Decision Referrals in Human-Automation Teams

Kesav Kaza, Jerome Le Ny and Aditya Mahajan

***Abstract*— We consider a model for optimal decision referrals in human-automation teams performing binary classification tasks. The automation observes a batch of independent tasks, analyzes them, and has the option to refer a subset of them to a human operator. The human operator performs fresh analysis of the tasks referred to him. Our key modeling assumption is that the human performance degrades with workload (i.e., the number of tasks referred to human). We model the problem as a stochastic optimization problem. We first consider the special case when the workload of the human is pre-specified. We show that in this setting it is optimal to myopically refer tasks which lead to the largest reduction in the conditional expected cost until the desired workload target is met. We next consider the general setting where there is no constraint on the workload. We leverage the solution of the previous step and provide a search algorithm to efficiently find the optimal set of tasks to refer. Finally, we present a numerical study to compare the performance of our algorithm with some baseline allocation policies.**

I. Introduction

In recent years there has been a significant interest in developing collaborative systems where automated agents team with human operators to perform a collaborative task such as monitoring an environment [1] or an industrial control system [2], manipulating objects [3], identifying dynamic threats or searching for objects in military and public safety applications [4], etc. To enable effective collaboration between an automated agent and a human operator, it is critical to design appropriate information sharing strategies that can lead to optimal decisions by the team. In this paper, we develop such a strategy for binary classification tasks performed by a human operator with the help of an automated decision support system (DSS).

The literature developing quantitative models for collaborative decision-making and workload distribution in mixed human-automation teams with hierarchical structure is relatively limited, despite the increasing use of such systems [5]–[7]. In contrast, much work has been done on distributed and collaborative decision-making in purely automated systems [8]–[11]. In the classical distributed or decentralized hypothesis testing problem, multiple sensors or decision makers receive observations depending on the true state of nature, which is the same for all of them. They might transmit either all or part of this information to a central entity making the final decisions, or they might try to achieve a consensus decision without a central coordinator. However, this literature does not take into account the impact on the decision maker performance of various human characteristics such as cognitive workload, fatigue, trust or belief in automation capability [2], [12]–[14].

Important factors to consider when forming human-automation teams are the dependence of human performance on workload [15] and an appropriate allocation of tasks between the automation and human operators [16]. There has been some work on decision queues where a sequence of general tasks arrive at a human operator, modeled as a server with utilization-dependent performance [17]. Two important classes of problems in this context relate to task release policies stabilizing the queues [18]–[21] and optimal time or attention allocation policies for human operators [22]. However, in these papers the nature of the tasks is abstracted, so that the solutions proposed do not typically apply to collaborative decision-making problems.

In [12], a problem of allocating independent classification tasks between a human and an automated decision system is considered. The human's decision performance depends on her workload, whereas her willingness to follow the task allocation suggested by the DSS depends on her trust in the capability of the automation. A collaborative decision-making architecture closer to the situation considered in this paper is studied in [23], [24]. A mixed-initiative team consisting of a DSS and a human performs classification tasks by following a two-step strategy: for a given task the DSS first examines the data, then either immediately makes a classification decision or refers the task to the human operator. The human acts as a second classifier, whose performance is task dependent. Such an architecture allows the optimization of the overall system performance by focusing the limited cognitive resources of the human operator on the tasks that are most difficult for the DSS.

The human-automation decision-making architecture considered in this paper is similar to that of [23]. One important difference is that we consider a finite set of tasks available from the start, whereas [23] focuses on a generic task in a steady-state environment. As a result, our definition of workload is different: it corresponds to all the tasks actually handled by the operator in our set-up, whereas it can be more easily interpreted as a probability of task referral by the DSS in [23]. Moreover, the analysis we perform here is general and does not depend on assuming specific probabilistic models (e.g., Gaussian or binary observations), as in [23]. The decision-making algorithm for the DSS that we develop

This work was supported by the Canadian Department of National Defence under the Innovation for Defence Excellence and Security (IDEaS) Micro-Net program.

Kesav Kaza and Jerome Le Ny are with the Department of Electrical Engineering, Polytechnique Montreal, Canada. Aditya Mahajan is with the Department of Electrical and Computer Engineering, McGill University, Canada. The authors are also with GERAD, Montreal, Canada. {kesav.kaza,jerome.le-ny}@polymtl.ca, aditya.mahajan@mcgill.ca

here is also provably optimal, whereas [23] optimizes the parameters of a specific task referral heuristic, for specific classification error costs.

The rest of the paper is organized as follows. In Section II, we present the system model and problem formulation for decision referrals. In Section III, we present an algorithm to obtain the optimal policy for decision referrals. In Section IV, we compare the performance of the proposed algorithm with a few baseline allocation schemes. We conclude in Section V.

II. System Model and Problem Formulation

We consider a DSS consisting of an automation and a human performing binary classification on a batch of tasks. The batch consists of K independent identically distributed tasks and we use $\mathcal{K} := \{1, \dots, K\}$ to denote the entire batch. Each task $k \in \mathcal{K}$ has a binary state $H_k \in \{\mathcal{H}_0, \mathcal{H}_1\}$. The automation and the human receive an observation for each task. We use $Y_{1,k} \in \mathcal{Y}_1$ and $Y_{2,k} \in \mathcal{Y}_2$ to denote the observations received by the automation and the human, respectively, for task $k \in \mathcal{K}$. These observations are random variables which depend on the true state H_k.

