On Distinguishing between Reliable and Unreliable Sensors Without a Knowledge of the Ground Truth

Anis Yazidi
Dept. of Computer Science
University College of Oslo and Akershus
Oslo, Norway
Email: anis.yazidi@hioa.no

B. John Oommen
School of Computer Science
Carleton University, Ottawa, Canada
oommen.scs@carleton.ca

Morten Goodwin
Dept. of ICT
University of Agder, Grimstad, Norway
morten.goodwin@uia.no

Abstract—In many applications, data from different sensors are aggregated in order to obtain more reliable information about the process that the sensors are monitoring. However, the quality of the aggregated information is intricately dependent on the reliability of the individual sensors. In fact, unreliable sensors will tend to report erroneous values of the ground truth, and thus degrade the quality of the fused information. Finding strategies to identify unreliable sensors can assist in having a counter-effect on their respective detrimental influences on the fusion process, and this has has been a focal concern in the literature. The purpose of this paper is to propose a solution to an extremely pertinent problem, namely, that of identifying which sensors are unreliable *without any knowledge of the ground truth*. This fascinating paradox can be formulated in simple terms as trying to identify *stochastic* liars without any additional information about the truth. Though apparently impossible, we will show that it is feasible to solve the problem, a claim that *is counter-intuitive in and of itself*. To the best of our knowledge, this is the first reported solution to the aforementioned paradox. Legacy work and the reported literature have merely addressed assessing the reliability of a sensor by comparing its reading to the ground truth either in an online or an offline manner. The informed reader will observe that the so-called Weighted Majority Algorithm is a representative example of a large class of such legacy algorithms. The essence of our approach involves studying the agreement of each sensor with the rest of the sensors, and not comparing the reading of the individual sensors with the ground truth – as advocated in the literature. Under some mild conditions on the reliability of the sensors, we can prove that we can, indeed, filter out the unreliable ones. Our approach leverages the power of the theory of Learning Automata (LA) so as to gradually learn the identity of the reliable and unreliable sensors. To achieve this, we resort to a team of LA, where a distinct automaton is associated with each sensor. The solution provided here has been subjected to rigorous experimental tests, and the results presented are, in our opinion, both novel and conclusive.

Keywords: *Sensor Fusion, Learning Automata*

I. INTRODUCTION

In many applications, data from different sources is received, processed and then fused, to obtain more reliable information about the process being monitored. This is often the case in industrial applications where multiple redundant sensors are used to measure the same quantities [11], [10], and for example, in nuclear or space applications, where human intervention is not possible. Sensors usually provide imprecise and uncertain observations. The field of sensor fusion involves a set of redundant sensors measuring the same physical quantity. This redundancy permits the operators to obtain a robustness of sorts, whenever some sensors are prone to error.

Furthermore, fused data will reduce or eliminate the effects due to failures of a few sensors operating in the system. Most of the research on fusing multiple sensor data merely assume that the confidence levels in the measurements are known. The accuracy of an observation can be computed by comparing the current observation with the reference data set and/or by performing physical investigation. However, performing a physical investigation or having a reference data set is not practical in many monitoring scenarios, although it is possible to adopt such measures during training or within a limited scope. To the best of our knowledge, trying to assess the reliability of a sensor without any additional information about the ground truth is still an open research question that has not been addressed before, and our strategy for resolving this will be discussed in the body of this paper.

The first question to be addressed is whether the problem of detecting an unreliable sensor without knowing the ground truth is even a solvable problem. Our position is that if there is no other information, it is a futile venture. But if we consider the fact that there is a set of sensors, all of which are measuring the same quantity, the information provided by the *other* sensors can provide invaluable metrics about how good any specific sensor is. This, indeed, is the philosophy that we advocate. The question of how the information from the other sensors is to be gleaned and processed is really, in and of itself, unsolved. Suffice it to state that we emphasize that our solution to the problem lies in investigating the level of agreement between the various data sources/sensors, which, in turn, constitutes valuable information to fuse them in an efficient manner. In simple words, we assert the rather fascinating claim that given a group of sensors, we can find the sub-group of unreliable sensors without any knowledge of the ground truth, if we also permit each sensor to be compared to the others!

In order to position our work in relation with the existing work, we shall present a brief review of the state-of-the-art related to data fusion. The legacy research has focused on

fusing sensor information under either known or estimated confidence levels. Most current data fusion methods employ probabilistic descriptions of observations and processes, and use Bayesian principles to combine this information. Other approaches rely on principles derived from evidential reasoning including Dempster-Shafer inference theory [3] and subjective logic [13]. Elmenreich [7] presented a novel algorithm that uses the estimated variance of each sensor measurement in order to find the optimal averaging weights. Another theme akin to multi-sensor fusions involves "prediction using expert advice" [15], where the performance is always nearly as good as the best forecasting strategy. The fault-tolerant averaging algorithm was first introduced by Marzullo [16] in the context of time synchronization in distributed systems. Afterwards, it was used in the domain of information fusion to fuse a set of abstract sensors into a single reliable abstract sensor that is correct even when some of the original sensors are incorrect or faulty. Consensus algorithms, such as majority voting, are suitable for fusing binary measurements.

