Randomized Edge-Assisted On-Sensor Information Selection for Bandwidth-Constrained Systems

Igor Burago and Marco Levorato
Department of Computer Science
University of California, Irvine
{iburago, levorato}@uci.edu

***Abstract*—The problem of intelligent information selection in the Internet-of-Things systems with limited computational and communication resources is studied. One distinctive property of such systems is the clash of the computational complexity of the desired selection procedure and the low throughput of the wireless links between the devices acquiring information (sensors) and processing it (edge and cloud computing servers). To adaptively resolve that conflict, we propose a stochastic optimization algorithm for edge-assisted online learning of the optimal on-sensor observation classification and transmission decision rules. Using the stochastic Lyapunov function method, we prove that the resulting adaptive procedure can be used to adjust the parameters of the two local decision rules to asymptotically satisfy the constraint on channel access probability and to minimize the expected classification error.**

***Index Terms*—Internet of Things, wireless networks, information selection, stochastic optimization, penalty function method.**

I. Introduction

Recent rise of Machine-to-Machine (M2M) communications, reinforced by the advances in machine learning methods, have been solidifying the growing interest for the introduction of distributed forms of intelligence into the Internet-of-Things (IoT) systems arising in a multitude of applications [1]–[5].

Despite the differences in their implementation, all of them have a principal trait making optimal organization of information processing fundamentally challenging. Within the devices participating in an IoT system, there are two sides that they tend to form based on their access to the *information* and the capacity for *intelligence*. One side consolidates devices whose primary function revolves around data acquisition; they are directly involved with the system's environment and can immediately observe its dynamics. The other side consists of nodes whose specialty is in data processing; having a wider view of the system, they are necessarily distanced from the data sources, but are capable of running the complex processing algorithms required for the system to autonomously adapt to the changing conditions of its environment.

The disconnect between these groups of nodes has its reflection at the network level, with the boundary between the two lying along the edge of the wireless networks comprising the system's periphery. The former devices operating inside those wireless islands are the *sensors*; the latter ones located at or beyond the edge are the edge and cloud processors, which we will simply call the *edge*, for short. The entities of the sensors and the edge embody the two components necessary for any sufficiently advanced decision-making: relevant data to drive decisions and adequately capable algorithms to realize them (see Figure 1). Finding a way for the latter to be timely and efficiently applied to the former constitutes the key question of the IoT system design

This work was supported in part by UCOP under Grant LFR-18-548175 and National Science Foundation under Grant IIS-1724331 and ECCS-1611349.

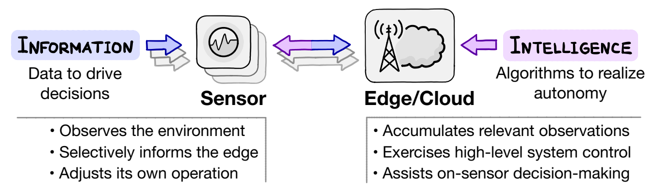

Figure 1. Principal conflict of optimal information processing in IoT systems.

Realization of this premise would be straightforward if not the major constraints dictated by the nature of the sensors' operation. In most IoT applications, sensors communicate with the rest of the system solely through the wireless link shared with other clients within their wireless island, and thus are limited in the amount of *bandwidth* they can use. Furthermore, sensors are additionally restricted in the complexity of *computation*, due to hardware and energy expenditure limitations (e.g., CPU voltage, battery capacity, etc.).

If both the communication and computation constraints are sufficiently relaxed, it is easy to organize the necessary processing by making the sensors communicate all their data in full to the edge or cloud servers where it will undergo the complete processing [6], [7]. If the communication constraint prevents that, but the sensors are powerful enough computationally, the processing may be offloaded to the sensors themselves, to the highest extent possible by the application [8]. If both constraints are tight enough, though, as is the case in many demanding IoT applications, neither approach is a viable solution. For some combinations of the constraints' severity, methods of compressed sensing and other compression-based approaches may provide a solution. However, being oblivious of the semantic consequences of an observation, these methods alone are often incapable of providing sufficient gains to satisfy both constraints without loss in decision quality.

Thus, there is a need for more intelligent information selection at the observation level, incorporating a deeper understanding of the observation's importance in the higher-level system control pipeline ran at the edge and in the cloud. In this paper, we propose one new approach to the edge-assisted on-sensor decision-making that can be learned on the fly and adapted in an online fashion, to make efficient data processing possible for constrained IoT systems.

II. Sensor-Edge Cooperation Problem

Let us consider an IoT system assigned with a typical task of monitoring the state of its environment in order to detect and react to anomalous situations. In other words, for each new observation z_t obtained by the sensors in time slot t, the system is to make a binary classification decision whether z_t is of interest for triggering subsequent higher-level control within the system (i.e., is a *positive*) or not (i.e., is a *negative*). We will assume that the edge processor is provided with some rather complex decision rule $\delta(z)$, which has been optimally trained for that purpose beforehand. Further, we will assume that, due to the computational constraint, the sensors are too weak to implement δ locally and, therefore, it can only be used at the edge. Finally, we will also assume that the bandwidth available to the sensors for communicating with the edge is insufficient for transmitting all of their observations. Thus, the sensors have the complete information, while the edge has the perfect intelligence, without them both having the ability to exchange their knowledge fully.

To resolve this contradiction, we propose for each sensor to have its own, simpler decision rule $\widehat{\delta}$ to approximate the inaccessible ground-truth δ. Ideally, the edge would like the two decisions be the same for any observation z, so that the limited channel would be used only for those z that do require a reaction from the system, with the rest being omitted to save the bandwidth. When the distribution of the sensor's observations is stationary and some initial period of the system's operation can be sacrificed for pre-training $\widehat{\delta}$, the sensor can subsequently use that decision for selecting data points to transmit. However, when $\widehat{\delta}$ has to be continuously tuned as the system gets new observations, a secondary rule to be used only for transmission decisions has to be introduced.

Indeed, imagine that decision $\widehat{\delta}_t$ happens to be (temporarily) suboptimal at some early time slot t, e.g., it correctly finds a small number of positives, but also rejects a lot of positives together with negatives. In doing so, it significantly hampers the representativeness of the sample reaching the edge, depriving it of the chance to correct this negative tendency. If, however, in addition to $\widehat{\delta}$, the sensor is provided with some other decision rule $\widetilde{\delta}$ that is allowed to deviate from the trajectory of the optimal fit of δ (within the limits of the communication constraint), then that destructive self-reinforcement loop can be broken by letting some number of extra negatives through for correcting $\widehat{\delta}$.

Thus, we propose the following workflow of the sensor-edge cooperation (see Figure 2). At every time slot t, the sensor has two decision rules $\widehat{\delta}_t$ and $\widetilde{\delta}_t$. As it acquires a new observation z_t, it computes the features $x_t = \chi(z_t)$ and evaluates $\widetilde{\delta}_t(x_t)$ to decide whether to transmit z_t to the edge. In parallel, if the logic of the sensor itself entails the need to trigger a time-sensitive action if z_t is, say, likely to be a positive, the sensor may evaluate $\widehat{\delta}_t(x_t)$ for that. We assume that those actions do not affect the next observation z_{t+1}.

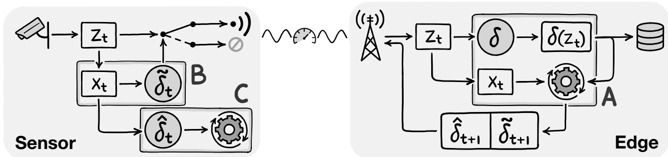

Figure 2. The operational diagram of the proposed sensor-edge workflow. **(A)** Edge-assisted update of the current on-sensor decision policies $\widehat{\delta}_t$ and $\widetilde{\delta}_t$ supervised by the reference decision rule δ. **(B)** On-sensor transmission filter of the observation stream based on the decision rule $\widetilde{\delta}_t$. **(C)** Optional on-sensor processing triggered by the approximate classification decision rule $\widehat{\delta}_t$.

