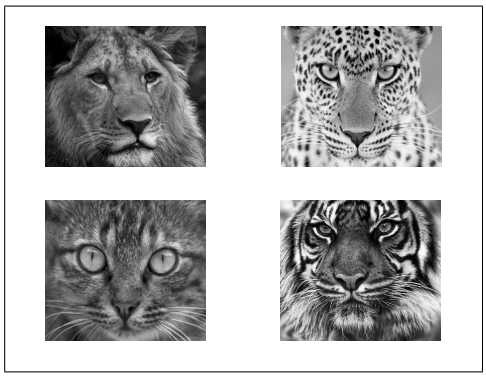

Fig. 1. Gray-scale test images of lion, leopard, cat and tiger that are used in our study.

performance is considerably superior to that of the robot. We assume that each sensing sample of the robot (an image in this case) is first corrupted by an AWGN noise with a known variance but an unknown mean, and then undergo a truncation process that is unknown to the robot. While this is a toy example, it captures a possible realistic scenario. For instance, the robot may be able to assess its noise variance based on its distance to the target on the field but may not know the mean of the added noise or the nonlinear truncation that has happened at the pixel level. We further emphasize that our proposed framework is applicable to any other example of a visual task in the presence of uncertainty when the robot cannot use the optimum detector.

Let $s \in S = \{s_1, ..., s_T\}$ represent the original true image in vector form (pixel values stacked up), which belongs to a set of T possible images. Let $\hat{r} = s + n$ represent the noisy image vector after it gets corrupted by an AWGN noise $n \sim N(\mu, \sigma^2 I)$, where μ and σ^2 denote the corresponding mean and variance respectively, and I is the identity matrix. Let r represent the final measured vector after truncation of each pixel to the interval $[0,1]$: $r = \text{trunc}(\hat{r})$, where trunc(.) is the element-wise truncation function. We next discuss the detector that the robot is using, based on its modeling and knowledge of the sources of uncertainty.

A. Robot's Classifier

From robot's perspective, the only source of uncertainty is an AWGN noise with a zero mean and known variance (σ^2 in this case). Then, it uses the optimal MAP (Maximum A Posteriori) detector for this case, which can be easily shown to result in the following classifier: $i = \arg\max_{s_i \in S} r^T s_i - \mathscr{E}_i/2 + \sigma^2 \ln P_i$, where s_i is the i^{th} original image from S, $\mathscr{E}_i = ||s_i||_2^2$ is the energy of the i^{th} image, P_i is the prior probability that the i^{th} image will appear, and all images are assumed equal size. Under the assumption of equal a priori probabilities, the MAP detector then reduces to the minimum distance detector: $i = \arg\min_{s_i \in S} ||r - s_i||_2^2$.

B. Experimental Study of Human and Robot Performances

Fig. 2 (solid line) shows the true probability of correct classification of the robot when it is using the minimun distance detector but the noise is not zero mean and there is an additional nonlinearity due to truncation as we discussed

earlier. The mean of the noise is $\mu = 0.3$ in this case. Fig.3 shows a sample corrupted image with the noise variance of 3, which results in the probability of correct classification of 0.5 for the robot. It can be seen that the information is fairly corrupted by the uncertainty in this case, due to the high level of noise.

Fig. 2 (dashed line) further shows the performance of the human. In order to get human's performance curve, we have utilized Amazon Mechanical Turk (MTurk) to present several noisy images to anonymous human subjects. More specifically, for each variance, 20 noisy images are generated and presented to 50 MTurk subjects. It can be seen that the human outperforms the robot considerably, as expected, especially if the variance is not too small or too high. For instance, for the sample case of Fig. 3, the human can achieve an average probability of correct classification of 0.744, which is considerably higher than robot's performance (0.5). Both performance curves will eventually drop to 0.25 as we continue to increase the variance.

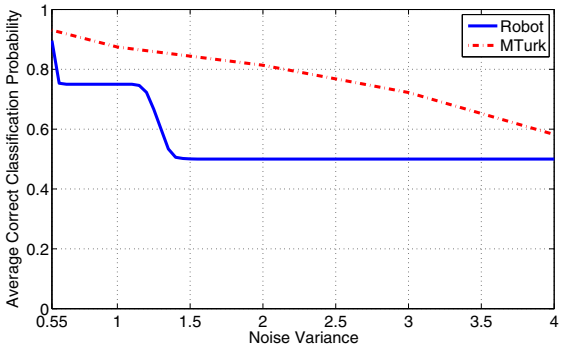

Fig. 2. Study of the performance of human and robot in target classification. The human data is acquired using Amazon MTurk. It can be seen that the human outperforms the robot considerably as expected.