Most approaches rely on accessing the ground truth to compute the accuracy of a sensor. The work by Hossain [12] is representative of such approaches: It computes the accuracy by comparing the observations provided by the sensor with the ground truth in the training data. This can be experimentally computed by comparing the outcome of the online sensor observations with the ground truth, and by repeating this process multiple times. The approach of Hossain and his colleagues considers the opinions of the sensors in performing a common observation, and proceeds to group the opinions into two subgroups, namely those which support the occurrence of the event and those which oppose it. The scheme then determines the winning group and increases the confidence of the sensors in that group (by considering this event as a "reward"), while, at the same time, it decreases the confidence of the sensors of the other group.

The Condorcet Jury Theorem demonstrates that the Majority group is always better at selecting superior alternatives than any single individual member [4]. There are some limitations to the hypotheses governing the theorem. In fact, it requires that each individual makes the right decision with a probability $p > 0.5$, and that all individuals are homogenous in p. Probably the most notable extension of this is the scenario when the population is not homogenous. Boland [4] assumes that the voters can be divided into two groups. The first group consists of individuals whose "true" interest lie in one direction, while the other group consists of those whose "true" interests lie in the other. When mapped to the case of sensor aggregation, we again have two groups, where the first group consists of reliable sensors that possess the "true" interest of reporting the ground truth, while the alternate group of unreliable sensors possess a "true" interest in misreporting it. Determining ways to solve the ($ATPP$) [26] and thus counter the detrimental influence of unreliable agents on a Reputation System, has been a focal concern of a number of very interesting studies [5], [6], [17], [22], [28], [29].

It is worth noting that the task of combining reports from different witnesses is akin to the problem of fusing possibly conflicting sources of information [2], [8]. Buchegger and Le Boudec [5] tackled the latter issue as follows: They proposed a Bayesian reputation mechanism in which each

node isolates malicious nodes by applying a so-called *deviation test* methodology. Their approach requires each agent to have enough *direct* experience with the services so that he can evaluate the trustworthiness of the reports of the witnesses. While this is a desirable option, unfortunately, in real life, such an assumption does not always hold, specially when the number of possible services is large.

This problem, of separating reliable and unreliable sensors, is called the Sensor-Type Partitioning Problem ($STPP$). Put in a nutshell, in this paper, we propose to solve the above-mentioned paradoxical $STTP$ using tools provided by Learning Automata (LA), which have proven powerful potential in efficiently and quickly learning the optimal action when operating in unknown stochastic environments. It adaptively, and in an on-line manner, gradually learns the identity and characteristics of the sensors that are reliable and those that are unreliable. In addition, we will provide two approaches for fusing the sensor readings which leverage the convergence result of our LA-based partitioning. Rigorous theoretical results and a host of empirical results will be presented in this paper. Our work differs from the aforementioned research since we aim to infer the confidence of the measurements based on their level of agreements *in the absence of knowledge of the ground truth*.

A. Paper Organization

Earlier, in Section I we introduced the research problem and presented a brief survey of the available solutions for dealing with reliable and unreliable sensors. The rest of the paper is organized as follows. First of all, in Section II, we submit a formal statement of the problem. Then, in Section III we present a brief overview of the field of LA. Thereafter, in Section IV we present our solution, which is the LA-based scheme for identifying unreliable sensors in a stochastic environment in the absence of knowledge of the ground truth. Experimental results obtained by rigorously testing our solution for a variety of scenarios and for agents with different characteristics, are presented in Section V. Section VI concludes the paper and discusses open avenues for future work.

II. MODELING THE PROBLEM

We consider a population of N sensors, $\mathcal{S} = \{s_1, s_2, \ldots, s_N\}$. Let the real situation of the environment at the time instant t be modeled by a binary variable $T(t)$, which can take one of two possible values, 0 and 1. The value of T is unknown and can only be inferred through measurements from sensors. The output from the sensor s_i is referred to as x_i. Let π be the probability of the state of the ground truth, i.e., $T = 0$ with probability π.

To formalize the scenario, we record four possibilities:

- $x_i = T$ (where $x_i = 0$ or 1): This is the case when the sensor correctly reports the ground truth.

- $x_i \neq T$ (where $x_i = 0$ or 1): This is the case when the sensor faultily reports the ground truth.

In our discussions, we make one simplifying assumption: The probability of the sensor reporting a value erroneously is symmetric. In other words, in terms of the binary detection

problem, we assume that the probability of a false alarm and the so-called miss probability are both equal. Formally, we assume that:

$$Prob(x_i = 0|T = 1) = Prob(x_i = 1|T = 0). \quad (1)$$

Further, let q_i denote the Fault Probability (FP) of sensor s_i, where:

$$q_i = Prob(x_i = 0|T = 1) = Prob(x_i = 1|T = 0).$$

Similarly, we define the Correctness Probability (CP) of sensor s_i as $p_i = 1 - q_i$.

It is easy to prove that the total probability $Prob(x_i = T)$ is, indeed, p_i, since, in fact:

$$
\begin{aligned}
Prob(x_i = T) &= Prob(T = 0)Prob(x_i = 0|T = 0) \\
&\quad + Prob(T = 1)Prob(x_i = 1|T = 1) \\
&= \pi p_i + (1 - \pi)p_i \\
&= p_i. \quad (2)
\end{aligned}
$$

Thus, the quantity $p_i = Prob(x_i = T)$ can be re-rewritten as $p_i = Prob(I\{x_i = T\} = 1)$, where $I\{.\}$ is the Indicator function.