At the edge, whenever a new observation z_t is received, it is first put through the reference decision $\delta(z_t)$ and, if declared positive, is stored for further processing in higher-level control. At the same time, that $\delta(z_t)$ is used as a supervisory decision to make a stochastic update of the policies $\widehat{\delta}_t$ and $\widetilde{\delta}_t$ currently installed at the sensor. Corrected decision rules $\widehat{\delta}_{t+1}$ and $\widetilde{\delta}_{t+1}$ are then communicated back to the sensor (in parameterized form), repeating the processing cycle.

More formally, let $Z_0 \subseteq Z$ and $Z_1 = Z \setminus Z_0$ be a partition of the abstract set of all possible observations Z, according to the reference decision rule δ. By some deterministic feature-extraction procedure $x = \chi(z)$, computable at both the sensor and the edge, observations from Z_0 and Z_1 are mapped to (possibly overlapping) subsets X_0 and X_1 within some feature space. The complete set X encompassing feature vectors for all z observable by the sensor is assumed to be bounded.

The decision rule $\widehat{\delta}$ for on-sensor classification approximates that partition of Z by fitting a separating surface breaking the feature space into $\widehat{X}_0$ and $\widehat{X}_1$ to match the ground-truth images X_0 and X_1 as close as possible (having, in the ideal case of best separation, $\widehat{X}_0 \cap X_1 = \widehat{X}_1 \cap X_0 = \varnothing$). Similarly, the decision rule $\widetilde{\delta}$ partitions the feature space into subsets $\widetilde{X}_{\mathrm{T}}$ and $\widetilde{X}_{\mathrm{W}}$ of observations to be transmitted by the sensor and to be withheld from transmission, respectively. Formally,

$$\widehat{\delta}(x):\ f(x,\theta) \underset{x\in\widehat{X}_0}{\overset{x\in\widehat{X}_1}{\gtrless}} \eta_1; \qquad \widetilde{\delta}(x):\ f(x,\theta) \underset{x\in\widetilde{X}_{\mathrm{W}}}{\overset{x\in\widetilde{X}_{\mathrm{T}}}{\gtrless}} \eta_0. \tag{1}$$

Here the function $f(x, \theta)$ defines the shape of a separating surface to be used, chosen from some parameterized family such that the limited computing resource of the sensor is sufficient for it to be computable for every observation. Of course, for the best possible quality, it is in our best interests to choose that family to be as complex as possible without violating the computational constraint. However, more complex functions generally require more parameters, increasing the amount of data to be sent from the edge to the sensor in the form of parameters. At some point, when the size of the parameter vector θ gets comparable to the mean size of a

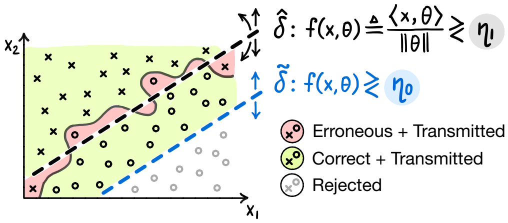

Figure 3. Implementation of the on-sensor decision rules for approximate classification ($\widehat{\delta}$) and transmission selection ($\widetilde{\delta}$) in an illustrative two-dimensional feature space, where crosses and circles depict feature vectors $x = \chi(z)$ of true positive ($z \in Z_1$) and negative ($z \in Z_0$) observations, respectively.

raw observation z, it will become more economical to resort to the straightforward scheme of consulting the edge for all observations. The communication factor is also the reason why it is reasonable for both decision rules $\widehat{\delta}$ and $\widetilde{\delta}$ to reuse the same parameters θ, as shown in (1). That is because the transmission filtering decision $\widetilde{\delta}$ does not have to be as precise as $\widehat{\delta}$. When adapted promptly, the threshold η_0 has enough resolution to keep the expected channel use bounded.

Thus, most efficient choices of f are bracketed by the two constraints and cannot be too flexible. For that reason, one universally practical choice for f that we can concentrate on, is the linear decision function

$$f(x,\theta) \triangleq \frac{\langle x,\, \theta\rangle}{\|\theta\|}, \tag{2}$$

where the value $f(x,\theta)$ can be interpreted as a signed distance from a point x to the separating hyperplane $\{f(x,\theta) = 0\}$, defined by the normal vector θ. Note that the fact of linearity of f in x does not eliminate the option of nonlinear surfaces completely, as the feature space can always be extended with extra coordinates being computed as some nonlinear transformations of the original features in x.

The resulting decision structure of $\widehat{\delta}$ and $\widetilde{\delta}$ in the space of features is shown in Figure 3. Given those decision rules, the sensor-edge system is faced with the problem of learning and updating the combined vector of parameters $\tau \triangleq \mathrm{vec}[\theta, \eta_0, \eta_1]$.

The main objective consists in achieving the minimum expected error, in which the role of the loss function for each separating surface is played by the distance from it to an erroneously classified observation in the feature space. Namely, we will work with the risk function

$$U(\tau) \triangleq \mathbb{E}\big[\hat{I}_{\eta_0}(u_\theta)\, I_{\mathrm{E}}(z,\tau)\, |u_\theta - \eta_0|\, |u_\theta - \eta_1|\big], \tag{3}$$

$$I_{\mathrm{E}}(z,\tau) \triangleq \big(1 - I_1(z)\big)\hat{I}_{\eta_1}(u_\theta) + I_1(z)\big(1 - \hat{I}_{\eta_1}(u_\theta)\big), \tag{4}$$

where $u_\theta \triangleq f(x,\theta)$, $\hat{I}_\eta(u) \triangleq \mathbb{1}[u > \eta]$, and $I_1(z) \triangleq \mathbb{1}[z \in Z_1]$, so that the indicator $\hat{I}_{\eta_0}(u_\theta)$ signals the event of an observation z with decision value u_θ being transmitted, while the combination of indicators in $I_{\mathrm{E}}(z,\tau)$ define an event of misclassification, i.e., the mismatch between the supervisory decision $I_1(z)$ produced by δ and its approximation $\hat{I}_{\eta_1}(u_\theta)$ produced by $\widehat{\delta}$. Noticing that

$$\hat{I}_{\eta_0}(u_\theta)\, |u_\theta - \eta_0| = \hat{I}_{\eta_0}(u_\theta)\, (u_\theta - \eta_0), \tag{5}$$

$$I_{\mathrm{E}}(z,\tau)\, |u_\theta - \eta_1| = \big(\hat{I}_{\eta_1}(u_\theta) - I_1(z)\big)(u_\theta - \eta_1), \tag{6}$$

we can conveniently rewrite the risk term U replacing the absolute values with to the corresponding signed differences:

$$U(\tau) = \mathbb{E}\big[\hat{I}_{\eta_0}(u_\theta)\big(\hat{I}_{\eta_1}(u_\theta) - I_1(z)\big)(u_\theta - \eta_0)(u_\theta - \eta_1)\big]. \tag{7}$$

For the formal definition of the communication constraint, we assume that, on average, the amount of bandwidth necessary to transmit any particular observation z is stable, so the bandwidth limit imposed on the sensor can be instead viewed as a constraint on its transmission frequency. That is, the edge is assumed to be aware of the highest probability φ_* corresponding to the maximum channel bandwidth that the sensor should utilize on average. Hence, we have the problem:

$$U(\tau) \longrightarrow \min_\tau, \quad \text{s.t.} \quad \mathbb{E}[\hat{I}_{\eta_0}(u_\theta)] \le \varphi_*. \tag{8}$$

To achieve both of these aims and obtain the guarantee of the asymptotic priority of the constraint, we define a combined criterion of the kind used by the penalty function method:

$$V(\tau,\zeta) \triangleq \frac{1}{\zeta}U(\tau) + W(\tau) \longrightarrow \min_\tau, \tag{9}$$

for a growing penalty coefficient ζ and the penalty term

$$W(\tau) \triangleq \frac{1}{2}\big(\mathbb{E}\big[\hat{I}_{\eta_0}(u_\theta)\big] - \varphi_*\big)^2. \tag{10}$$

We assume that the limit φ_* is set with some slack, so that every observation z_t, if chosen for transmission, successfully reaches the edge within the same time slot t. As a consequence, we define W to penalize not only for overusing, but also for underusing the channel, since not sending observations when there is bandwidth available may only worsen the promptness of the decision rules' adaptation.