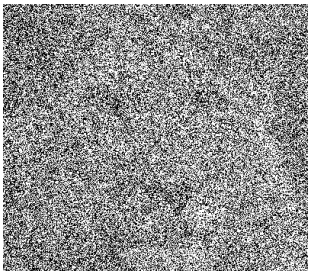

Fig. 3. A sample corrupted image (lion) with the noise variance of 3.

Curves like Fig. 2 will provide the basis for our human robot collaboration. There are two pieces key to obtaining these curves. First, there is a metric (or a number of metrics) that the robot will use, such as the noise variance that we have utilized in this paper. More specifically, a metric is a parameter that the robot can assess on the field. For instance, it can assess a variance in this case possibly based on its distance to the target. The next important piece is to get the performance curves as a function of the metric. This is what the robot will do either online (via a training phase) or from past human data. The robot's ongoing task with the human will then serve to further add to the pool of available data. It should be noted that, in principle, the data for the curves do not have to be all acquired for the same exact set of images. More exploration of this is a subject of our future work.

In the next section, we utilize these curves and propose our framework for optimizing human-robot collaboration.

III. Optimizing Human-Robot Collaboration in a Target Classification Task

In this section, we propose a new approach for human-robot collaboration in a target classification task. We consider a simple setup in which the robot has an initial sensing (in the form of an acquired image) of N given sites. The robot is given a total motion energy budget as well as a total number of questions to ask the human (dictated by the frequency of the availability of communication with the human). For each site, the robot has three choices: 1) rely on its own classification (based on the initial sensing), 2) use a question and present the data of the site to the human, or 3) spend motion energy to go to the site and sense better. Our setup of this section is simple to allow for bringing an understanding of the optimal collaboration and will be extended to a surveillance and path planning scenario in Section IV.

A. Problem Setup

Consider the case where we have a total number of N sites. The sensing model of the robot is as explained in the previous section. In summary, each site contains one of T a priori known targets (see Fig. 1 for an example with $T = 4$ targets). The sensing of the robot is then corrupted by an additive Gaussian noise with an unknown mean but a known variance, and is truncated. The robot uses the classifier of Section II-A, which would have been ideal for the case of a zero-mean additive Gaussian noise with a known variance. The robot is allowed to query the human M times and has a total motion energy budget of $\mathscr{E}_{\max}$. The probabilities of correct target classification of the k^{th} site, for $k \in \{1,...,N\}$, are denoted by $p_{r,k}$ and $p_{h,k}$ for the robot and human respectively. These probabilities are obtained from Fig. 2, based on the variance assessed by the robot. Let $\mathscr{E}_k$ denote the motion energy cost to visit the k^{th} site, which can be evaluated by the robot. If the robot chooses to visit a site, the probability of correct classification increases to a high value of $\tilde{p}$. The objective of the robot is then to decide which sites to present to the human, which sites to visit and which sites to rely on its own classification based on its initial sensing, in order to maximize the overall average probability of correct classification under resource constraints. Let P_{corr} denote the average probability of correct classification of a site. We have

$$\begin{aligned} P_{\text{corr}} &= \frac{1}{N}\Big(\sum_{k=1}^{N}\gamma_k p_{h,k} + \sum_{k=1}^{N}\eta_k\tilde{p} + \sum_{k=1}^{N}(1-\gamma_k)(1-\eta_k)p_{r,k}\Big), \\ &= \frac{1}{N}\Big(\sum_{k=1}^{N}\gamma_k(p_{h,k}-p_{r,k}) + \sum_{k=1}^{N}\eta_k(\tilde{p}-p_{r,k}) + \sum_{k=1}^{N}p_{r,k}\Big), \end{aligned}$$

where γ_k is 1 if the robot seeks human's help for the k^{th} site and is 0 otherwise. Similarly, $\eta_k = 1$ indicates that the robot

will physically visit the k^{th} site and $\eta_k = 0$ denotes otherwise. We then have the following optimization problem:

$$\begin{aligned} \max_{\gamma,\eta} \quad & \gamma^T(p_h - p_r) + \eta^T(\tilde{p}\mathbf{1} - p_r) \\ \text{s.t.} \quad & \eta^T \mathscr{E} \le \mathscr{E}_{\max}, \quad \mathbf{1}^T\gamma \le M, \\ & \gamma, \eta, \gamma+\eta \in \{0,1\}^N, \end{aligned} \tag{1}$$

where $p_h = [p_{h,1},...,p_{h,N}]^T$, $p_r = [p_{r,1},...,p_{r,N}]^T$, $\gamma = [\gamma_1,...,\gamma_N]^T$, $\eta = [\eta_1,...,\eta_N]^T$, $\mathscr{E} = [\mathscr{E}_1,...,\mathscr{E}_N]^T$, and $\mathbf{1}$ is the vector with all 1s.

It can be seen that $(p_{h,k} - p_{r,k})$ and $(\tilde{p} - p_{r,k})$ are important parameters as they represent the performance gains by asking the human and visiting the k^{th} site respectively. Note that we do not allow the robot to both query the human and make a visit for the same site. This is because we already assume a high probability of correct classification when the robot visits a site. Thus, allowing the robot to both visit and ask about the same site will be a waste of resources in this case. In Section IV we consider a more general setting and relax this assumption as there may be cases where it makes sense for the robot to visit a site and then ask the human.

B. Properties of Optimal Bugging

In this section, we discuss some properties of the optimization problem (1). We first consider a special case where the exact solution can be derived and then show a few properties of the general optimum solution.

1) Zero Motion Energy Budget: If $\mathscr{E}_{\max} = 0$, problem (1) reduces to the case where the robot needs to decide between asking the human and relying on its initial classification. It can then be easily confirmed that γ_k has to be one for the M sites with the largest $p_{h,k} - p_{r,k}$ in this case.

2) Considering the General Case: In general, optimization problem (1) is a Mixed Integer Program, which makes theoretical analysis challenging. In order to better understand the optimum solution of problem (1), Fig. 4 shows an example of the optimum decisions for the case of $N = 2000$ sites, with $M = 500$ allowed questions and an energy budget equal to 25% of the total energy needed to visit all the sites. The optimum decision for each site is marked based on solving problem (1). Interesting behavior emerges as can be seen. For instance, we can observe that there are separations between different decisions. The clearest patterns are two transition points that mark when the robot asks human for help, as shown with the dashed vertical lines in Fig. 4. Basically, the figure suggests that the robot should not bug the human if the variance is smaller than a threshold or bigger than another threshold, independent of the motion cost of a site. This makes sense as the robot itself will perform well for low variances and human does not perform well for high variances, suggesting an optimum query range. Furthermore, it shows that the robot only visits the sites where the energy cost is not too high and relies more on itself for the sites with both high variance and energy cost.

In the rest of this section, our goal is to bring more analytical understanding to the properties of the optimum collaboration. We consider a relaxed approximated version

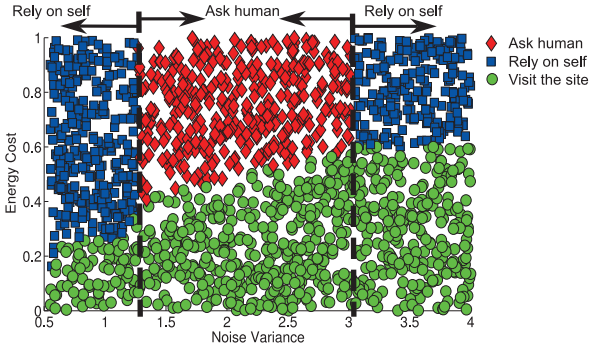

Fig. 4. An example of the optimum decisions with 2000 sites, 500 questions and an energy budget of 25% of the total energy needed to visit all the sites.

of problem (1) where the integers γ, η and $\gamma+\eta$ can take continuous values from $[0,1]^N$, as follows:

$$\begin{aligned} \max_{\gamma,\eta} \quad & \gamma^T(p_h - p_r) + \eta^T(\tilde{p}\mathbf{1} - p_r) \\ \text{s.t.} \quad & \eta^T \mathscr{E} \le \mathscr{E}_{\max}, \quad \mathbf{1}^T\gamma \le M, \\ & 0 \preceq \gamma, \eta, \omega = \gamma+\eta \preceq 1. \end{aligned} \tag{2}$$