We refer to a sensor as being reliable when it has a FP $q_i < 0.5$. Conversely, the sensor is unreliable when it has a FP $q_i > 0.5$. Equivalently, we can define a reliable sensor to be one that has a CP $p_i > 0.5$ and an unreliable sensor as one that has a CP of $p_i < 0.5$.

Observe that as a result of this model, a reliable sensor will probabilistically tend to report 0 when the ground truth is 0, and 1 when the ground truth is 1. Otherwise, it is clearly, unreliable. Our aim, then, is to partition the sensors as being reliable or unreliable. Furthermore, once partitioned, our aim is to use the partitioning as a basis for better fusion.

To simplify the analysis[1], we assume that every p_i can assume one of two possible values from the set $\{p_R, p_U\}$, where $p_R > 0.5$ and $p_U < 0.5$. Then, a sensor s_i is said to be reliable if $p_i = p_R$, and is said be unreliable if $p_i = p_U$. To render the problem non-trivial and interesting, we assume that p_R and p_U are unknown to the algorithm.

Based on the above, the set of reliable sensors is $\mathcal{S}_R = \{s_i|p_i = p_R\}$, and the set of unreliable sensors is $\mathcal{S}_U = \{s_i|p_i = p_U\}$.

We now formalize the Sensor-Type Partitioning Problem (STPP). The STPP involves a set of N sensors[2], $\mathcal{S} = \{s_1, s_2, \ldots, s_N\}$, where each sensor s_i is characterized by a fixed but unknown probability p_i of it sensing the ground truth correctly. The STPP involves partitioning $\mathcal{S}$ into 2 mutually exclusive and exhaustive groups so as to obtain a 2-partition

$\mathbb{G} = \{G_U, G_R\}$, such that each group, G_R, of size, N_R, and G_U, of size N_U, exclusively contains only the sensors of its own type, i.e., which are either reliable or unreliable respectively.

We define $P_{(N_R-1, N_U)}$ as the probability of a deterministic majority voting scheme, which involves the opinions of N_R-1 reliable sensors and N_U unreliable ones, to yield the correct decision using the majority rule. In other words, this is the probability that a majority of more than $(N_R - 1 + N_U)/2$ of the sensors will advocate the ground truth. Similarly, we define $P_{(N_R, N_U-1)}$ as the probability of a deterministic majority voting scheme, which involves the opinions of N_R reliable sensors and $N_U - 1$ unreliable ones, to yield the correct decision using the majority rule. As one can see, this quantity is the same: It too is the probability that a majority of more than $(N_R + N_U - 1)/2$ of the sensors will, in turn, advocate the ground truth.

To render the problem meaningful and solvable[3], we will assume that:

$$(N_R - 1)p_R + N_U p_U > (N_R + N_U)/2.$$

This mild condition that we require in this paper, is founded on a fundamental premise that has to hold in any sustainable society, where telling the "truth" is considered a virtue, while "lying" is considered detrimental and harmful to the society. The rationale for invoking this is the following: A reliable sensor will tend to agree with the averaged/aggregated opinion of the rest of other sensors, and thus by comparing the reading of any specific sensor with the rest of other sensors, we hypothesize that we will be able to detect sensors that deviate from the accepted norm even without knowing the ground truth.

III. STOCHASTIC LEARNING AUTOMATA

Learning Automata(LA) have been used in systems that have incomplete knowledge about the Environment in which they operate [1], [18], [20], [24]. The learning mechanism attempts to learn from a *stochastic Teacher* which models the Environment. In his pioneering work, Tsetlin [25] attempted to use LA to model biological learning. In general, a random action is selected based on a probability vector, and these action probabilities are updated based on the observation of the Environment's response, after which the procedure is repeated.

The term "Learning Automata" was first publicized by Narendra and Thathachar [18]. The goal of LA is to "determine the optimal action out of a set of allowable actions" [1]. The distinguishing characteristic of automata-based learning is that the search for the optimizing parameter vector is conducted in the space of probability *distributions* defined over the parameter space, rather than in the parameter space itself [23].

In the first LA designs, the transition and the output functions were time invariant, and for this reason these LA were considered "Fixed Structure Stochastic Automata" (FSSA). Tsetlin, Krylov, and Krinsky [25] presented notable examples of this type of automata.

---

[1]This assumption, however, does not simplify the problem. Indeed, p_R can be assigned to be the smallest value of all the values of p_i for the reliable sensors, and p_U can be assigned to be the largest value of all the values of p_i for the unreliable ones.

[2]Throughout this paper, since we will be invoking majority-like decisions, we assume that $N = N_R + N_U$ is an even number.

[3]If this condition is not satisfied, it means that we are dealing with a system from which no meaningful measurements can be inferred.