Note that, unlike the penalty function method, we will not be solving a sequence of separate problems (9) with different values of ζ. Instead, we will optimize $V(\tau, \zeta_t)$ stochastically within the framework of a single problem, where the penalty coefficient ζ_t will be increasing at every time slot t. To construct an algorithm for this problem, we need to first tackle some nontrivial complications in estimating the gradients of the risk term $U(\tau)$ and, especially, the constraint term $W(\tau)$.

III. Bandwidth-Constraint Challenges

In order to introduce the techniques that will be necessary for constructing the stochastic gradient of the combined criterion (9), let us first consider an illustrative problem for the constraint term (10) alone, with the parameters θ being fixed:

$$W(\tau) = W(\eta_0) \longrightarrow \min_{\eta_0}. \tag{11}$$

Here the aim is to find an optimal transmission threshold η_0 for a certain fixed configuration θ of the separating surface. Note that this is not the same as the constraint from the original problem, and will not be used to achieve it as a separate goal.

Since the probability distribution of the feature vectors x_t is typically unknown and the value of the gradient $g(\eta_{0,t}) \triangleq \frac{d}{d\eta_0}W(\eta_0)$ cannot be computed exactly, it is only fitting to use a stochastic gradient descent algorithm for solving (11):

$$\eta_{0,t+1} = \eta_{0,t} - \gamma_t\, \widehat{g}(x_{t+1}, \eta_{0,t}), \tag{12}$$

where, in general, the stochastic gradient $\widehat{g}(x, \eta_0)$ defines a random variable estimating the true gradient $g(\eta_0)$ for a given observation x, so that, ideally,

$$\mathbb{E}[\widehat{g}(x, \eta_0)] = g(\eta_0) = \mathbb{E}[\hat{I}_{\eta_0}(u_\theta) - \varphi_*] \frac{d}{d\eta_0} \mathbb{E}[\hat{I}_{\eta_0}(u_\theta)]. \quad (13)$$

Construction of such estimator is not straightforward, though, and involves the following challenges.

A. Derivative of the Probability Function

Let us denote the factors comprising the gradient (13) as

$$\varphi(\eta_0) \triangleq \mathbb{E}[\hat{I}_{\eta_0}(u_\theta)] \quad \text{and} \quad \psi(\eta_0) \triangleq \frac{d}{d\eta_0}\varphi(\eta_0). \quad (14)$$

The probability $\varphi(\eta)$ is easy to estimate, yet it is not the case for its derivative $\psi(\eta)$. To make it possible, it is necessary to find a way to represent it as an expectation of a random variable that can be computed in a stochastic algorithm. Luckily, there are established results for differentiating integrals [9]–[13]that we can employ here. Below we state two of them.

Lemma 1. *Given*

$$F(\tau) \triangleq \int\limits_{\{x: f(x,\tau)>0\}} Q(x,\tau)\, p(x)\, dx, \quad (15)$$

where functions f, Q, and p are such that

1) *p is continuously differentiable with a compact support;*
2) *Q and f are continuously differentiable in both x and τ;*
3) *for any τ, $\nabla_x f(x,\tau) \neq 0$ for all x such that $f(x,\tau) = 0$;*

$$\nabla_\tau F(\tau) = \int\limits_{\{x: f(x,\tau)>0\}} \nabla_\tau Q(x,\tau)\, p(x)\, dx - \int\limits_{\{x: f(x,\tau)=0\}} \frac{\nabla_\tau f(x,\tau)}{\|\nabla_x f(x,\tau)\|} Q(x,\tau)\, p(x)\, d\sigma. \quad (16)$$

Lemma 2. *In the definitions and assumptions of Lemma 1,*

$$\nabla_\tau F(\tau) = \int\limits_{\{x: f(x,\tau)>0\}} \big(\nabla_\tau Q(x,\tau)\, p(x) - \operatorname{div}_x[\Lambda(x,\tau)\, p(x)]\big)\, dx, \quad (17)$$

where

$$\Lambda(x,\tau) \triangleq \frac{\nabla_\tau f(x,\tau)(\nabla_x f(x,\tau))^\intercal}{\|\nabla_x f(x,\tau)\|^2} Q(x,\tau), \quad (18)$$

and, for a matrix $H(x)$ of elements $h_{i,j}(x)$, $i = \overline{1,m}$, $j = \overline{1,n}$,

$$\operatorname{div}_x H(x) \triangleq \operatorname{vec}\left[\sum_{i=1}^n \frac{\partial h_{1,i}(x)}{\partial x_i}, \dots, \sum_{i=1}^n \frac{\partial h_{m,i}(x)}{\partial x_i}\right]. \quad (19)$$

Being applied to $\psi(\eta_0)$, Lemma 2 gives us the following:

$$\psi(\eta_0) = -\mathbb{E}\big[\hat{I}_{\eta_0}(u_\theta)\big(\operatorname{div}_x \Lambda(x,\tau) + \Lambda(x,\tau)\nabla_x[\log p(x)]\big)\big]. \quad (20)$$

Making this expectation stochastic would require from us to either know the p.d.f. $p(x)$ or be able to estimate it. Clearly, the former option is unavailable by the nature of the problem, and the latter is too expensive both in communication resources and time, as it necessitates a large body of samples to be collected at the edge. It is impractical even when parameters θ are fixed, and is effectively impossible when θ will be allowed to change on iterations (in the original problem from Section II).

Lemma 1 allows us to instead represent $\psi(\eta_0)$ as a surface integral that we can then approximate via a volume integral:

$$\psi(\eta_0) = -\int\limits_{\{x: f(x,\theta)=\eta_0\}} \frac{p(x)}{\|\nabla_x f(x,\theta)\|}\, d\sigma = -\lim_{\varepsilon\to 0} \frac{1}{\varepsilon} \int\limits_{S(\tau,\varepsilon)} p(x)\, dx, \quad (21)$$

that is taken over the ε-wide band along the separating surface:

$$S(\tau,\varepsilon) \triangleq \{x : \eta_0 < f(x,\theta) \le \eta_0 + \varepsilon\}. \quad (22)$$

B. Estimation of the Gradient Product

Thus, thanks to Lemma 1, we can now write single-observation stochastic estimators both for the transmission probability $\varphi(\eta_0)$ and its derivative $\psi(\eta_0)$:

$$\widehat{\varphi}(x,\eta_0) \triangleq \mathbb{1}\big[f(x,\theta) > \eta_0\big] = \hat{I}_{\eta_0}(u_\theta), \quad (23)$$

$$\widehat{\psi}(x,\eta_0,\varepsilon) \triangleq -\frac{1}{\varepsilon}\mathbb{1}\big[x \in S(\eta_0,\varepsilon)\big] = -\frac{1}{\varepsilon} I_{S(\tau,\varepsilon)}(u_\theta), \quad (24)$$

where (24) is constructed by direct translation from (21) (for the first time, an idea of this kind was proposed by Raik [9]).

Separately, both of these estimators are unbiased and asymptotically sound, assuming ε tends to zero. However, an estimator for the gradient $g(\eta_0)$, required by the algorithm (12), necessitates having an estimation of the product $\varphi(\eta_0)\psi(\eta_0)$. Obviously, $\widehat{\varphi}(x,\eta_0)\widehat{\psi}(x,\eta_0,\varepsilon)$ cannot serve as such, because, generally, it will violate condition (13).