The states $H_{1:K}$ are independent across tasks, the observations $Y_{1,1:K}$ and $Y_{2,1:K}$ are independent and identically distributed across tasks, and for every task $k \in \mathcal{K}$, the observations $Y_{1,k}$ and $Y_{2,k}$ are conditionally independent given H_k.

The DSS operates as follows. First, the automation sees the observations $Y_{1,1:K}$ of the entire batch and, for each task $k \in \mathcal{K}$, decides to either classify the task as $\mathcal{H}_0$ or $\mathcal{H}_1$ or to refer (i.e., transfer) the task to the human. Let $\mathcal{N} \subseteq \mathcal{K}$ denote the indices of the tasks referred to the human. The human then sees the observations $\{Y_{2,n}\}_{n\in\mathcal{N}}$ of the tasks referred to her and classifies each of the referred task n as $\mathcal{H}_0$ or $\mathcal{H}_1$.

A. Observation models

We assume that the automation has a static observation model $P_1 : \{\mathcal{H}_0, \mathcal{H}_1\} \to \Delta(\mathcal{Y}_1)$. In particular,

$$\mathbb{P}(Y_{1,1}, \dots, Y_{1,K}) = \prod_{k\in\mathcal{K}} \sum_{i\in\{0,1\}} \pi_i P_1(Y_{1,k}|\mathcal{H}_i)$$

where $\pi_i = \mathbb{P}(H_k = \mathcal{H}_i)$ denotes the prior on the state of task $k \in \mathcal{K}$, and is same for all tasks.

In contrast, the human observation model depends on his workload, which is defined as the fraction $w = |\mathcal{N}|/|\mathcal{K}| \in [0, 1]$ of tasks referred to her by the automation. In particular, we assume that there is an observation model $P_2 : \{\mathcal{H}_0, \mathcal{H}_1\} \times [0, 1] \to \Delta(\mathcal{Y}_2)$ such that

$$\mathbb{P}(\{Y_{2,n}\}_{n\in\mathcal{N}}) = \prod_{n\in\mathcal{N}} \sum_{i\in\{0,1\}} \pi_i P_2(Y_{2,n}|\mathcal{H}_i, w)$$

where $w = |\mathcal{N}|/|\mathcal{K}|$ is the workload (or taskload). The workload dependent observation model can capture a degradation of observation performance for the human as the workload increases, since the operator must dedicate less time or cognitive resources to each individual task.

We now present two examples of observation models for the automation and the human.

Example 1: Suppose $\mathcal{H}_0 = 0$ and $\mathcal{H}_1 = d_0$. The observations of the automation are given by

$$Y_{1,k} = H_k + N_{1,k}, \quad k \in \mathcal{K}, \tag{1}$$

where $N_{1,1:K}$ is an independent Gaussian process, independent of $H_{1:K}$, with $N_{1,1:K} \sim \mathcal{N}(0, \sigma_1^2)$. In contrast, the observations of the human are given by

$$Y_{2,n} = H_n + N_{2,n}, \quad n \in \mathcal{N}, \tag{2}$$

where $\{N_{2,n}\}_{n\in\mathcal{N}}$ is also an independent Gaussian process, which is independent of $H_{1:K}$ as well as $N_{1,1:K}$. To capture the performance degradation of the human with workload, we assume that for some σ_2 such that $\sigma_2^2 \le \sigma_1^2 < 2\sigma_2^2$,

$$N_{2,n} \sim \mathcal{N}(0, (1 + w)\sigma_2^2), \quad n \in \mathcal{N}.$$

The above model has two salient features. First, the observation noise of the human increases as a function of workload. Second, under low workload, the observation noise of the human is no worse than that of the automation. But at high workload, the automation has lower observation noise than the human. Thus, it is not universally better to allocate all the tasks to the automation or the human.

Example 2: As in Example 1, we consider $\mathcal{H}_0 = 0$, $\mathcal{H}_1 = 1$, and the observation model of the automation is the same as in Example 1. However, the observation model of the human is now given as follows. For any $n \in \mathcal{N}$,

$$Y_{2,n}|\{H_n = \mathcal{H}_0\} \sim \mathcal{N}(0, \sigma_2^2)$$
$$Y_{2,n}|\{H_n = \mathcal{H}_1\} \sim \mathcal{N}(d_0(1 - w), \sigma_2^2).$$

As in Example 1, in Example 2 it becomes harder for the human to differentiate between the two hypothesis as the workload increases, but the actual mechanism is different.