Later, Vorontsova and Varshavskii [18] introduced a class of stochastic automata known in the literature as Variable Structure Stochastic Automata (VSSA). The solution we present here, essentially falls within this family and so we shall explain this family in greater detail in Section IV. In the definition of a VSSA, the LA is completely defined by a set of actions (one of which is the output of the automaton), a set of inputs (which is usually the response of the Environment) and a learning algorithm, T. The learning algorithm [18] operates on a vector (called *the Action Probability vector*)

$$P(t) = [p_1(t), \ldots, p_R(t)]^T,$$

where $p_i(t)$ ($i = 1, \ldots, R$) is the probability that the automaton will select the action α_i at time 't',

$$p_i(t) = Pr[\alpha(t) = \alpha_i], i = 1, \ldots, R, \text{ and it satisfies}$$

$$\sum_{i=1}^{R} p_i(t) = 1 \ \forall \ t.$$

Note that the algorithm $T : [0,1]^R \times A \times B \to [0,1]^R$ is an updating scheme where $A = \{\alpha_1, \alpha_2, \ldots, \alpha_R\}$, $2 \leq R < \infty$, is the set of output actions of the automaton, and B is the set of responses from the Environment. Thus, the updating is such that

$$P(t+1) = T(P(t), \alpha(t), \beta(t)),$$

where $P(t)$ is the action probability vector, $\alpha(t)$ is the action chosen at time t, and $\beta(t)$ is the response it has obtained.

IV. THE SOLUTION

A. Overview of Our Solution

In this paper, we provide a novel solution to the $STTP$, based on the field of LA that was briefly surveyed above. We intend to take advantage of the fact that LA combine rapid and accurate convergence with low computational complexity. In addition to its computational simplicity, unlike most reported approaches, as mentioned earlier, our scheme does not require prior knowledge of the ground truth. Rather, it adaptively, and in an on-line manner, gradually learns the identity and characteristics of the sensors which tend to provide reliable readings, and of those which tend to provide unreliable ones.

Our solution involves a team of LA where each LA is uniquely attached to (or rather, associated with) a specific sensor, on a one-to-one basis. Each automaton $\mathcal{A}^i$ attached to sensor s_i, has two actions.

By suitably modeling the agreement or disagreement of the opinions about the sensed ground truth between each sensor and the rest of the other sensors, we can appropriately model these as responses from the corresponding "Environment". Using these synthesized responses, our scheme will intelligently group the sensors according to the readings that they report about the ground truth. Since a sensor is reliable if it reports the ground truth correctly with a probability $p_i > 0.5$ (and unreliable otherwise), we will design our scheme so that it can infer the similar sensors and collect them into their respective groups. In other words, we will infer the crucial sensor identities from the random stream of sensor reports.

The fusion part of our scheme will be based on the result of a prior partitioning phase. Ultimately, the aim behind identifying the set of unreliable sensors, $\mathcal{S}_U$, is to improve the performance of the fusion process for inferring the ground truth. The result of the convergence of the team of LA, which results in a partitioning that infers the identity of the sensor, will serve as an input to the fusion process. In this vein, we shall present two approaches for fusing the results, and study their performances in the section that describes the experimental result. The first fusion approach only considers the measurements from the reliable sensors as being informative, and simultaneously discards measurements from the unreliable sensors. As opposed to this, the second approach attempts to intelligently combine (or fuse) the measurements from both the reliable and the unreliable sensors to yield an accurate value of the ground truth. In this approach, the reading from an unreliable sensor is modified so that it can be considered informative.

The first formal result concerning the performance of the LA is given below.

Theorem 1: Consider the scenario when $(N_R - 1)p_R + N_U p_U > (N_R + N_U)/2$ and when $N_R + N_U - 1 \geq 3$. Let $s_i \in \mathcal{S}_R$. Consider now the agreement between the opinion of a reliable sensor s_i and the opinion of the majority formed by all the rest of the sensors $S\backslash\{s_i\} = (\mathcal{S}_R\backslash\{s_i\}) \cup \mathcal{S}_U$. Let $y_{(N_R-1,N_U)}$ be the decision of a majority voting scheme $S\backslash\{s_i\}$, based on the responses of $N_R - 1$ reliable and N_U unreliable sensors.

Then, if x_i is the output of s_i: $Prob(x_i = y_{(N_R-1,N_U)}) > 0.5$.

Proof: Sketch of Proof The proof is quite involved and so we include only the sketch of the proof in the interest of space and brevity. The complete proof is included in the unabridged version of this paper [27]. The proof relies on considering $\bar{p}_{(N_R-1,N_U)}$, the mean competence of individual s_i in a heterogeneous group $S\backslash\{s_i\} = \mathcal{S}_R\backslash\{s_i\} \cup \mathcal{S}_U$. Then we apply Theorem 4 due to Boland [4], which is an extension of the Condorcet Jury theorem for heterogeneous groups, to demonstrate that:

$$P_{(N_R-1,N_U)} > \bar{p}_{(N_R-1,N_U)} > 1/2.$$

The next key element of the proof is to express $Prob(x_i = y_{(N_R-1,N_U)})$ as a convex function of $P_{(N_R-1,N_U)}$ and p_R. By studying the dynamics of the convex function, we can prove that $Prob(x_i = y_{(N_R-1,N_U)}) > 1/2$.

∎

We shall now consider the converse case of omitting an unreliable sensor, and prove the analogous result.