One way to resolve this issue is to make observations x that go into $\widehat{\varphi}(x,\eta_0)$ and $\widehat{\psi}(x,\eta_0,\varepsilon)$ different and independent. For that, we can switch to viewing the stream of independent observations as a sequence of batches $\overline{z}_t = \{z_{t,i}\}_{i=1}^n$, each consisting of n subsequent data points. The batch $\overline{z}_t$, then, can be further split into two independent instrumental portions of $\widetilde{x}_t^\varphi$ and $\widetilde{x}_t^\psi$, containing the feature vectors for observations that will be used exclusively for estimating $\varphi(\eta_0)$ and $\psi(\eta_0)$, respectively. In the simplest case, it is sufficient to set $n = 2$ and choose $\widetilde{x}_t^\varphi = \{\chi(z_{t,1})\}$ and $\widetilde{x}_t^\psi = \{\chi(z_{t,2})\}$. In general, for $\widetilde{x}_t^\varphi = \{x_{t,i}^\varphi\}_{i=1}^{n_\varphi}$, $\widetilde{x}_t^\psi = \{x_{t,i}^\psi\}_{i=1}^{n_\psi}$ and $u_{t,i}^\varphi = f(x_{t,i}^\varphi,\theta)$, $u_{t,i}^\psi = f(x_{t,i}^\psi,\theta)$, estimators (23) and (24) turn into

$$\widehat{\varphi}(\widetilde{x}_t^\varphi,\eta_{0,t}) \triangleq \frac{1}{n_\varphi}\sum_{i=1}^{n_\varphi} \hat{I}_{\eta_0}(u_{t,i}^\varphi), \quad (25)$$

$$\widehat{\psi}(\widetilde{x}_t^\psi,\eta_{0,t},\varepsilon_t) \triangleq -\frac{1}{\varepsilon_t n_\psi}\sum_{i=1}^{n_\varphi} I_{S(\tau_t,\varepsilon_t)}(u_{t,i}^\psi). \quad (26)$$

Then, taking advantage of the instrumental sub-batches $\widetilde{x}_t^\varphi$ and $\widetilde{x}_t^\psi$ being independent, we can correctly use the product of the two estimators as an estimator for the gradient:

$$\widehat{g}(\widetilde{x}_t^\varphi,\widetilde{x}_t^\psi,\eta_{0,t},\varepsilon_t) = \big(\widehat{\varphi}(\widetilde{x}_t^\varphi,\eta_{0,t}) - \varphi_*\big)\,\widehat{\psi}(\widetilde{x}_t^\psi,\eta_{0,t},\varepsilon_t). \quad (27)$$

(In general, the requirement of the mutual independence of $\widetilde{x}_t^\varphi$ and $\widetilde{x}_t^\psi$ can be relaxed to the independence in pairs of $x_{t,i}^\varphi$ and $x_{t,i}^\psi$ — with $\widehat{g}$ changed accordingly.)

C. Mitigating Possibly Slow Convergence

Estimators (25), (26), (27), together with (12) form a working stochastic optimization procedure solving the illustrative problem (11). However, specifically for identifying the threshold η_0 in the original problem from Section II, the algorithm may not be adequately adaptive due to slow convergence of η_0.

Indeed, by definition, estimation (26) successfully converges to the exact gradient of the probability function $\varphi(\eta_0)$:

$$\lim_{t\to\infty} \widehat{\psi}(\widetilde{x}_t^{\psi}, \eta_{0,t}, \varepsilon_t) = \psi(\eta_0), \quad \text{if} \quad \lim_{t\to\infty} \varepsilon_t = 0. \tag{28}$$

To keep that guarantee, however, we have to steadily decrease ε_t on iterations. Theoretically this requirement is inconsequential, but practically it might be quite limiting, as it ties the frequency of parameter updates to the occurrences of $\widehat{x}_{t,i}^{\psi} \in S(\eta_{0,t}, \varepsilon_t)$, whose probability decreases as ε_t gets smaller. Lagging values of $\eta_{0,t}$ negatively affect representativeness of the sample reaching the edge, which, in turn, harms the quality of the decision rules produced by the edge for the sensor.

Hence, in order to prevent $\eta_{0,t}$ from getting stale between updates as ε_t reaches sufficiently small values, it is useful to switch the algorithm to a looser estimation $\widehat{\psi}$ for the purpose of breaking its direct dependence on ε_t. As long as it maintains the correct mean direction of parameter updates, the overall convergence will not be affected. For scalar parameters like η_0, it means that even the constant estimator of the correct sign

$$\widehat{\psi}(\widetilde{x}_t^{\psi}, \eta_{0,t}, \varepsilon_t) \triangleq \widehat{\psi}_0 = \text{const} < 0, \tag{29}$$

is sufficient, that is, the same algorithm (12) with stochastic gradient (27) in which (26) is replaced with (29) also converges to a solution of problem (11).

Below, in Algorithm 1, we use this technique with a constant $\widehat{\psi}$ for updating the transmission threshold η_0, but keep the more advanced ε-dependent estimator for updating the parameters θ of the separating surface.

IV. Stochastic Optimization Algorithm

Let us now return to the full scope of the original problem (9) for the combined criterion V. For achieving the goal of minimizing it, we propose the following stochastic procedure.

Algorithm 1. *Given current parameters $\tau_t = \mathrm{vec}[\theta_t, \eta_{1,t}, \eta_{0,t}]$ obtained after t time slots, collect a new batch of observations $\overline{z}_{t+1}$, extract sub-batches $\widetilde{x}_{t+1}^{\varphi}$, $\widetilde{x}_{t+1}^{\psi}$, and then compute:*

$$\begin{aligned}\tilde{\theta}_{t+1} = \theta_t - \gamma_{1,t}\, \widehat{G}_\theta(\overline{z}_{t+1}, \tau_t) \\ - \gamma_{2,t}\, \widehat{g}_\theta(\widetilde{x}_{t+1}^{\varphi}, \widetilde{x}_{t+1}^{\psi}, \tau_t, \varepsilon_t),\end{aligned} \tag{30}$$

$$\eta_{0,t+1} = \eta_{0,t} - \gamma_{3,t}\, \widehat{g}_{\eta_0}(\widetilde{x}_{t+1}^{\varphi}, \tau_t), \tag{31}$$

$$\tilde{\eta}_{1,t+1} = \eta_{1,t} - \gamma_{1,t}\, \widehat{G}_{\eta_1}(\overline{z}_{t+1}, \tau_t), \tag{32}$$

$$\theta_{t+1} = \tilde{\theta}_{t+1}/\|\tilde{\theta}_{t+1}\|, \tag{33}$$

$$\eta_{1,t+1} = \max\{\eta_{0,t+1},\, \tilde{\eta}_{1,t+1}\}, \tag{34}$$

for some learning rates $\gamma_{1,t}$, $\gamma_{2,t}$, $\gamma_{3,t} > 0$, $\varepsilon_t > 0$; where

$$\widehat{G}_\theta(\{z_i\}_{i=1}^n, \tau) \triangleq \frac{1}{n}\sum_{i=1}^n J(z_i, \tau)(2u_i - \eta_0 - \eta_1)\nabla_\theta f(x_i, \theta), \tag{35}$$

$$\widehat{G}_{\eta_1}(\{z_i\}_{i=1}^n, \tau) \triangleq -\frac{1}{n}\sum_{i=1}^n J(z_i, \tau)(u_i - \eta_0), \tag{36}$$

$$J(z_i, \tau) \triangleq \hat{I}_{\eta_0}(u_i)\big(\hat{I}_{\eta_1}(u_i) - I_1(z_i)\big), \tag{37}$$

for $x_i \triangleq \chi(z_i)$, $u_i \triangleq f(x_i, \theta)$; and

$$\widehat{g}_\theta(\widetilde{x}^{\varphi}, \widetilde{x}^{\psi}, \tau, \varepsilon) \triangleq \big(\widehat{\varphi}(\widetilde{x}^{\varphi}, \tau) - \varphi_*\big)\, \widehat{\psi}_\theta(\widetilde{x}^{\psi}, \tau, \varepsilon), \tag{38}$$

$$\widehat{g}_{\eta_0}(\widetilde{x}^{\varphi}, \tau) \triangleq \big(\widehat{\varphi}(\widetilde{x}^{\varphi}, \tau) - \varphi_*\big)\, \widehat{\psi}_0, \tag{39}$$

for some constant $\widehat{\psi}_0 < 0$, and

$$\widehat{\varphi}\big(\{x_i\}_{i=1}^{n_\varphi}, \tau\big) \triangleq \frac{1}{n_\varphi}\sum_{i=1}^{n_\varphi} \hat{I}_{\eta_0}(u_i), \tag{40}$$

$$\widehat{\psi}_\theta\big(\{x_i\}_{i=1}^{n_\psi}, \tau, \varepsilon\big) \triangleq \frac{1}{\varepsilon n_\psi}\sum_{i=1}^{n_\psi} I_{S(\tau,\varepsilon)}(u_i)\nabla_\theta f(x_i, \theta), \tag{41}$$

where $I_{S(\tau,\varepsilon)}(u) \triangleq \mathbb{1}[\eta_0 < u \le \eta_0 + \varepsilon]$.