B. Cost and Performance

Let D_k denote the final classification decision on task $k \in \mathcal{K}$ (made either by the automation or the human). We assume that for each task $k \in \mathcal{K}$ the system incurs a cost $\bar{C}(D_k, H_k)$ where

$$\bar{C}(D_k, H_k) = \begin{cases} c_{tp} & \text{if } (H_k, D_k) = (\mathcal{H}_1, \mathcal{H}_1), \\ c_{fp} & \text{if } (H_k, D_k) = (\mathcal{H}_0, \mathcal{H}_1), \\ c_{tn} & \text{if } (H_k, D_k) = (\mathcal{H}_0, \mathcal{H}_0), \\ c_{fn} & \text{if } (H_k, D_k) = (\mathcal{H}_1, \mathcal{H}_0). \end{cases} \tag{3}$$

In addition, the system incurs a cost c_m for each task referred to the human.

Thus, if the automation decides to refer the set $\mathcal{N} \subseteq \mathcal{K}$ of tasks to the human and makes a classification decision D_k on tasks $k \in \mathcal{K} \setminus \mathcal{N}$ and the human makes a classification decision D_n on tasks $n \in \mathcal{N}$ referred to it, then the system incurs a cost (from the point of view of the DSS) given by

$$J(D_{1:K}, \mathcal{N}, Y_{1,1:K}) = \sum_{k\in\mathcal{K}\setminus\mathcal{N}} \sum_{i\in\{0,1\}} p^1_{i,k} \bar{C}(D_k, H_i) + |\mathcal{N}|c_m + \sum_{n\in\mathcal{N}} \sum_{i\in\{0,1\}} p^1_{i,n} \bar{C}(D_n, H_i) \tag{4}$$

where $p^1_{i,k}$ is the posterior on the state H_k given the observation $Y_{1,k}$, *i.e.*,

$$p^1_{i,k} = \mathbb{P}(H_k = \mathcal{H}_i | Y_{1,k}), \quad i \in \{0,1\}, k \in \mathcal{K}.$$

Note that the automation does not know the decisions $\{D_n\}_{n\in\mathcal{N}}$ made by the human and needs to form a posterior belief on them. For that matter, we introduce a model for the decisions made by the human.

C. Human Decision Model

We impose the following assumption on the decision making of the human.

Assumption 1: For each task $n \in \mathcal{N}$, the human decides between $\mathcal{H}_0$ and $\mathcal{H}_1$ based only on the observation $Y_{2,n}$. In particular, the human does not have access to the automation's observation $Y_{1,n}$. The human also does not account for the fact that the automation referred the task to her after observing the batch $Y_{1,1:K}$.

Assumption 1 can be justified based for example on the limited time that the operator has to make a decision. Note also that even in situations where the raw data received for classification by the automation and the human could be the same (for example, a picture or a text message), $Y_{1,n}$ and $Y_{2,n}$ would still typically differ. These random variables generally represent the high-level features detected in the raw data by the automation's signal processing pipeline or the human's cognitive process respectively, on which the final decision would be based. Hence, it is reasonable to assume that the human does not have access to or means to interpret the features $Y_{1,n}$ detected by the automation.

Let $P_{2,tp}(w)$ and $P_{2,fp}(w)$ denote the false and true positive probabilities of the human's decision when operating at a workload of w, i.e.,

$$P_{2,tp}(w) = \mathbb{P}(D_{2,n} = \mathcal{H}_1 | \mathcal{H}_n = \mathcal{H}_1, w), \quad \forall n \in \mathcal{N}, \quad (5a)$$
$$P_{2,fp}(w) = \mathbb{P}(D_{2,n} = \mathcal{H}_1 | \mathcal{H}_n = \mathcal{H}_0, w), \quad \forall n \in \mathcal{N}. \quad (5b)$$

In practice, these probabilities can be obtained through preliminary calibration experiments with the human operator [12], [25]. Alternatively, we discuss below through two examples how these functions can be obtained from first-principle reasoning.

1) Threshold-Based Classification Rule: Assume $\mathcal{H}_0 < \mathcal{H}_1$ are scalar values and that the human uses a threshold-based decision rule to make a decision, i.e., for every workload $w \in [0,1]$ there exists a threshold $\tau(w)$ such that for task $n \in \mathcal{N}$, the human's decision with observation $Y_{2,n}$ is given by

$$D_n = \begin{cases} \mathcal{H}_0, & \text{if } Y_{2,n} < \tau(w), \\ \mathcal{H}_1, & \text{if } Y_{2,n} \geq \tau(w). \end{cases} \quad (6)$$

Then, the true and false positive probabilities are given by

$$P_{2,tp}(w) = \mathbb{P}(Y_{2,n} \geq \tau(w) | H_n = \mathcal{H}_1),$$
$$P_{2,fp}(w) = \mathbb{P}(Y_{2,n} \geq \tau(w) | H_n = \mathcal{H}_0),$$

which can be computed by the automation from the knowledge of the threshold τ and the observation model $\mathbb{P}(Y_{2,n}|H_n)$. We present one such example below.