Theorem 2: Consider the scenario when $(N_R - 1)p_R + N_U p_U > (N_R + N_U)/2$ and when $N_R + N_U - 1 \geq 3$. Let $s_i \in \mathcal{S}_U$. Consider now the agreement between the opinion of an unreliable sensor s_i and the opinion of the majority formed by all the rest of the sensors $S\backslash\{s_i\} = \mathcal{S}_R \cup \mathcal{S}_U\backslash\{s_i\}$. Let $y_{(N_R,N_U-1)}$ be the decision of a majority voting scheme formed of $S\backslash\{s_i\}$, based on the responses of N_R reliable and $N_U - 1$ unreliable sensors.

Then, if x_i is the output of s_i: $Prob(x_i = y_{(N_R,N_U-1)}) > 0.5$.

Proof: The proof follows the same line as the proof of the previous theorem. The proof is found in the unabridged version of this paper [27]. ∎

B. Construction of the Learning Automata

The results that we presented in the previous section form the basis of our LA-based solution. We explain this below, including the strategy by which the majority vote is invoked. In the partitioning strategy, with each sensor s_i we associate a 2-action L_{RI} automaton $\mathcal{A}^i$, $(\Sigma^i, \Pi^i, \Gamma^i, \Upsilon^i, \Omega^i)$, where Σ^i is the set of actions, Π^i is the set of action probabilities, Γ^i is the set of feedback inputs from the Environment, and Υ^i is the set of action probability updating rules.

1) *The set of actions of the automaton*: (Σ^i)
   The two actions of the automaton are α_k^i, for $k \in \{0, 1\}$, i,e, α_0^i and α_1^i

2) *The action probabilities*: (Π^i)
   $P_k^i(n)$ represent the probabilities of selecting the action α_k^i, for $k \in \{0, 1\}$, at step n. Initially, $P_k^i(0) = 0.5$, for $k = 0, 1$.

3) *The feedback inputs from the Environment to each automaton*: (Γ^i)
   Let the automaton select either the the action α_0^i or α_1^i. Then, the responses from the Environment and the corresponding probabilities are tabulated below. For a chosen action, the Environment will respond by a "Reward", or a "Penalty". The conditional probabilities of the "Reward", and "Penalty" are also specified in the tables.

ACTION	ASSOCIATED PROBABILITY	
	REWARD	PENALTY
α_0^i	$Prob(x_i = y_{(N_R-1, N_U)})$	$1 - Prob(x_i = y_{(N_R-1, N_U)})$
α_1^i	$1 - Prob(x_i = y_{(N_R-1, N_U)})$	$Prob(x_i = y_{(N_R-1, N_U)})$

TABLE I: Reward and Penalty probabilities for sensor $s_i \in \mathcal{S}_R$

ACTION	ASSOCIATED PROBABILITY	
	REWARD	PENALTY
α_0^i	$Prob(x_i = y_{(N_R, N_U-1)})$	$1 - Prob(x_i = y_{(N_R, N_U-1)})$
α_1^i	$1 - Prob(x_i = y_{(N_R, N_U-1)})$	$Prob(x_i = y_{(N_R, N_U-1)})$

TABLE II: Reward and Penalty probabilities for sensor $s_i \in \mathcal{S}_U$

A brief explanation about the equations in these tables could be beneficial.
   a) The LA system is rewarded if it chooses action α_0^i, in which case the reading of the sensor s_i agrees with the opinion of the majority voting scheme associated with $S \backslash \{s_i\}$. This occurs with probability $Prob(x_i = y_{(N_R-1, N_U)})$ whenever $s_i \in \mathcal{S}_R$ and with probability $Prob(x_i = y_{(N_R, N_U-1)})$ whenever $s_i \in \mathcal{S}_U$.
   b) Alternatively, the system is rewarded if it chooses action α_1^i, in which case the reading of the sensor s_i disagrees with the opinion of the majority voting scheme associated with $S \backslash \{s_i\}$. This occurs with probability

$1 - Prob(x_i = y_{(N_R-1, N_U)})$ whenever $s_i \in \mathcal{S}_R$ and with probability $1 - Prob(x_i = y_{(N_R, N_U-1)})$ whenever $s_i \in \mathcal{S}_U$.
   c) The penalty scenarios are the reversed ones.

4) *The action probability updating rules*: (Υ^i)
   First of all, since we are using the L_{RI} scheme, we ignore all the penalty responses. Upon reward, we obey the following updating rule : If α_k^i for $k \in \{0, 1\}$ was rewarded then,

$$P_{1-k}^i(n+1) \leftarrow \theta \times P_{1-k}^i(n)$$
$$P_k^i(n+1) \leftarrow 1 - \theta \times P_{1-k}^i(n),$$

   where $0 \ll \theta < 1$ is the L_{RI} reward parameter.

Before we prove the properties of the overall system, we first state a fundamental result of the L_{RI} learning schemes which we will repeatedly allude to in the rest of the paper.

Lemma 1: An L_{RI} learning scheme with parameter $0 \ll \theta < 1$ is ϵ-optimal, whenever an optimal action exists. In other words, $\lim_{\theta \to 1} \lim_{n \to \infty} P_k^i(n) \to 1$.