Algorithm 1 exhibits the recognizable structure of a recurrent stochastic optimization algorithm. On each iteration, parameters τ_t are subjected to the additive corrections (30)–(32), which are defined by the combination of the step size factors $\gamma_{1,t}$, $\gamma_{2,t}$, $\gamma_{3,t}$ with functions $\widehat{G}_\theta$, $\widehat{G}_{\eta_1}$ and $\widehat{g}_\theta$, $\widehat{g}_{\eta_0}$ as the components of the quasi-gradients of the expected error term U and the constraint term W, respectively. The distribution of these components across the three correction steps reflects a notable trait of the two-threshold decision structure. While the updates (31) and (32) of the thresholds η_0 and η_1 are driven, respectively, by W and U alone (via $\widehat{g}_{\eta_0}$ and $\widehat{G}_{\eta_1}$), the update (30) of the separating surface parameters θ_t is affected by both U and W (via $\widehat{G}_\theta$ and $\widehat{g}_\theta$ combined), thus tying them all into a single cohesive process.

The estimators $\widehat{G}_\theta$, $\widehat{G}_{\eta_1}$ of the quasi-gradient of U have a rather predictable form (35) and (36). The estimators for the quasi-gradient of W, in turn, reflect some challenges that have been principally explained in Section III above. Namely, the estimator $\widehat{g}_\theta$ assembled from (38) and (41) implements the ideas from Sections III-A and III-B, while the estimator $\widehat{g}_{\eta_0}$ given by (39) is additionally motivated by Section III-C.

Finally, by the end of each iteration, the corrected parameters are brought to their regular form, so that the invariants $\|\theta\| = 1$ and $\eta_0 \le \eta_1$ are maintained. The former takes advantage of the definition (2) of the decision function f to eliminate the redundancy and counteract excessive drift in the space of θ by normalizing it in (33). The latter invariant naturally follows from the assumption that the maximal allowed transmission probability φ_* is at least theoretically sufficient to let through all positive observations. Hence, step (34) enforces $\eta_{0,t}$ as the lower bound for $\eta_{1,t}$, so that the transmission constraint prevails over fluctuations in classification threshold. Furthermore, transformations (33) and (34) do not worsen the values of both criterion terms.

Lemma 3. *Let $\widetilde{\tau}_t = \mathrm{vec}[\tilde{\theta}_t, \eta_{0,t}, \tilde{\eta}_{1,t}]$. Then, for all t,*

$$U(\tau_t) \le U(\widetilde{\tau}_t) \quad \textit{and} \quad W(\tau_t) = W(\widetilde{\tau}_t). \tag{42}$$

Proof. By definition (2), $f(x, a\theta) = f(x, \theta)$ for any real a, so $U(\widetilde{\tau}_t) = U(\mathrm{vec}[\theta_t, \eta_{0,t}, \tilde{\eta}_{1,t}])$, $W(\widetilde{\tau}_t) = W(\mathrm{vec}[\theta_t, \eta_{0,t}, \tilde{\eta}_{1,t}])$.

In turn, by definition (10), $W(\tau)$ depends only on θ and η_0 but not η_1. Hence, we immediately have $W(\tau_t)=W(\widetilde{\tau}_t)$.

When $\eta_{0,t}\le\tilde{\eta}_{1,t}$, the maximum in (34) assigns $\eta_{1,t}=\tilde{\eta}_{1,t}$, therefore making $U(\tau_t)=U(\widetilde{\tau}_t)$.

Otherwise, when $\eta_{0,t}>\tilde{\eta}_{1,t}$, step (34) sets $\eta_{1,t}=\eta_{0,t}$, and

$$U(\widetilde{\tau}_t)=\mathbb{E}\big[\hat{I}_{\eta_{0,t}}(u_{\theta_t})(u_{\theta_t}-\eta_{0,t})\big(\hat{I}_{\tilde{\eta}_{1,t}}(u_{\theta_t})-I_1(z)\big)(u_{\theta_t}-\tilde{\eta}_{1,t})\big],$$
$$U(\tau_t)=\mathbb{E}\big[\hat{I}_{\eta_{0,t}}(u_{\theta_t})(u_{\theta_t}-\eta_{0,t})\big(\hat{I}_{\eta_{0,t}}(u_{\theta_t})-I_1(z)\big)(u_{\theta_t}-\eta_{0,t})\big].$$

If $\hat{I}_{\eta_{0,t}}(u_{\theta_t})=0$, then $U(\tau_t)=U(\widetilde{\tau}_t)=0$; alternatively, if $\hat{I}_{\eta_{0,t}}(u_{\theta_t})=1$, then $\hat{I}_{\tilde{\eta}_{1,t}}(u_{\theta_t})=1$, too, so either way,

$$U(\tau_t)-U(\widetilde{\tau}_t)=\mathbb{E}\big[\hat{I}_{\eta_{0,t}}(u_{\theta_t})(u_{\theta_t}-\eta_{0,t}) \cdot\big(1-I_1(z)\big)(\tilde{\eta}_{1,t}-\eta_{0,t})\big]\le 0, \quad (43)$$

since $\hat{I}_{\eta_{0,t}}(u_{\theta_t})(u_{\theta_t}-\eta_{0,t})\ge 0$. □

Altogether, the update (30)–(34) realizes a convergent algorithm that asymptotically satisfies the constraint on the transmission probability and achieves the necessary condition of criterion extremum, in the following formal sense.

Theorem 4. *Given*

1) *a sample of batches of i.i.d. observations* $\{\overline{z}_t\}$;
2) *mutually independent sub-batches* $\widetilde{x}_t^{\varphi}$ *and* $\widetilde{x}_t^{\psi}$ *of feature vectors obtained from every batch* $\overline{z}_t$;
3) *a continuously differentiable decision function* $f(x,\theta)$;
4) *a compact subset* X *in the feature space, such that features* $\chi(z)\in X$ *for all possibly observable* z;
5) *a continuous p.d.f.* $p(x)$ *of observations in the feature space that has compact support on* X *and additionally guarantees* $\Pr[\Gamma(\theta,\eta)]>0$ *for all* θ *and* η *such that* $\Gamma(\theta,\eta)\triangleq\{x\in\operatorname{int}(X): f(x,\theta)=\eta\}\ne\varnothing$;
6) *a sequence* ε_t, *learning rates* $\gamma_{1,t}$, $\gamma_{2,t}$, $\gamma_{3,t}$, *and the ratio* $\gamma_{1,t}/\gamma_{2,t}$ *which are all decreasing in such a way that*

$$\textstyle\sum_{t=1}^{\infty}\gamma_{i,t}^2<\infty,\quad i\in\{1,2,3\}; \quad (44)$$
$$\textstyle\sum_{t=1}^{\infty}\gamma_{j,t}\,\varepsilon_t<\infty,\quad j\in\{1,2\}; \quad (45)$$
$$\textstyle\sum_{t=1}^{\infty}\gamma_{1,t}\,\gamma_{3,t}/\gamma_{2,t}<\infty; \quad (46)$$
$$\textstyle\sum_{t=1}^{\infty}\gamma_{1,t}^2/\gamma_{2,t}=\infty; \quad (47)$$
$$\textstyle\sum_{t=1}^{\infty}\gamma_{3,t}=\infty; \quad (48)$$

Algorithm 1

1) *exhibits the criterion convergence, i.e.,*
$$\lim_{t\to\infty}\left[\frac{\gamma_{1,t}}{\gamma_{2,t}}U(\tau_t)+W(\tau_t)\right]\overset{\text{a.s.}}{=}V_*,\quad \textit{for}\ \ \mathbb{E}V_*<\infty; \quad (49)$$
2) *achieves the global minimum of the constraint penalty* W *on a subsequence, i.e.,*
$$\liminf_{t\to\infty}|\varphi(\tau_t)-\varphi_*|=0; \quad (50)$$
3) *achieves the the necessary condition of extremum in the classification threshold* η_1 *for the expected error penalty* U *on a subsequence, i.e.,*
$$\liminf_{t\to\infty}G_{\eta_1}(\tau_t)=0; \quad (51)$$
4) *achieves the the necessary condition of extremum for the criterion* V *on a subsequence, i.e.,*
$$\liminf_{t\to\infty}\left[G_\theta(\tau_t)+\frac{\gamma_{2,t}}{\gamma_{1,t}}\nabla_\theta W(\tau_t)\right]=0. \quad (52)$$

Proof. To establish the convergence of the algorithm, in the following we employ the stochastic Lyapunov function method and martingale theory [14]–[17]. For the Lyapunov function we use the criterion (9) with $\zeta_t=\gamma_{2,t}/\gamma_{1,t}$.