Example 3: Consider the observation models of Example 1 and 2. Let $\mathcal{Q}(x) := \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-z^2/2} dz$ denote the tail distribution function of the standard normal distribution. For the model of Example 1,

$$P_{2,fp}(w) = \mathcal{Q}\left(\frac{\tau(w)}{\sigma_2\sqrt{1+w}}\right), P_{2,tp}(w) = \mathcal{Q}\left(\frac{\tau(w)-d_0}{\sigma_2\sqrt{1+w}}\right).$$

For the model of Example 2,

$$P_{2,fp}(w) = \mathcal{Q}\left(\frac{\tau(w)}{\sigma_2}\right), P_{2,tp}(w) = \mathcal{Q}\left(\frac{\tau(w)-d(w)}{\sigma_2}\right).$$

A threshold-based rule does not have to be implemented on $Y_{2,n}$ however. For example, if the human behaved as an optimal Bayes classifier, she would choose between $\mathcal{H}_0$ and $\mathcal{H}_1$ by using a standard Bayes likelihood ratio test [26]. One way to write the resulting decision rule is

$$D_n = \begin{cases} \mathcal{H}_0, & \text{if } p^2_{1,n}(w) < \rho_{th}, \\ \mathcal{H}_1, & \text{if } p^2_{1,n}(w) \geq \rho_{th}, \end{cases} \quad (7)$$

where $p^2_{1,n}(w)$ denotes the posterior probability of task n having state $\mathcal{H}_1$ given the observation $Y_{2,n}$ and workload w, *i.e.*, $p^2_{i,n}(w) := \mathbb{P}(H_n = \mathcal{H}_i | Y_{2,n}, w)$ and

$$\rho_{th} = \frac{c_{fp} - c_{tn}}{c_{fp} - c_{tn} + c_{fn} - c_{tp}},$$

where c_{tp}, c_{fp} and so on are the classification decision costs given by (3). As a result, we can write the probabilities (5) as

$$P_{2,tp}(w) = \mathbb{P}(p^2_{1,n}(w) \geq \rho_{th} | H_k = \mathcal{H}_1),$$
$$P^2_{2,fp}(w) = \mathbb{P}(p^2_{1,n}(w) \geq \rho_{th} | H_k = \mathcal{H}_0).$$

For example, in the observation model of Example 1, standard calculations show that the decision rule (7) can be rewritten as a threshold-based rule (6) on $Y_{2,k}$ with

$$\tau(w) = \frac{d_0}{2} + \frac{(1+w)\sigma_2^2}{d_0} \ln\left(\frac{(c_{fp} - c_{tn})\pi_0}{(c_{fn} - c_{tp})\pi_1}\right).$$

and hence $P_{2,fp}$ and $P_{2,tp}$ can be expressed again in terms of the $\mathcal{Q}$ function.

2) Softmax Decision Rule: Human decision making under uncertainty need not be Bayesian or deterministic with a fixed classification threshold. A possible model discussed in [13], [27] is the softmax rule, which is probabilistic and depends on the relative expected costs of various actions. This decision rule for the human is written as follows. For task $n \in \mathcal{N}$, define

$$P_n(w) = \mathcal{S}\left(p^2_{1,n}(w) c_{tp} + p^2_{0,n}(w) c_{fp}\right),$$

with $\mathcal{S}$ the sigmoid function $\mathcal{S}(x) = \frac{1}{1+e^{-x}}$. Then, the decision of the human is

$$D_n = \begin{cases} \mathcal{H}_1, & \text{w.p. } P_n(w), \\ \mathcal{H}_0, & \text{w.p. } 1 - P_n(w). \end{cases}$$

Again, it is possible to compute the functions in (5) from this rule and an observation model for $Y_{2,n}|H_n$.

D. The optimization problem

The DSS does not need to know the complete model according to which the human makes the decision; rather it simply needs to know the probabilities $P_{2,tp}(w)$ and $P_{2,fn}(w)$ given by (5). Given the probabilities, the expected performance (averaged over the decisions made by the human) is given by

$$\begin{aligned}\bar{J}(\mathcal{N}, \{D_k\}_{k\in\mathcal{K}\setminus\mathcal{N}}, \{p^1_{i,1:K}\}_{i\in\{0,1\}}) \\ = \sum_{k\in\mathcal{K}\setminus\mathcal{N}} \sum_{i\in\{0,1\}} p^1_{i,k}\bar{C}(D_k, H_i) \\ + |\mathcal{N}|c_m + \Gamma_2(\mathcal{N}, |\mathcal{N}|/|\mathcal{K}|),\end{aligned} \tag{8}$$

where

$$\begin{aligned}\Gamma_2(\mathcal{N}, w) = \sum_{n\in\mathcal{N}} \Big(p^1_{1,n}[P_{2,tp}(w)c_{tp} + (1 - P_{2,tp}(w))c_{fn}] \\ + p^1_{0,n}[P_{2,fp}(w)c_{fp} + (1 - P_{2,fp}(w))c_{tn}]\Big).\end{aligned} \tag{9}$$

We are interested in the following optimization problem.

Problem 1: Given the posterior beliefs $\{p^1_{i,k}\}_{k\in\mathcal{K}}, i \in \{0,1\}$ of the automation, and the decision distribution $P_{2,tp}; P_{2,fp} : [0,1] \to [0,1]$ of the human, determine $\mathcal{N}$ and $\{D_k\}_{\mathcal{K}\setminus\mathcal{N}}$ so as to maximize $\bar{J}(\mathcal{N}, \{D_k\}_{k\in\mathcal{K}\setminus\mathcal{N}}, \{p^1_{i,1:K}\}_{i\in\{0,1\}})$ given by (8).