The above result is well known [14], [18], [21]. By virtue of this property, we are guaranteed that for any L_{RI} scheme with the two actions $\{\alpha_0, \alpha_1\}$, if $\exists \ k \in \{0, 1\}$ such that $c_k^i < c_{1-k}^i$, then the action α_k^i is optimal, and for this action $P_k^i(n) \to 1$ as $n \to \infty$ and $\theta \to 1$, where the $\{c_k^i\}$, are the penalty probabilities for the two actions of the automaton $\mathcal{A}^i$. By invoking the property of the L_{RI} learning scheme, we state and prove the convergence property of the overall system.

Theorem 3: Consider the scenario when $(N_R - 1)p_R + N_U p_U > (N_R + N_U)/2$ and when $N_R + N_U - 1 \geq 3$. If each of the LA in the system uses the L_{RI} scheme with a parameter θ which is arbitrarily close to unity, the following is true:

$$\text{If } s_i \in \mathcal{S}_R, \quad \text{then } \lim_{\theta \to 1} \lim_{n \to \infty} P_1^i(n) \to 1;$$
$$\text{If } s_i \in \mathcal{S}_U, \quad \text{then } \lim_{\theta \to 1} \lim_{n \to \infty} P_0^i(n) \to 1.$$

Proof: The proof of this theorem is a direct consequence of Theorem 1 and Theorem 2. It is omitted here due to the space limitations. ∎

1) Remarks and some Additional Notation: For the case when $N_R p_R + (N_U - 1)p_U > (N_R + N_U)/2$, once the partitioning has taken place, all the LA will have converged to their appropriate partitions. From the results of Theorem 3, we see that the reliable sensors will have converged to action α_1^i, while the unreliable ones will have converged to action α_0^i – both with an arbitrarily large probability.

- Partitioning when $N_R p_R + (N_U - 1)p_U > (N_R + N_U)/2$
  - $G_R = \{s_i \in S \text{ such that } \lim_{n \to \infty} P_1^i(n) = 1\}$
  - $G_U = \{s_i \in S \text{ such that } \lim_{n \to \infty} P_0^i(n) = 1\}$.

Indeed, since the results are ϵ-optimal results, if θ is not arbitrarily close to unity, some of the LA might fail to converge to the optimal action and thus the set $\mathcal{S}_R$ may not necessary be equivalent to G_R. However, as θ is arbitrarily close to unity, G_R will converge exactly to $\mathcal{S}_R$.

C. Fusion approaches

We now present two simple fusion schemes that make use of the partitionings in order to improve the quality of the aggregated opinion from the different sensors for guessing the ground truth.

1) Fusion Scheme with Exclusion: Discarding the opinions of the unreliable sensors: A possible strategy to increase the accuracy of the fusion process is to employ a simple majority voting strategy that excludes all the sensors whose LA converged to the action G_U during the partitioning phase. This means that the prediction of the ground truth will be exclusively based on the sensors whose LA converged to the action G_R.

2) Fusion Scheme with Inversion: Inverting the opinions of the unreliable sensors: In this subsection, instead of excluding the readings of the unreliable sensors, we propose intelligently combining the readings from both the reliable and unreliable sensors when evaluating the ground truth. In fact, we opt to invert the decision of the unreliable sensors as inferred by the LA algorithm, rendering them to be informative. Thus, for every reading x_i from a sensor s_i whose LA has converged to the action G_U, we record the reverse of the reading.

Indeed, the majority voting scheme will be equivalent to one that aggregates the votes from a group of sensors consisting of:

- N_R reliable sensors, each possessing a correctness probability p_R, and

- N_U unreliable that have been rendered reliable and that possess a correctness probability $p_U' = 1 - p_U$ (where $p_U' = 1 - p_U > 0.5$). By the phrase, rendered reliable, we mean that we are inverting the respective readings of the sensors in G_U .

We now report the experimental results that we have obtained by testing the strategies explained in the previous sections.

V. EXPERIMENTAL RESULTS

The performance of the LA-based partitioning as well as the two fusion schemes that make use of the partitioning, have been rigorously tested by simulation in a variety of parameter settings, and the results that we have obtained are truly conclusive. In the interest of brevity, we merely report a few representative (and typical) experimental results, so that the power of our proposed methodology can be justified. In the experiments, the settings were chosen so that the condition $N_R p_R + (N_U - 1) p_U > (N_R + N_U)/2$ was met, reflecting the phenomenon where "the truth prevails over lying".

A. Performance of the Partitioning

We first examine the convergence speed of the LA algorithm. Since a LA is associated with every sensor (whether it is reliable or unreliable), where each possesses its own distinct reward probabilities for its respective actions, they will, clearly, have different convergence speeds, as is well-known in the theory of LA. Observe that the convergence of the individual LA is defined in terms of its ϵ-convergence, where the LA were

(p_R, p_U)	Average Convergence time for $s_i \in S_R$	Average Convergence time for $s_i \in S_U$
$(0.8, 0.1)$	40.91	43.37
$(0.8, 0.2)$	36.41	44.11
$(0.85, 0.1)$	31.16	30.84
$(0.85, 0.2)$	29.53	35.82
$(0.9, 0.1)$	26.14	26.71
$(0.9, 0.2)$	25.90	32.81
$(0.95, 0.1)$	23.51	25.88
$(0.95, 0.2)$	23.47	32.27

TABLE III: Average convergence time for the case when $(N_R, N_U) = (20, 10)$.

deemed to have converged if one of its action probabilities attained the value $1 - \epsilon$[4]. Formally:

- If $P_0^i(n) \geq 1 - \epsilon$, then the LA has converged to the action α_0^i;

- If $P_1^i(n) \geq 1 - \epsilon$, then the LA has converged to the action α_1^i.