This results in a linear program, which we can analyze by applying Karush-Kuhn-Tucker (KKT) conditions [10]. We then have the following expression for the Lagrangian:

$$\begin{aligned} \mathscr{L}(\gamma,\eta,\omega,\lambda_1,\lambda_2,\theta,\zeta,\kappa,\tau,\xi,\psi) &= \gamma^T(p_r - p_h) + \eta^T(p_r - \tilde{p}\mathbf{1}) \\ &- \lambda_1(M - \mathbf{1}^T\gamma) - \lambda_2(\mathscr{E}_{\max} - \eta^T\mathscr{E}) - \zeta^T\gamma - \theta^T(1-\gamma) - \tau^T\eta \\ &- \kappa^T(1-\eta) - \psi^T\omega - \xi^T(1-\omega), \end{aligned}$$

where $\lambda_1, \lambda_2, \theta, \zeta, \kappa, \tau, \xi, \psi$ are non-negative Lagrange multipliers. The optimum solution (marked by $\star$) then satisfies the following KKT conditions, in addition to the primal/dual feasibility conditions:

1) Gradient condition, for $k \in \{1,...,N\}$:

$$\nabla_{\gamma_k^\star}\mathscr{L} = p_{r,k} - p_{h,k} + \lambda_1^\star - \zeta_k^\star + \theta_k^\star - \psi_k^\star + \xi_k^\star = 0, \tag{3}$$

$$\nabla_{\eta_k^\star}\mathscr{L} = p_{r,k} - \tilde{p} + \lambda_2^\star\mathscr{E}_k - \tau_k^\star + \kappa_k^\star - \psi_k^\star + \xi_k^\star = 0. \tag{4}$$

2) Complementary slackness: $\theta^\star(1-\gamma^\star) = \mathbf{0}$, $\zeta^\star\gamma^\star = \mathbf{0}$, $\kappa^\star(1-\eta^\star) = \mathbf{0}$, $\tau^\star\eta^\star = \mathbf{0}$, $\xi^\star(1-\omega^\star) = \mathbf{0}$, $\psi^\star\omega^\star = \mathbf{0}$, $\lambda_1^\star(M - \mathbf{1}^T\gamma^\star) = 0$, $\lambda_2^\star(\mathscr{E}_{\max} - \mathscr{E}^T\eta^\star) = 0$, where $\mathbf{0}$ is the all 0 vector.

Next, we characterize necessary conditions for the optimum solution of problem (1).

Lemma 1: Consider two sites k and l. Let $\eta^\star$ and $\gamma^\star$ denote the optimum decision vectors for these sites.

1) If $\gamma_k^\star = 1, \eta_k^\star = 0, \gamma_l^\star = 0$ and $\eta_l^\star = 0$, then $p_{h,k} - p_{r,k} \ge p_{h,l} - p_{r,l}$.

2) If $\gamma_k^\star = 0, \eta_k^\star = 1, \gamma_l^\star = 0$ and $\eta_l^\star = 0$, then $(\tilde{p} - p_{r,k})/\mathscr{E}_k \ge (\tilde{p} - p_{r,l})/\mathscr{E}_l$.

Proof: 1) Suppose that we have two sites k and l that satisfy the 1st decision condition in Lemma 1. Applying the complementary slackness results in $\zeta_k^\star = \psi_k^\star = \theta_l^\star = \xi_l^\star = 0$. Then, the gradient condition gives $p_{r,k} - p_{h,k} + \theta_k^\star + \xi_k^\star = p_{r,l} - p_{h,l} - \zeta_l^\star - \psi_l^\star$. Since $\theta_k^\star, \xi_k^\star, \zeta_l^\star, \psi_l^\star$ are all nonnegative, we need to have $p_{h,k} - p_{r,k} \ge p_{h,l} - p_{r,l}$.