1. Parametric boundedness. Let us ensure that Algorithm 1 does not allow parameters τ_t to escape to infinity in magnitude.

First of all, step (33) of the algorithm assures that $\|\theta_t\|=1$ holds after each iteration. Moreover, since the decision function $f(x,\theta)$ and its gradients in both arguments are continuous, and all observed feature vectors x come from the compact X, there is a constant C_f such that, for any $x\in X$,

$$\max\big\{|f(x,\theta_t)|,\ \|\nabla_x f(x,\theta_t)\|,\ \|\nabla_\theta f(x,\theta_t)\|\big\}\le C_f. \quad (53)$$

As to the thresholds η_0 and η_1, step (34) of the algorithm guarantees that the order $\eta_{0,t}\le\eta_{1,t}$ is maintained, so it is only sufficient to show that $\eta_{0,t}$ and $\eta_{1,t}$ cannot indefinitely decrease toward $-\infty$ and increase toward $+\infty$, respectively.

If we assume the contrary for η_0, there must come some moment t when $\eta_{0,t}<f_{\inf}\triangleq\inf f(x,\theta)$, with the infimum taken over all x and θ such that $x\in X$ and $\|\theta\|=1$. Then, in the next iteration of the algorithm, $\widehat{\varphi}(\widetilde{x}_{t+1}^{\varphi},\tau_t)=1$ regardless of the observations in the new batch $\widetilde{x}_{t+1}^{\varphi}$ and the value of θ_t. Hence, since $\widehat{\varphi}(\widetilde{x}_{t+1}^{\varphi},\tau_t)-\varphi_*=1-\varphi_*\ge 0$ and $\widehat{\psi}_0<0$, from (31) and (39) it is clear that η_0 cannot decrease any further:

$$\eta_{0,t+1}=\eta_{0,t}-\gamma_{3,t}(1-\varphi_*)\widehat{\psi}_0\ge\eta_{0,t}. \quad (54)$$

Likewise, when $\eta_{0,t}>f_{\sup}\triangleq\sup f(x,\theta)$ or $\eta_{1,t}>f_{\sup}$,

$$\eta_{0,t+1}=\eta_{0,t}+\gamma_{3,t}\varphi_*\widehat{\psi}_0\le\eta_{0,t}, \quad (55)$$
$$\tilde{\eta}_{1,t+1}=\eta_{1,t}-\frac{\gamma_{1,t}}{n}\sum_{i=1}^{n}I_1(z_i)\hat{I}_{\eta_{0,t}}(u_i)(u_i-\eta_{0,t})\le\eta_{1,t}, \quad (56)$$

respectively. Since both $\eta_{0,t}$ and $\tilde{\eta}_{1,t}$ are bounded from above, so must be $\eta_{1,t}$ as the maximum of the two.

Thus, from (54), (56), (55), and the fact that $\|\theta_t\|=1$, we have established that there must exist a constant C_τ such that

$$\|\tau_t\|\le C_\tau. \quad (57)$$

2. Risk term variation. Herein, for the sake of brevity, let $\tau=\operatorname{vec}[\theta,\eta_0,\eta_1]$, $\tau'=\operatorname{vec}[\theta',\eta_0',\eta_1']$, and $\tau''=\operatorname{vec}[\theta'',\eta_0'',\eta_1'']$ be, respectively, the current parameters at the start of an iteration of Algorithm 1, the corrected parameters after steps (30)–(32), and them in the regular form enforced by steps (33), (34).

As it has been established in Lemma 3, $U(\tau'')\le U(\tau')$. For that reason, we can now concentrate on the difference $U(\tau')-U(\tau)$ and split it into two variations: one in which the only changing factors are those involving η_0, and the other in which those are the ones involving η_1. Replacing the differences $(u_{\theta'}-\eta_i')-(u_\theta-\eta_i)$ with their Taylor series expansions in τ, we will then have:

$$U(\tau'')-U(\tau)\le U(\tau')-U(\tau) \quad (58)$$

$$= E_{1,1}(\tau,\tau') + \langle \tau' - \tau,\, E_{1,2}(\tau) + E_{1,3}(\tau,\tau')\rangle$$
$$+ E_{2,1}(\tau,\tau') + \langle \tau' - \tau,\, E_{2,2}(\tau)\rangle + O\big(\|\tau' - \tau\|^2\big),$$

where, for $\hat{J}^{(i)}_{\tau',\tau} \triangleq \hat{I}_{\eta_i'}(u_{\theta'}) - \hat{I}_{\eta_i}(u_\theta)$, $J^{(1)}_\tau \triangleq \hat{I}_{\eta_1}(u_\theta) - I_1(z)$,

$$E_{1,1}(\tau,\tau') \triangleq \mathbb{E}\big[(u_{\theta'} - \eta_0')(u_{\theta'} - \eta_1')\hat{J}^{(0)}_{\tau',\tau} J^{(1)}_{\tau'}\big], \tag{59}$$

$$E_{1,2}(\tau) \triangleq \mathbb{E}\big[\nabla_\tau[u_\theta - \eta_0](u_\theta - \eta_1)\hat{I}_{\eta_0}(u_\theta) J^{(1)}_\tau\big], \tag{60}$$

$$E_{1,3}(\tau,\tau') \triangleq \mathbb{E}\big[\nabla_\tau[u_\theta - \eta_0](u_\theta - \eta_1)\hat{I}_{\eta_0}(u_\theta)\hat{J}^{(1)}_{\tau',\tau}\big], \tag{61}$$

$$E_{2,1}(\tau,\tau') \triangleq \mathbb{E}\big[(u_\theta - \eta_0)(u_{\theta'} - \eta_1')\hat{I}_{\eta_0}(u_\theta)\hat{J}^{(1)}_{\tau',\tau}\big], \tag{62}$$

$$E_{2,2}(\tau) \triangleq \mathbb{E}\big[(u_\theta - \eta_0)\nabla_\tau[u_\theta - \eta_1]\hat{I}_{\eta_0}(u_\theta) J^{(1)}_\tau\big]. \tag{63}$$

Because both $f(x,\theta)$ and $\nabla_\tau f(x,\theta)$ are bounded by (53),

$$|E_{1,3}(\tau,\tau')| \le C\,\mathbb{E}\big[|\hat{I}_{\eta_1'}(u_{\theta'}) - \hat{I}_{\eta_1}(u_\theta)|\big], \tag{64}$$

for some constant C. It can be shown that the expectation on the right side in (64) is such that

$$E_{1,3}(\tau,\tau') = O(\|\tau' - \tau\|). \tag{65}$$

Further, it can be shown that $E_{1,1}(\tau,\tau') \le O\big(\|\tau' - \tau\|^2\big)$ and $E_{2,1}(\tau,\tau') \le O\big(\|\tau' - \tau\|^2\big)$, so from (58) together with (65) we can finally reach the following bound:

$$U(\tau'') - U(\tau) \le \langle \tau' - \tau,\, E_{1,2}(\tau) + E_{2,2}(\tau)\rangle + O\big(\|\tau' - \tau\|^2\big). \tag{66}$$