We also initialized all the LA at time instant $t = 0$, to have the values: $P_0^i(t) = P_1^i(t) = 0.5$. To render the results meaningful, we took an ensemble average of $1,000$ experiments, and computed the average convergence times for the LA associated with the sensors in S_R and for those in S_U. Although the experiments related to the convergence speeds were performed for different settings, we only report some representative results in which we fixed N_R to 20, N_U to 10 and $\theta = 0.8$, and where we also simultaneously varied p_R and p_U. In fact, it turns out that these parameters will influence the agreement probability (reward probability), and consequently the speed of convergence as per the theoretical results reported earlier. The results obtained are given in Table III.

By examining this table, we observe:

1) Remarkably, the LA converge very rapidly. In fact, on the average, the LA were able to determine the optimal partition in less than 44.11 time instances, which, incidentally, was the largest value in the table.
2) Earlier, we proved that the probability of a reward is a decreasing function of p_U whenever we deal with an unreliable sensor. As we fix p_R and vary p_U, we observe that the convergence speed decreases, which, in this case, translates into a decreased reward probability.
3) In addition, we p_R is increased towards unity and as p_U is decreased closer to 0, the convergence speed increases for both the individual LA and for those included in S_R. This reflects the concept that the environment becomes "easier" when the sensor is less noisy (i.e, (p_R, p_U) approaches $(0, 1)$) and consequently, the LA converge faster to the optimal actions. By "easier", we mean that the difference between the reward probabilities of the actions of the LA becomes larger, and thus, the LA will converge both faster and with a higher probability to the optimal action. This is consistent with the well-known results in the field of LA.
4) Consider the case when $(p_R, p_U) = (0.95, 0.1)$ as reported in the table. The respective convergence speeds

---

[4]The value of ϵ was set to be 0.01.

(p_R, p_U)	$P(C_C)$ for Fusion Scheme with Exclusion	$P(C_C)$ for MV for all sensors
(0.75, 0.45)	0.921	0.766
(0.75, 0.4)	0.921	0.87
(0.75, 0.35)	0.921	0.599
(0.75, 0.3)	0.921	0.5
(0.8, 0.45)	0.972	0.84
(0.8, 0.4)	0.972	0.775
(0.8, 0.35)	0.972	0.574
(0.8, 0.3)	0.9672	0.604

TABLE IV: Comparisons of the value of $P(C_C)$, the probabilities of the consensus being correct for different values of (p_R, p_U), and for the different approaches for $N_R = 10$ and $N_U = 10$.

(p_R, p_U)	$P(C_C)$ for Fusion Scheme with Exclusion	$P(C_C)$ for MV for all sensors
(0.75, 0.45)	0.9986	0.943
(0.75, 0.4)	0.9986	0.919
(0.75, 0.35)	0.9986	0.888
(0.75, 0.3)	0.9986	0.85
(0.8, 0.45)	0.997	0.98
(0.8, 0.4)	0.997	0.97
(0.8, 0.35)	0.997	0.954
(0.8, 0.3)	0.997	0.934

TABLE V: Comparisons of $P(C_C)$, the probabilities of the consensus being correct for different values of (p_R, p_U), and for the different approaches for $N_R = 20$ and $N_U = 10$.

for the LA associated with the reliable and unreliable sensors are 23.51 and 25.88 respectively. However, as the sensors became more noisy by decreasing p_R to 0.8, the task of differentiating between the partitions became more difficult. Indeed, the convergence speed for LA associated with a reliable sensor dropped down to 36.41, and the speed of the LA associated with unreliable sensor became 44.11.

B. Fusion Scheme with Exclusion

We now compare the "Fusion Scheme with Exclusion" explained in Section IV-C1 with the deterministic Majority Voting (MV) strategy that incorporates all the sensors in S. As detailed earlier, the latter scheme relies exclusively on the decision of the vote of the majority of the sensors that converged to the G_R partition. Let $P(C_c)$ denote the probability of the consensus being correct, i.e, that the probability that the vote of the majority coincides with the ground truth. Table IV reports the result of the comparison for the case when N_R and N_U are both equal to 10. We observe from the table:

1) As one can see, the results we report are conclusive. In fact, we were able to increase the value of $P(C_c)$ quite remarkably. For example, for the case when $(p_R, p_U) = (0.75, 0.3)$, our scheme yielded a value of 0.921 for $P(C_C)$, while the scheme which operated with the majority voting involving all the sensors yielded the value of only 0.5.

2) The value of $P(C_C)$ for our Fusion Scheme with Exclusion is immune to the variation of p_U. For example, for the entries corresponding to $p_R = 0.75$, we see that $P(C_C)$ is equal to 0.921 even if p_U changes, for example, by taking the values 0.45, 0.35 and 0.3.

Consider now the case when we double the value N_R from 10 to 20 while the value of N_U is equal to 10. As expected, we see from Table V, the value of $P(C_C)$ for our scheme increases and approaches unity.