2) Suppose that we have two sites k and l, which satisfy the 2nd decision condition in Lemma 1. We have $\tau_k^\star = \psi_k^\star = \kappa_l^\star = \xi_l^\star = 0$ from complementary slackness. Eq. (4) then

becomes: $(p_{r,k}-\tilde{p})/\mathscr{E}_k+\lambda_2^\star+\kappa_k'+\xi_k'=0$, where $\kappa_k'=\kappa_k^\star/\mathscr{E}_k$ and $\xi_k'=\xi^\star/\mathscr{E}_k$. Similarly, we have $(p_{r,l}-\tilde{p})/\mathscr{E}_l+\lambda_2^\star-\tau_l'-\psi_l'=0$ for $\nabla_{\eta_l^\star}\mathscr{L}$, resulting in $(p_{r,k}-\tilde{p})/\mathscr{E}_k+\lambda_2^\star+\kappa_k'+\xi_k'=(p_{r,l}-\tilde{p})/\mathscr{E}_l+\lambda_2^\star-\tau_l'-\psi_l'$. Since $\kappa_k^\star,\xi_k^\star,\tau_l^\star,\psi_l^\star$ are all nonnegative, we have the condition in part 2 of the lemma. ■

This lemma says that if we have two sites k and l, for which the robot will ask the human and rely on its initial sensing respectively, then $p_{h,k}-p_{r,k}\geq p_{h,l}-p_{r,l}$; there should be a greater benefit by asking the human. Similarly, if we have two sites k and l, for which the robot will explore and rely on its initial sensing respectively, then $(\tilde{p}-p_{r,k})/\mathscr{E}_k\geq(\tilde{p}-p_{r,l})/\mathscr{E}_l$; the visited site should have a higher information gain normalized by the energy cost.

C. Numerical Results

In this part, we show the performance of our proposed collaboration approach for target classification. By using the collected MTurk data of Fig. 2, the robot solves the optimization problem (1), using the Mixed Integer Linear Program solver of MATLAB. The numerical results and curves are based on MTurk human data with actual human decisions.

We show the energy and bandwidth savings of our proposed approach as compared to a possible state-of-the-art methodology where human's collaboration is not fully optimized, to which we refer as the benchmark method. In the benchmark approach, the robot optimizes its given energy budget to best explore the field based on site variances, i.e. it chooses the sites that maximize the sum of noise variances. It then randomly chooses from the remaining sites to ask the human, given the total number of questions. In other words, the robot optimizes its energy usage without any knowledge of the human's performance.

1) Energy Saving: Table I shows the amount of motion energy the robot saves, by using our approach, for achieving a target probability of correct classification. The first column shows the target average probability of correct classification, while the second column shows the percentage reduction of the needed energy by using our proposed approach compared to the benchmark method. In this case, there is a total of $N=10$ sites and $M=4$ given queries. The noise variance of each site is randomly assigned from the interval $[0.55,4]$. $\tilde{p}$ is set to 0.896 - the best achievable robot performance based on Fig. 2. The motion energy cost to visit each site is also assigned randomly and the total given energy budget is taken to be a percentage of the total energy required to visit all the sites. It can be seen that the robot can reduce its energy consumption considerably by properly taking advantage of its collaboration. For instance, it can achieve an average probability of correct classification of 0.7 with 66.67% less energy consumption. The term "Inf" denotes the cases where the benchmark cannot simply achieve the given target probability of classification.

2) Bandwidth Saving: Next, we show explicitly how our proposed approach can also result in a considerable communication bandwidth saving by reducing the number of questions. More specifically, consider the benchmark approach with "large bandwidth" and "zero bandwidth". In the first case, the robot has no communication limitation and can probe the human with as many questions as it wants to (10 in this case) after exploring the field. In the latter, no access to a human operator is available and thus the robot has to rely on itself to classify the gathered data after it surveys the field. Fig. 5 compares the performance of our proposed approach with these two cases. The robot is given an energy budget of 30% of the total energy needed to visit all the sites.

TABLE I

Target Ave. Correct Classification Prob.	% Energy Saving
0.7	66.67%
0.75	44.30%
0.8	27.83%
0.85	6.3%
0.9	0.71%
0.915	Inf

Energy saving as compared to the case of no proper collaboration.

As expected, the case of "no bandwidth" performs considerably poorly as the robot could not seek human's help in classification. On the other hand, the case of "large bandwidth" performs considerably well as the robot fully relies on the human for all the classification after it explores the field. This, however, comes at a cost of excessive communication and thus a high bandwidth usage. It can be seen that our proposed approach can achieve a performance very close to this upper bound with a much less bandwidth usage. For instance, we can see that by asking only 6 questions (40% bandwidth reduction), the robot can achieve an average probability of correct classification of 0.888, which is only 4.3% less than the case of large bandwidth (0.928 in this case). Overall, we can see that by using our proposed collaboration approach, considerable amount of motion energy and bandwidth can be preserved.