3. Constraint term variation. By the virtue of Lemma 3, we know that $W(\tau'') = W(\tau')$. Therefore,

$$W(\tau'') - W(\tau) = W(\tau') - W(\tau) \tag{67}$$
$$= \frac{1}{2}\big(\varphi(\tau') - \varphi(\tau)\big)^2 + \big(\varphi(\tau) - \varphi_*\big)\big(\varphi(\tau') - \varphi(\tau)\big),$$

where $\varphi(\tau) \triangleq \mathbb{E}\big[\hat{I}_{\eta_0}(u_\theta)\big]$. Denoting $\psi(\tau) \triangleq \nabla_\tau\varphi(\tau)$ and substituting the Taylor series expansion for $\varphi(\tau)$ into (67), we straightforwardly obtain:

$$W(\tau'') - W(\tau) = (\varphi(\tau) - \varphi_*)\,\psi(\tau) + O\big(\|\tau' - \tau\|^2\big). \tag{68}$$

In order to advance with the gradient $\psi(\tau)$, we can invoke Lemma 1 introduced earlier in Section III. With its help, we can express $\psi(\tau)$, the gradient of the volume integral $\varphi(\tau)$, as a surface integral. With the partials for θ and η_0 written separately, it will look as follows.

$$\psi_\theta(\tau) \triangleq \nabla_\theta\varphi(\tau) = \int\limits_{\{x: f(x,\theta) = \eta_0\}} \frac{\nabla_\theta f(x,\theta)}{\|\nabla_x f(x,\theta)\|}\, p(x)\, d\sigma, \tag{69}$$

$$\psi_\eta(\tau) \triangleq \frac{d}{d\eta_0}\varphi(\tau) = -\int\limits_{\{x: f(x,\theta) = \eta_0\}} \frac{1}{\|\nabla_x f(x,\theta)\|}\, p(x)\, d\sigma. \tag{70}$$

4. Combined criterion variation. For a pair of penalty coefficients $\zeta \le \zeta'$, from (58) and (67), we have:

$$V(\tau'',\zeta') - V(\tau,\zeta) \le V(\tau',\zeta') - V(\tau,\zeta)$$
$$\le \frac{1}{\zeta}\big(U(\tau') - U(\tau)\big) + W(\tau') - W(\tau). \tag{71}$$

In parts 3 and 4 of the proof, we have separately obtained inequalities (66) and (68) for the changes induced by a single step of the algorithm in the expected penalty term U and the constraint term W. Substituting those results into the right side of (71), we can now reveal the following structure in the difference of the combined criterion V:

$$V(\tau',\zeta') - V(\tau,\zeta) \le \left\langle \theta' - \theta,\, \frac{1}{\zeta}G_\theta(\tau) + g_\theta(\tau)\right\rangle \tag{72}$$
$$+ (\eta_0' - \eta_0)\left(\frac{1}{\zeta}G_{\eta_0}(\tau) + g_{\eta_0}(\tau)\right)$$
$$+ (\eta_1' - \eta_1)\,\frac{1}{\zeta}G_{\eta_1}(\tau) + O\big(\|\tau' - \tau\|^2\big),$$

where, for $J(x,\tau)$ defined as in (37),

$$G_\theta(\tau) \triangleq \mathbb{E}\big[J(x,\tau)(2u_\theta - \eta_0 - \eta_1)\nabla_\theta f(x,\theta)\big], \tag{73}$$

$$G_{\eta_0}(\tau) \triangleq -\mathbb{E}\big[J(x,\tau)(u_\theta - \eta_1)\big], \tag{74}$$

$$G_{\eta_1}(\tau) \triangleq -\mathbb{E}\big[J(x,\tau)(u_\theta - \eta_0)\big], \tag{75}$$

$$g_\theta(\tau) \triangleq \mathbb{E}\big[(\varphi(\tau) - \varphi_*)\psi_\theta(\tau)\big], \tag{76}$$

$$g_{\eta_0}(\tau) \triangleq \mathbb{E}\big[(\varphi(\tau) - \varphi_*)\psi_\eta(\tau)\big]. \tag{77}$$

5. Supermartingale inequalities. Let $V_t \triangleq V(\tau_t, \gamma_{2,t}/\gamma_{1,t})$. Taking the expectation of (72) conditioned on the trajectory of the algorithm for a sequence of batches $\overline{z}_1,\dots,\overline{z}_t$, we get:

$$\mathbb{E}\big[V_{t+1} - V_t \mid \overline{z}_{1:t}\big]$$
$$\le \left\langle \mathbb{E}[\tilde{\theta}_{t+1} - \theta_t \mid \overline{z}_{1:t}],\, \frac{\gamma_{1,t}}{\gamma_{2,t}}G_\theta(\tau_t) + g_\theta(\tau_t)\right\rangle$$
$$+ \mathbb{E}[\eta_{0,t+1} - \eta_{0,t} \mid \overline{z}_{1:t}]\left(\frac{\gamma_{1,t}}{\gamma_{2,t}}G_{\eta_0}(\tau_t) + g_{\eta_0}(\tau_t)\right)$$
$$+ \mathbb{E}[\tilde{\eta}_{1,t+1} - \eta_{1,t} \mid \overline{z}_{1:t}]\,\frac{\gamma_{1,t}}{\gamma_{2,t}}G_{\eta_1}(\tau_t)$$
$$+ C_1\,\mathbb{E}\big[\|\tau_{t+1} - \tau_t\|^2 \mid \overline{z}_{1:t}\big], \tag{78}$$

for some constant C_1. It can be readily seen that

$$\mathbb{E}\big[\widehat{G}_\theta(\overline{z},\tau)\big] = G_\theta(\tau), \tag{79}$$

$$\mathbb{E}\big[\widehat{G}_{\eta_1}(\overline{z},\tau)\big] = G_{\eta_1}(\tau), \tag{80}$$

$$\mathbb{E}\big[\widehat{\varphi}(\widetilde{x}^{\varphi},\tau)\big] = \varphi(\tau). \tag{81}$$

Using Lemma 1 for taking a derivative of a volume integral once again, we establish a similar link between $\widehat{\psi}_\theta$ and ψ_θ:

$$\mathbb{E}\big[\widehat{\psi}_\theta(\widetilde{x}^{\psi},\tau,\varepsilon)\big] = \frac{1}{\varepsilon}\int\limits_{S(\tau,\varepsilon)} \nabla_\theta f(x,\theta)\, p(x)\, dx$$
$$= \frac{1}{\varepsilon}\int\limits_{\{x: f(x,\theta) > \eta_0\}} \nabla_\theta f(x,\theta)\, p(x)\, dx - \frac{1}{\varepsilon}\int\limits_{\{x: f(x,\theta) > \eta_0 + \varepsilon\}} \nabla_\theta f(x,\theta)\, p(x)\, dx$$
$$= -\left[\frac{d}{d\varepsilon}\int\limits_{\{x: f(x,\theta) > \eta_0 + \varepsilon\}} \nabla_\theta f(x,\theta)\, p(x)\, dx\right]_{\varepsilon = 0} + b(\varepsilon)$$
$$= \psi_\theta(\tau) + b(\varepsilon), \tag{82}$$

where $\|b(\varepsilon)\| = O(\varepsilon)$. Since $\widetilde{x}^{\varphi}$ and $\widetilde{x}^{\psi}$ are independent,

$$\mathbb{E}\big[\widehat{g}_\theta(\widetilde{x}^{\varphi},\widetilde{x}^{\psi},\tau,\varepsilon)\big] = \mathbb{E}\big[\big(\widehat{\varphi}(\widetilde{x}^{\varphi},\tau) - \varphi_*\big)\widehat{\psi}_\theta(\widetilde{x}^{\psi},\tau,\varepsilon)\big]$$

$$
\begin{aligned}
&= \mathbb{E}\big[\widehat{\varphi}(\widetilde{x}^{\varphi}, \tau) - \varphi_*\big]\, \mathbb{E}\big[\widehat{\psi}_\theta(\widetilde{x}^{\psi}, \tau, \varepsilon)\big] \\
&= \big(\varphi(\tau) - \varphi_*\big)\big(\psi_\theta(\tau) + b(\varepsilon)\big) \\
&= g_\theta(\tau) + O(\varepsilon), \qquad (83)
\end{aligned}
$$

$$
\mathbb{E}\big[\widehat{g}_{\eta_0}(\widetilde{x}^{\varphi}, \tau)\big] = (\varphi(\tau) - \varphi_*)\, \widehat{\psi}_0. \qquad (84)
$$