III. Main Results

In this section, we present an algorithm to find the optimal solution to Problem 1. The main idea of our algorithm is as follows. In Section III-A we present an efficient solution for the sub-problem where the constraint $|\mathcal{N}|$ to a pre-specified value. Then, in Section III-B, we present an algorithm that searches for the optimal solution by iterating over $|\mathcal{N}| \in \{0, 1, \ldots, K\}$.

A. Optimal Allocation under Fixed Workload

Suppose the workload w (or equivalently, the number $|\mathcal{N}|$ of tasks to be referred) is pre-specified, and the DSS has to choose which $|\mathcal{N}|$ tasks to refer to the human. For ease of notation, let $p^1_k = [p^1_{0,k}, p^1_{1,k}]$, and define

$$\bar{C}_1(D_k, p^1_k) = \sum_{\{0,1\}} p^1_{0,k}\bar{C}(D_k, H_i) \tag{10}$$

and

$$\begin{aligned}\bar{\Gamma}_2(p^1_k, w) = p^1_{1,k}\left[P_{2,tp}(w)c_{tp} + (1 - P_{2,tp}(w))c_{fn}\right] \\ + p^1_{0,k}\left[P_{2,fp}(w)c_{fp} + (1 - P_{2,fp}(w))c_{tn}\right].\end{aligned} \tag{11}$$

First, observe that if a task k is not referred to the human, the the automation Bayes optimal decision D_k for task k is given by

$$D_k = \begin{cases}\mathcal{H}_0, & \text{if } \bar{C}_1(p^1_k, \mathcal{H}_0) \le \bar{C}_1(p^1_k, \mathcal{H}_1), \\ \mathcal{H}_1, & \text{otherwise.}\end{cases} \tag{12}$$

and the expected cost for task k is

$$\bar{C}^*_1(p^1_k) = \min\{\bar{C}_1(p^1_k, \mathcal{H}_0), \bar{C}_1(p^1_k, \mathcal{H}_1)\}. \tag{13}$$

Thus,

$$\begin{aligned}&\bar{J}(\mathcal{N}, \{D_k\}_{k\in\mathcal{K}\setminus\mathcal{N}}, p^1_{1:K}) \\ &\ge \sum_{k\in\mathcal{K}\setminus\mathcal{N}} \bar{C}^*_1(p^1_k) + |\mathcal{N}|c_m + \sum_{n\in\mathcal{N}} \bar{\Gamma}_2(p^1_n, w) \\ &= \sum_{k\in\mathcal{K}} \bar{C}^*_1(p^1_k) - \sum_{n\in\mathcal{N}} G(p^1_n, |\mathcal{N}|/|\mathcal{K}|),\end{aligned} \tag{14}$$

where

$$G(p^1_k, w) = \bar{C}^*_1(p^1_k) - \bar{\Gamma}_2(p^1_k, w) - c_m. \tag{15}$$

Note, the equality is achieved when the decisions $\{D_k\}_{k\in\mathcal{K}\setminus\mathcal{N}}$ are chosen according to (12). Thus, for a fixed $|\mathcal{N}|$, $\bar{J}(\mathcal{N}, \{D_k\}_{k\in\mathcal{K}\setminus\mathcal{N}}, p^1_{1:K})$ is minimized when $\mathcal{N}$ is chosen to maximize

$$\bar{G}(\mathcal{N}) := \sum_{n\in\mathcal{N}} G(p^1_n, |\mathcal{N}|/|\mathcal{K}|). \tag{16}$$

Here, $\bar{G}(\mathcal{N})$ is the total cost of 'offloading' the set of tasks $\mathcal{N}$ to the human. The total cost $\bar{J}(\mathcal{N}, \{D_k\}_{k\in\mathcal{K}\setminus\mathcal{N}}, p^1_{1:K})$ is minimized when $\mathcal{N}$ is the set of states with the $|\mathcal{N}|$ highest $G-$indices. This gives us the following result.

Lemma 1: For a pre-specified workload $w = |\mathcal{N}|/|\mathcal{K}|$, it is optimal to allocate the tasks with the highest $|\mathcal{N}|$ $G-$indices given by (15).