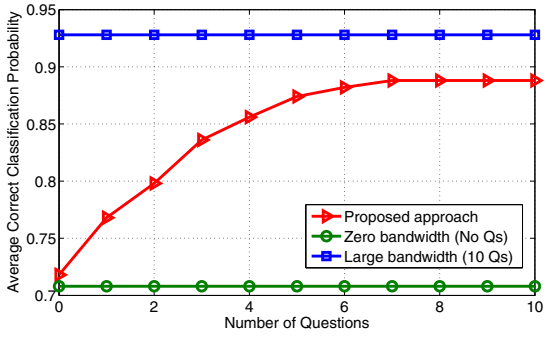

Fig. 5. Average probability of correct classification in a human-robot collaboration as a function of the total number of given queries. In this example, there are a total of 10 sites and the given motion energy budget is 30% of what is needed to visit all the sites.

Table II shows the amount of bandwidth the robot can save by using our approach, when trying to achieve a given target probability of correct classification.[1] The first column shows the target average probability of correct classification and the second column shows the percentage reduction of

[1]Bandwidth usage is taken proportional to the number of questions.

the needed bandwidth by using our proposed approach as compared to the benchmark. In this case, the robot is given an energy budget of 30% of the total energy needed to visit all the sites. It can be seen that the robot can reduce its bandwidth consumption considerably. For instance, it can achieve an average probability of correct classification of 0.75 with 48.61% less bandwidth usage.

TABLE II

Target Ave. Correct Classification Prob.	**% Bandwidth Saving**
0.7	37.04%
0.75	48.61%
0.8	33.18%
0.85	7.33%
0.875	Inf

Bandwidth saving as compared to the case of no proper collaboration.

IV. Optimizing Human-Machine Collaboration in a Surveillance Task

In this section, we extend our proposed human-robot collaboration approach of Section III to a surveillance scenario, where the robot explicitly plans its trajectory in the operation. More specifically, the robot is given a task of surveying a field and classifying a number of targets, under a motion energy budget constraint. The robot is also given a limited number of questions to ask the human. Thus the robot needs to plan its trajectory to best survey the site while deciding on which sites to ask the human. The trajectory of the robot affects its variance in sensing a site and thus the corresponding $p_{r,k}$ and $p_{h,k}$ values, making the path planning dependent on the human's and robot's performances. The purpose of this setup is then to bring an understanding to how collaboration with the human affects the site surveillance.

A. Problem Setup

Consider the case that there are N targets located in a field. The robot's goal is to survey the field and maximize the correct classification probability of the targets, by properly planning its trajectory under a limited energy budget $\mathscr{E}_{\max}$. The robot can also ask M questions to the human. The sensing model is the same as explained in Section III. The noise variance is distance dependent and can be explicitly expressed as follows for sensing the i^{th} site: $\sigma_i(d) = \alpha_i d^2 + \beta_i$, where d is the distance from the robot to the i^{th} site and α_i and β_i are positive constants. α_i and β_i model the fact that some sites may be easier to sense than others.

We then have the following optimization problem:

$$\begin{aligned} \max_{x,\gamma} \quad & \frac{1}{N}\Big(\sum_{k=1}^{N}(1-\gamma_k)\max_{x_i}\ p_{r,k}(x_i)+\gamma_k\max_{x_i}\ p_{h,k}(x_i)\Big) \\ \text{s.t.} \quad & \sum_{k=1}^{N}\gamma_k \le M, \quad \mathscr{E}(x) \le \mathscr{E}_{\max}, \\ & ||x_k - x_{k+1}||_2 \le \delta_r, \quad \forall k = 1,2,...,x_{\text{num}}-1, \\ & \gamma_k \in \{0,1\}, \quad \forall k = 1,2,...,N, \end{aligned} \tag{5}$$

where $\gamma_k = 1$ if the robot asks the human about the k^{th} site and 0 otherwise, $x = [x_1, x_2, ...]^T$ are the way-points of the trajectory with x_{num} denoting the number of points, $p_{r,k}(x_i)$ and $p_{h,k}(x_i)$ denote the robot's and human's probability of correctly classifying target k if sensed from position x_i, and $\mathscr{E}(x)$ evaluates the motion energy cost of trajectory x. Furthermore, $||x_k - x_{k+1}||_2 \le \delta_r$ imposes a speed limit on the robot by limiting its step size.