Substituting (30)–(34) into (78), taking into account (79), (80), (83), (84), and removing the parentheses, we have the bound:

$$
\begin{aligned}
&\mathbb{E}[V_{t+1} - V_t \mid \overline{z}_{1:t}] \\
&\quad \le -\frac{\gamma_{1,t}^2}{\gamma_{2,t}} \left\| G_\theta(\tau_t) + \frac{\gamma_{2,t}}{\gamma_{1,t}} g_\theta(\tau_t) \right\|^2 - \frac{\gamma_{1,t}^2}{\gamma_{2,t}} \big(G_{\eta_1}(\tau_t)\big)^2 \\
&\qquad - \gamma_{3,t} (\varphi(\tau_t) - \varphi_*)^2\, \psi_\eta(\tau_t)\, \widehat{\psi}_0 + R_t. \qquad (85)
\end{aligned}
$$

Here the remainder term equals to

$$
R_t \triangleq C_3 \left(\textstyle\sum_{i=1}^{3} \gamma_{i,t}^2 + (\gamma_{1,t} + \gamma_{2,t})\, \varepsilon_t + \dfrac{\gamma_{1,t}\gamma_{3,t}}{\gamma_{2,t}} \right), \qquad (86)
$$

where $C_3 \triangleq C_2 \max\{1,\, \|G_\theta(\tau)\|^2,\, \|g_\theta(\tau)\|^2,\, |\widehat{\psi}_0|\, \|G_{\eta_0}(\tau)\|\}$ is constant, as the quasi-gradients are bounded via (53), and C_2 is another constant.

6. Convergence results. Given that the first three terms on the right side of (85) are all negative, it follows that, up to the remainder term R_t, the sequence V_t forms a supermartingale:

$$
\mathbb{E}[V_{t+1} \mid \overline{z}_{1:t}] \le \mathbb{E}[V_t \mid \overline{z}_{1:t}] + R_t. \qquad (87)
$$

Accordingly, by standard reasoning [16], V_t must satisfy (49).

On the other hand, constructing a telescoping sum out of (85) and taking the unconditional expectation, eliminating the observation history $\overline{z}_{1:t}$ from consideration, we will have

$$
\begin{aligned}
&\sum_{t=1}^{\infty} \frac{\gamma_{1,t}^2}{\gamma_{2,t}} \mathbb{E} \left\| G_\theta(\tau_t) + \frac{\gamma_{2,t}}{\gamma_{1,t}} g_\theta(\tau_t) \right\|^2 + \sum_{t=1}^{\infty} \frac{\gamma_{1,t}^2}{\gamma_{2,t}} \mathbb{E}\big(G_{\eta_1}(\tau_t)\big)^2 \\
&\qquad + \widehat{\psi}_0 \sum_{t=1}^{\infty} \gamma_{3,t} \mathbb{E}\big[(\varphi(\tau_t) - \varphi_*)^2\, \psi_\eta(\tau_t)\big] \\
&\qquad\qquad \le \mathbb{E}V_* - \mathbb{E}V_1 + \sum_{t=1}^{\infty} R_t. \qquad (88)
\end{aligned}
$$

Due to the conditions (44)–(48) of the theorem, the right side of (88) is finite and, therefore, all three series on its left side are summable, which, in turn, implies the desired conclusions (51) and (52), as well as

$$
\liminf_{t\to\infty} \mathbb{E}\big[(\varphi(\tau_t) - \varphi_*)^2\, |\psi_\eta(\tau_t)|\big] = 0. \qquad (89)
$$

From the inequalities (54) and (55), the requirement of divergence (48), and the structure of the update (31) together with its step size (39), it can be seen that, on the trajectory of the algorithm, when $t \to \infty$, the event of $\psi_\eta(\tau_t) \neq 0$ occurs infinitely often, and the conclusion (50) must also hold. □

V. Conclusions

In this paper, we considered the problem of optimal organization of the sensor-edge cooperation in the setting where the available wireless channel is inadequate for the volume of observations, and existing approaches are insufficient due to the computation and communication constraints. We proposed a way to solve this problem by introducing a pair of decision rules, where one is responsive for detecting observations of interest, and the other is responsive for controlling channel access. We constructed a stochastic optimization algorithm for finding optimal parameters of these decision rules, so that, taken together, they deliver a solution for both uses. The asymptotic behavior of the resulting adaptation procedure was studied and shown to exhibit criterial convergence, realizing on the convergence set the necessary conditions of optimality, which correspond to the probability of channel access approaching a desired level, and the expected loss of misdetections tending to minimum, provided the former holds. The proposed approach is of interest for autonomous IoT systems, in which problems of online adaptation need to be solved in the presence of probability function constraints.

References

[1] L. Atzori, A. Iera, and G. Morabito, “The Internet of Things: A survey,” *Computer networks*, vol. 54, no. 15, pp. 2787–2805, 2010.

[2] R. K. Kodali, V. Jain, S. Bose, and L. Boppana, “IoT based smart security and home automation system,” in *International Conference on Computing, Communication and Automation (ICCCA)*. IEEE, 2016, pp. 1286–1289.

[3] T. Zhang, A. Chowdhery, V. Bahl, K. Jamieson, and S. Banerjee, “The design and implementation of a wireless video surveillance system,” in *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*. ACM, 2015, pp. 426–438.

[4] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, “Internet of things for smart cities,” *IEEE Internet of Things Journal*, vol. 1, no. 1, pp. 22–32, Feb. 2014.

[5] C. T. Barba, M. A. Mateos, P. R. Soto, A. M. Mezher, and M. A. Igartua, “Smart city for VANETs using warning messages, traffic statistics and intelligent traffic lights,” in *IEEE Intelligent Vehicles Symposium*. IEEE, 2012, pp. 902–907.

[6] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, “The case for VM-based cloudlets in mobile computing,” *IEEE Pervasive Computing*, vol. 8, no. 4, pp. 14–23, Oct. 2009.

[7] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, “Fog computing and its role in the Internet of Things,” in *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*. ACM, 2012, pp. 13–16.

[8] Y. Kim, E. Park, S. Yoo, T. Choi, L. Yang, and D. Shin, “Compression of deep convolutional neural networks for fast and low power mobile applications,” *CoRR*, vol. abs/1511.06530, 2015. [Online]. Available: http://arxiv.org/abs/1511.06530

[9] E. Raik, “The differentiability in the parameter of the probability function and optimization of the probability function via the stochastic pseudo-gradient method,” *Izvestiya Akad. Nayk Est. SSR*, vol. 24, no. 1, pp. 3–6, 1975, in Russian.

[10] A. I. Kibzun and Y. S. Kan, *Stochastic Programming Problems With Probability and Quantile Functions*, ser. Interscience Series In Systems and Optimization. Wiley, 1996, vol. 9.

[11] S. Uryas’ev, “Differentiability of the integral over a set defined by inclusion,” *Cybernetics*, vol. 24, no. 5, pp. 638–642, 1988.

[12] ——, “Derivatives of probability functions and integrals over sets given by inequalities,” *Journal of Computational and Applied Mathematics*, vol. 56, pp. 197–223, 1994.

[13] K. Marti, *Stochastic Optimization Methods*. Springer, 2005.

[14] J. L. Doob, *Stochastic Processes*, revised ed. New York: Wiley, 1990.

[15] H.-F. Chen, *Stochastic Approximation and Its Applications*, ser. Nonconvex Optimization and Its Applications. Dordrecht: Kluwer Academic Publishers, 2002, vol. 64.

[16] B. T. Polyak, *Introduction to Optimization*, ser. Translations Series in Mathematics and Engineering. New York: Optimization Software, 1987.

[17] A. N. Shiryaev, *Probability*, 2nd ed., ser. Graduate Texts in Mathematics. New York: Springer-Verlag, 1996, vol. 95.