Proof: The proof follows immediately from the discussion. ■

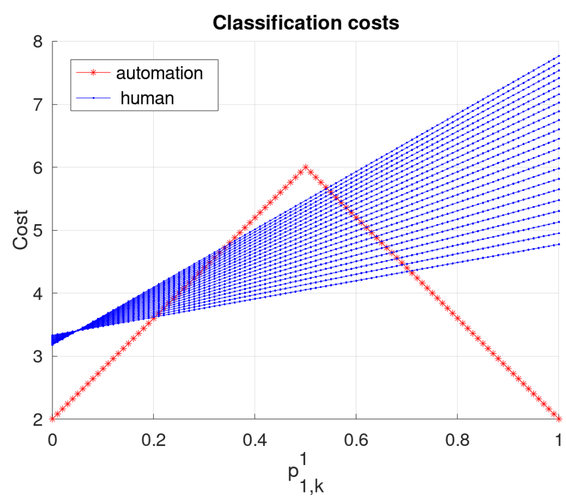

Fig. 1. *Numerical Example* (with $c_{tp} = c_{tn}$ and $c_{fp} = c_{fn}$): The red hill is the classification cost of the automation, $\bar{C}^*_1(p^1_k)$, as a function of posterior probability $p^1_{1,k}$ of hypothesis $\mathcal{H}_1$. The blue lines show the expected classification cost for the human, $\bar{\Gamma}_2(p^1_k, w)$, $w \in \{1/K, ..., K/K\}$. Batch size $K = 20$. The cost reduction for offloading is $G(p^1_k, w)$, which is the difference between the red and blue functions.

Fig. 1 can be used to illustrate the idea of $G-$index. Each task is a point $p^1_{1,k} \in [0,1]$, and the $G-$index is difference between the automation's classification cost (in red) and the human classification cost (in blue) which depends on workload. As the workload increases the cost of human decisions also goes up (slope increases). The total cost would be reduced if tasks with greater $G-$index values are referred to the human.

B. Optimal Workload Distribution

The task allocation strategy of Lemma 1 is optimal for a pre-specified workload w, but is not optimal in general. However, since there are only $K+1$ possible values of workload w, we can simply do a brute force search over all of them to find the optimal w. As argued above, minimizing the total expected cost is equivalent to minimizing

$$\bar{G}^*(w) = \min_{\mathcal{N}:|\mathcal{N}|} = wK, \tag{17}$$

over $w \in \{0, 1/K, \ldots, 1\}$. Thus, the optimal workload w can be identified by evaluating $G^*(w)$ for all choices of w. The complete algorithm is shown in Algorithm 1.

Algorithm 1: Optimal task allocation — Exhaustive workload search

Input: Posterior probabilities of tasks $\{p^1_{1:K}\}$
Output: w^*, $[a_1^{w^*}, a_2^{w^*}, \ldots, a_K^{w^*}]$
for $w \in \{0, 1/K, 2/K, \ldots, 1\}$ **do**
  Initialize $a_k^w = 0, \forall k = 1, 2, \ldots, K$.
  Compute $\left[G(p_k^1, w)\right]_{k\in\mathcal{K}}$.
  Indices $[i_1, ..., i_K] \leftarrow$ Sort $\Downarrow \left([G(p_k^1, w)]_{k\in\mathcal{K}}\right)$.
  Allocate first wK tasks to human:
  $a_{i_1}^w = 1, \ldots, a_{i_{wK}}^w = 1$.
  Compute total cost reduction for workload w:
  $\bar{G}(w) := \sum_{k=i_1}^{i_{wK}} G(p_k^1, w)$.
end
$w^* \leftarrow \arg\max \bar{G}(w)$.
return w^*; $[a_1^{w^*}, a_2^{w^*}, \ldots, a_K^{w^*}]$.

Theorem 1: For an arbitrary batch of K tasks with corresponding posterior probabilities $[p^1_{1:K}]$, Algorithm 1 gives the optimal workload w^* and task allocation $\left[a_k^{w^*}\right]_{k\in\mathcal{K}}$.

Proof: From Lemma 1 we see that for fixed workload w, the optimal allocation has wK tasks assigned to the human. As the first term $\sum_{k=1}^K C_1^*(p_k^1)$ in (14) remains constant, the workload value maximizing the term $\sum_{n\in\mathcal{N}} \bar{\Gamma}_2(p_n^1) + c_m$ over the entire batch minimizes the cost. This value w^* is found by Algorithm 1 by exhaustive search over all possible workload values. ■

To close this section, we consider a special situation where the optimal decision rule for the automation has an intuitive form.

Corollary 1: If the human behaves as an optimal Bayes classifier, the priors are uniform $\mathbb{P}(H_k = \mathcal{H}_1) = \mathbb{P}(H_k = \mathcal{H}_0)$ and the decision costs satisfy $c_{tp} = c_{tn} = 0$ and $c_{fn} = c_{fp}$, it is optimal for a given workload value w to allocate wK tasks to the human with the lowest values of $\left|p^1_{1,k} - \frac{1}{2}\right|, k \in \mathcal{K}$.

IV. Numerical Simulations

In this section we present a simulation experiment to compare the performance of the proposed algorithms with other baseline allocation schemes.