In general, solving Eq. (5) is challenging even without the binary constraints as it requires maximizing the objective function over all possible trajectories under a number of other constraints. Thus, we utilize Rapidly Exploring Random Tree Star (RRT*) algorithm [11], [12] to solve this problem. There are a number of reasons to use this algorithm: 1) it is a fast algorithm that can handle trajectory planning problems, which generally incur high computational complexity; 2) we can incorporate our binary constraints in it easily.

B. Using RRT* for Path Planning and Question Optimization

We slightly modify the original RRT* algorithm to fit our problem. More specifically, RRT* is a recursive graph-based path generation approach. In each iteration, it returns a trajectory for which we need to evaluate the optimal value. We next examine the objective function of problem (5) if the trajectory is given. It can be seen that, for a given trajectory x, the optimal solution is to choose the M sites with the highest $\max_{x_i \in S(x)} p_{h,k}(x_i) - \max_{x_i \in S(x)} p_{r,k}(x_i)$, where $S(.)$ is the set of way-points in x. Then, given that $p_{h,k}(.)$ and $p_{r,k}(.)$ are monotonic functions of the variance, this becomes the same as the special case of Section III-B.1. After a sufficient number of iterations, RRT* generates a large number of trajectories. By selecting the trajectory with the maximum average correct classification probability, a good solution to the original problem (5) can then be obtained.

C. Simulation Results

In this part, we show the performance of our approach in a surveillance scenario. We consider a total of 10 sites and 3 given queries. The motion energy is taken proportional to the distance traveled. The field has three types of targets in terms of sensing difficulty. More specifically, we have three αs (0.005, 0.024 and 0.123). Furthermore, $\beta = 0.55$ is assumed for all the sites. Fig. 6 shows the field where the sites are marked with circles, with a radius that is an indicator of the level of difficulty (thus three different sizes where the larger circle indicates a higher α).

Fig. 6 shows the resulting surveillance trajectory, starting from the blue diamond to the green dot. It can be seen that the robot has a tendency to travel near sites with a higher sensing difficulty. It also mainly asks questions on the sites with a medium sensing difficulty. This is in line with the observation of Fig. 4 where the robot asked questions when the variance was neither too high nor too low. Finally, letting $p_{h,k} = \max_{x_i \in S(x)} p_{h,k}(x_i)$ and $p_{r,k} = \max_{x_i \in S(x)} p_{r,k}(x_i)$, the shade of the square on each site indicates the value of $p_{h,k} - p_{r,k}$ of that site, while traversing this trajectory, with a lighter shade indicating a higher value. As expected, the robot asks

questions on the sites with the highest $p_{h,k} - p_{r,k}$ along the trajectory.

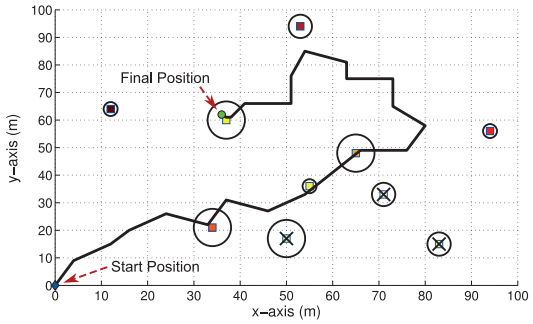

Fig. 6. Path planning for collaborative human-robot target classification using RRT*. There are 10 targets, 3 given questions and 150m allowed travel distance.

1) Energy Saving: Table III shows the percentage energy saved by properly seeking human's help, as compared to the case of no questions. In this case, the goal is to reach a target average correct classification probability of 0.8. It can be seen that the motion energy saving is significant, with nearly 45% reduction by allowing only two questions.

TABLE III

# of Queries	% Energy Saving as Compared to No Collaboration
2	44.96%
5	69.39%
10	93.99%

Energy saving as compared to the case of no proper collaboration.

2) Bandwidth Saving: Consider the two cases of "large bandwidth" and "zero bandwidth", introduced in Section III-C.2. Fig. 7 compares the performance of our proposed collaborative surveillance strategy with these two cases. It can be seen that our proposed approach can achieve a performance very close to the upper bound, but with a considerable reduction in the bandwidth usage. For instance, we can see that by asking only 6 questions (40% bandwidth reduction), the robot can achieve an average probability of correct classification of 0.885, which is only 3.8% less than that of the case of large bandwidth (0.920).