We consider the observation model of Example 2 with $d_0 = 3$ and $\pi = [0.8, 0.2]$. The other parameters are chosen randomly according to $\sigma_1 \sim U(1.5, 2), \sigma_2 \sim U(1, 1.5)$, $\{c_{fp}, c_{fn}\} \sim U(8, 12)$, $\{c_{tp}, c_{tn}\} \sim U(0, 2)$ and $c_m \sim U(0, 0.5)$. Here, $U(a, b)$ denotes uniform distribution over the interval $[a, b]$. We compare the performance of the optimal policy with the following baselines:

- **Blind allocation (BA)**, which decides on a workload w^*_{ba} before seeing the batch $Y_{1,1:K}$ and refers $w^*_{ba}|\mathcal{K}|$ tasks to the human at random. The choice of w^*_{ba} in this case is given by

$$w^*_{ba} = \arg\min_{w\in\mathcal{W}} \{(1-w)E_1 + wE_2(w)\},$$

where

$$\begin{aligned} E_1 &= \pi_1(P_{1,tp}c_{tp} + (1 - P_{1,tp})c_{fn}) \\ &\quad + \pi_0(P_{1,fp}c_{fp} + (1 - P_{1,fp})c_{tn}), \\ E_2(w) &= c_m + \pi_1(P_{2,tp}(w)c_{tp} + (1 - P_{2,tp}(w))c_{fn}) \\ &\quad + \pi_0(P_{2,fp}(w)c_{fp} + (1 - P_{2,fp}(w))c_{tn}). \end{aligned}$$

- **Static allocation (SA)**, which decides on the workload w^*_{sa} before seeing the batch $Y_{1,1:K}$, but then refers $w^*_{sa}|\mathcal{K}|$ tasks to the human according to Lemma 1. The choice of w^*_{sa} in this case is given by

$$\begin{aligned} w^*_{sa} &= \arg\min_{w\in\mathcal{W}} \mathbb{E}_{p^1_{1:K}} \left[\min_{[a_k^w]_{k\in\mathcal{K}}} \sum_{k=1}^K C(p_k^1, a_k^w) \right], \quad (18) \\ C(p^1_{1,k}, a_k^w) &= (1 - a_k^w)C_1(p^1_{1,k}) + a_k^w \bar{\Gamma}_2(p^1_{1,k}, w), \end{aligned}$$

where allocation vector $[a_k^w]_{k\in\mathcal{K}}$ is for workload w with $\sum_{k=1}^K a_k^w = wK$.

We choose 25 values of $\{\sigma_1, \sigma_2, c_{fp}, c_{fn}, \{c_{tp}, c_{tn}, c_m\}$ at random as described above. For each choice of these parameters, we generate 2000 batches of size $K = 20$. The expected performance as well as the standard deviation of performance for each problem instance is shown in Fig. 2. The workload allocation to humans under the different policies is shown in Fig. 3.

The results show that decision referral has better performance and less variance (about 17% less cost and 3% less standard deviation) than blind allocation. It is also interesting to note that across the different instances, blind allocation has a higher variation of human workload than decision referrals. Finally, static allocation performs almost as well as the optimal allocation in almost all cases, suggesting that it may be a useful practical alternative.

Fig. 3 shows the workload allocations to the human under various policies. Blind allocation policy assigns either very low or very high workloads to the human.

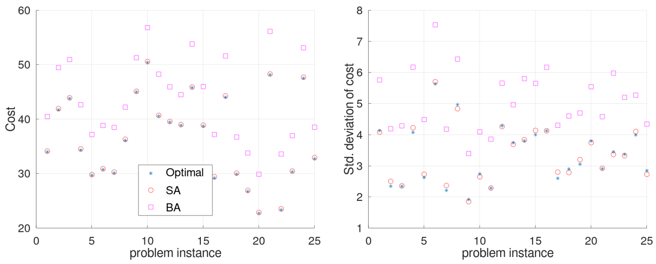

Fig. 2. Comparison of various policies for 25 distinct problem instances, for batch size $K = 20$. *[left]* Average cost *[right]* Standard deviations of costs

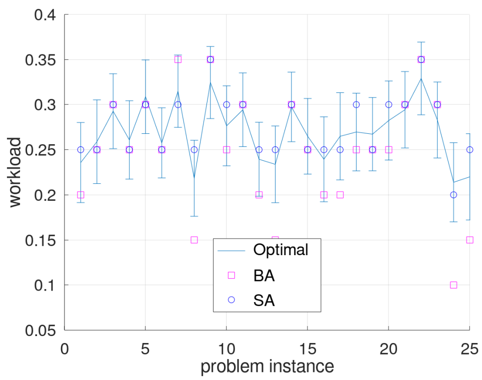

Fig. 3. Average workload allotted to human by various policies, over 25 distinct problem instances, for batch size $K = 20$.

V. Conclusion

In this paper a decision referral problem is formulated for a human-automation team jointly performing binary classification tasks. The decision support system needs to decide which tasks should be referred to the human for final classification decisions, after performing a first analysis of the data. An algorithm for finding the optimal referral decisions is presented. The proposed decision model only requires the true and false positive rates of the 'human classifier' as a function of time, and does not need any other information about human decision making process. Numerical simulations illustrate the benefits of the informed allocation policies over static blind task allocation scheme. Simulations also suggest that informed allocation heuristics which are close to optimal can be devised and employed based on convenience of implementation.

In the future we plan to validate the proposed model through experiments with human participants.

References

[1] S. D. Ramchurn, J. E. Fischer, Y. Ikuno, F. Wu, J. Flann, and A. Waldock, "A study of human-agent collaboration for multi-uav task allocation in dynamic environments," in *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2015, pp. 1184–1192.

[2] J. Gao and J. D. Lee, "Extending the decision field theory to model operators' reliance on automation in supervisory control situations," *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 36, no. 5, pp. 943–959, 2006.

[3] S. Music and S. Hirche, "Control sharing in human-robot team interaction," *Annual Reviews in Control*, vol. 44, pp. 342–354, 2017.

[4] Y. Wang, "Regret-based automated decision-making aids for domain search tasks using human-agent collaborative teams," *IEEE Transactions on Control Systems Technology*, vol. 24, no. 5, pp. 1680–1695, 2016.

[5] V. Riley, "A general model of mixed-initiative human-machine systems," in *Proceedings of the Human Factors Society Annual Meeting*, vol. 33, no. 2. SAGE Publications Sage CA: Los Angeles, CA, 1989, pp. 124–128.

[6] M. L. Cummings and S. Bruni, *Collaborative Human-Automation Decision Making*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 437–447.

[7] M. C. Gombolay, R. A. Gutierrez, S. G. Clarke, G. F. Sturla, and J. A. Shah, "Decision-making authority, team efficiency and human worker satisfaction in mixed human–robot teams," *Autonomous Robots*, vol. 39, no. 3, pp. 293–312, 2015.

[8] J. N. Tsitsiklis, "Decentralized detection," in *Advances in Statistical Signal Processing*, 1993, pp. 297—344.

[9] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential analysis: Hypothesis testing and changepoint detection*. CRC Press, 2014.

[10] A. Raghavan and J. S. Baras, "Binary hypothesis testing by two collaborating observers: A fresh look," in *2019 27th Mediterranean Conference on Control and Automation (MED)*. IEEE, 2019, pp. 244–249.

[11] B. Hyun, M. Faied, P. Kabamba, and A. Girard, "Mixed-initiative nested classification by optimal thresholding," in *2011 50th IEEE Conference on Decision and Control and European Control Conference*, 2011, pp. 7653–7658.

[12] C. Dubois and J. Le Ny, "Adaptive task allocation in human-machine teams with trust and workload cognitive models," in *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2020, pp. 3241–3246.

[13] M. Chen, S. Nikolaidis, H. Soh, D. Hsu, and S. Srinivasa, "Planning with trust for human-robot collaboration," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018, pp. 307–315.

[14] C. D. Wickens, J. G. Hollands, S. Banbury, and R. Parasuraman, *Engineering psychology and human performance*. Psychology Press, 2015.

[15] M. K. Tulga and T. B. Sheridan, "Dynamic decisions and work load in multitask supervisory control," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 10, no. 5, pp. 217–232, 1980.

[16] M. M. Cummings, "Man versus machine or man + machine?" *IEEE Intelligent Systems*, vol. 29, no. 5, pp. 62–69, 2014.

[17] V. Jog, R. J. La, and N. C. Martins, "Channels, remote estimation and queueing systems with a utilization-dependent component: A unifying survey of recent results," *CoRR*, vol. abs/1905.04362, 2021. [Online]. Available: http://arxiv.org/abs/1905.04362v2

[18] K. Savla and E. Frazzoli, "Maximally stabilizing task release control policy for a dynamical queue," *IEEE Transactions on Automatic Control*, vol. 5, no. 11, pp. 2655–2660, 2010.

[19] ——, "A dynamical queue approach to intelligent task management for human operators," *Proceedings of the IEEE*, vol. 100, no. 3, pp. 672–686, 2012.

[20] V. Srivastava, A. Surana, and F. Bullo, "Adaptive attention allocation in human-robot systems," in *2012 American Control Conference (ACC)*. IEEE, 2012, pp. 2767–2774.

[21] M. Lin, R. J. La, and N. C. Martins, "Stabilizing a queue subject to action-dependent server performance," *CoRR*, vol. abs/1903.00135v3, 2021. [Online]. Available: https://arxiv.org/abs/1903.00135v3

[22] V. Srivastava, R. Carli, C. Langbort, and F. Bullo, "Task release control for decision making queues," in *Proceedings of the 2011 American Control Conference*, 2011, pp. 1855–1860.

[23] B. Hyun, P. Kabamba, and A. Girard, "Optimal classification by mixed-initiative nested thresholding," *IEEE Transactions on Cybernetics*, vol. 45, no. 1, pp. 29–39, Jan 2015.

[24] Y. Chitalia, W. Zhang, B. Hyun, and A. Girard, "A revisit-based mixed-initiative nested classification scheme for unmanned aerial vehicles," in *2014 American Control Conference*. IEEE, 2014, pp. 1793–1798.

[25] T. D. Wickens, *Elementary signal detection theory*. Oxford University Press, USA, 2002.

[26] H. V. Poor, *An Introduction to Signal Detection and Estimation (2nd Ed.)*. Berlin, Heidelberg: Springer-Verlag, 1994.

[27] N. D. Daw, J. O'doherty, P. Dayan, B. Seymour, and R. Dolan, "Cortical substrates for exploratory decisions in humans," *Nature*, vol. 441, no. 7095, pp. 876–879, 2006.