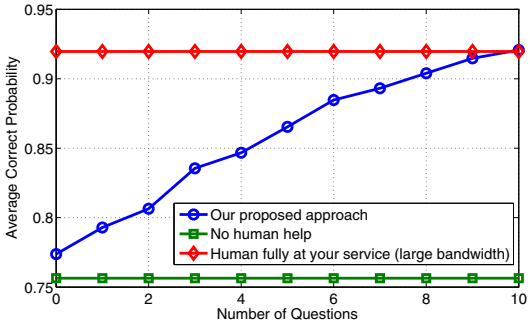

Fig. 7. Average probability of correct classification in a human-robot collaboration as a function of the total number of given queries. In this example, there are a total of 10 sites and the robot is allowed to travel 100m.

V. Conclusions and Future Directions

In this paper, we proposed a new framework for the optimization of human-robot collaboration that properly takes advantage of the strengths of both, i.e. human's superior visual performance and robot's exploration capability. Our goal was to bring an understanding to when the robot should ask for human's help, an area that has not been much explored. More specifically, we considered a robotic field exploration and target classification task with limited communications with a human operator and under a given energy budget. We developed a new foundation for "optimal bugging" through a co-optimization of sensing, navigation, and questioning, i.e. we showed when the robot should ask humans for help, when it should rely on its own judgment and when it should gather more information from the field. In order to show the performance of our framework, we then collect several human data using Amazon MTurk. Our simulation results with real data then confirmed that our approach can save the resources considerably via optimal bugging. They further revealed interesting behaviors in terms of when to ask humans for help, which we also mathematically characterized. There are several possible extensions for this work. For instance, we are currently working on understanding human and robot's performance in the presence of other sources of uncertainty. Furthermore, while we considered a specific task of target classification, our underlying methodology is applicable to any collaborative scenario.

Acknowledgement

The authors would like to thank Yuan Yan for the useful discussions.

References

[1] S. Branson, C. Wah, F. Schroff, B. Babenko, P. Welinder, P. Perona, and S. Belongie. Visual recognition with humans in the loop. In *European Conference on Computer Vision*, pages 438–451. 2010.

[2] M. Cao, A. Stewart, and N. Leonard. Integrating human and robot decision-making dynamics with feedback: models and convergence analysis. In *47th IEEE Conference on Decision and Control*, pages 1127–1132, 2008.

[3] A. Stewart, M. Cao, A. Nedic, D. Tomlin, and N. Leonard. Towards human–robot teams: model-based analysis of human decision making in two-alternative choice tasks with social feedback. *Proceedings of the IEEE*, 100(3):751–775, 2012.

[4] K. Savla, T. Temple, and E. Frazzoli. Human-in-the-loop vehicle routing policies for dynamic environments. In *47th IEEE Conference on Decision and Control*, pages 1145–1150, 2008.

[5] J. Burke, R. Murphy, M. Coovert, and D. Riddle. Moonlight in miami: Field study of human-robot interaction in the context of an urban search and rescue disaster response training exercise. *Human–Computer Interaction*, 19(1-2):85–116, 2004.

[6] M. B. Dias, B. Kannan, B. Browning, E. Jones, B. Argall, M. F. Dias, M. Zinck, M. Veloso, and A. Stentz. Sliding autonomy for peer-to-peer human-robot teams. In *10th International Conference on Intelligent Autonomous Systems*, pages 332–341, 2008.

[7] R. Chipalkatty. *Human-in-the-loop control for cooperative human-robot tasks*. PhD thesis, Georgia Institute of Technology, 2012.

[8] V. Srivastava. *Stochastic search and surveillance strategies for mixed human-robot teams*. PhD thesis, University of California, Santa Barbara, 2012.

[9] M. Eckstein, C. Abbey, and F. Bochud. A practical guide to model observers for visual detection in synthetic and natural noisy images. *Handbook of medical imaging*, 1:593–628, 2000.

[10] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2009.

[11] S. Karaman and E. Frazzoli. Incremental sampling-based algorithms for optimal motion planning. *arXiv preprint arXiv:1005.0416*, 2010.

[12] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research*, 30(7):846–894, 2011.