C. Fusion Scheme with Inversion

In Table VI, we report the results when we fix N_R to 20 and N_U to 10 and compare the result of a simple MV scheme involving all sensors with the Fusion Scheme with Inversion presented in Section IV-C2. We can make the following observations: Under a fixed value of p_R, a smaller value of p_U yields a higher value for $P(C_C)$ for the Fusion Scheme

with Inversion. For example, for a fixed value of $p_R = 0.8$, $P(C_C)$ increases from 0.929 to 0.986 as we decrease p_U from 0.45 to 0.3. This is due to the fact that a smaller value for p_U actually implies a higher value for $1 - p_U$. Thus, a sensor which is highly unreliable, can be transformed into one that is highly reliable – thanks to the operation of inverting its reading! The results for the case where we increase N_U to 20 is reported in Table VI. Indeed, in general we can affirm from Tables VI and VII that the Fusion Scheme with Inversion outperforms the simple majority voting involving all sensors in all the settings. However, by comparing both Tables VI and VII, we remark that $P(C_C)$ for the scheme with inversion does not necessarily increase as we increase N_R, the number of reliable sensors.

(p_R, p_U)	$P(C_C)$ for Fusion Scheme with Inversion	$P(C_C)$ for MV for all sensors
(0.75, 0.45)	0.974	0.943
(0.75, 0.4)	0.984	0.919
(0.75, 0.35)	0.99	0.888
(0.75, 0.3)	0.994	0.85
(0.8, 0.45)	0.992	0.98
(0.8, 0.4)	0.995	0.97
(0.8, 0.35)	0.997	0.954
(0.8, 0.3)	0.998	0.934

TABLE VI: Comparisons of $P(C_C)$ the probabilities of the consensus being correct for different values of (p_R, p_U), and for the different approaches for $N_R = 20$ and $N_U = 10$.

(p_R, p_U)	$P(C_C)$ for Fusion Scheme with Inversion	$P(C_C)$ for MV for all sensors
(0.75, 0.45)	0.974	0.8821
(0.75, 0.4)	0.985	0.804
(0.75, 0.35)	0.994	0.699
(0.75, 0.3)	0.997	0.571
(0.8, 0.45)	0.987	0.941
(0.8, 0.4)	0.994	0.892
(0.8, 0.35)	0.998	0.816
(0.8, 0.3)	0.998	0.71

TABLE VII: Comparisons of $P(C_C)$ the probabilities of the consensus being correct for different values of (p_R, p_U), and for the different approaches for $N_R = 20$ and $N_U = 20$.

VI. CONCLUSION

Sensor Fusion has become a prevalent research topic due to the wide deployment of sensor technology in the industry and in our daily life. In this paper, we have considered an extremely pertinent problem in the area of Sensor Fusion, namely the one of identifying unreliable sensors without knowing the ground truth. Although paradigms like the one that involves majority voting offer a generic prediction strategy for the ground truth

by aggregating sensor-provided readings, they are prone to error caused by unreliable sensors. Clearly, such unreliable sensors may degrade the quality of the aggregated information.

A large body of the research in sensor fusion deduces the reliability of the sensors either online or offline by assuming that one can access the ground truth. While this is a desirable option, unfortunately, in real life, such an assumption does not always hold. The question of whether a solution to the problem even exists in this scenario is open. In this paper, we have presented a novel solution for the problem using tools provided by the family of Learning Automata (LA). Unlike most reported approaches, our scheme does not require prior knowledge of the ground truth. Instead, our solution gradually learns the identity and characteristics of the sensors which provide reliable readings, and of those who provide unreliable measurements.

In addition to presenting rigorous theoretical results for the unsolved problem, we have also included comprehensive empirical results that demonstrate that our LA-based scheme achieves optimal partitioning with a high convergence speed.

A possible extension of this research, which we are currently working on, is to develop the analogous methodology for continuous sensor readings instead of boolean ones. In addition, we advocate that it is possible to render the two phases of partitioning and fusion to be interleaving by using the information contained in the all N vectors $P^i(t) = [P_0^i(t), P_0^i(t)]$ at time n. Thus, the fusion can take place at each time instant t, instead of delaying or postponing the execution of the proposed fusion (with/without Inversion/Exclusion) until all the LA have converged.

The entire issue of whether we can use the field of Random Races [19] to achieve a comparison between the various sensors also holds a great potential. Finally, the question of investigating the effect of adding more voters on $P(C_C)$ has not been considered here. Some details about this scenario can be found in [9].

REFERENCES

[1] M. Agache and B. J. Oommen. Generalized pursuit learning schemes: New families of continuous and discretized learning automata. *IEEE Transactions on Systems, Man, and Cybernetics-Part B: Cybernetics*, 32(6):738–749, December 2002.

[2] H. Altincay. On the independence requirement in dempster-shafer theory for combining classifiers providing statistical evidence. *Applied Intelligence*, 25:73–90, 2006.

[3] P. L. Bogler. Shafer-dempster reasoning with applications to multisensor target identification systems. *IEEE Transactions on Systems, Man and Cybernetics*, 17(6):968–977, 1987.

[4] P. J. Boland. Majority systems and the condorcet jury theorem. *The Statistician*, pages 181–189, 1989.

[5] S. Buchegger and J. Le Boudec. A robust reputation system for P2P and mobile ad-hoc networks. In *Proceedings of the Second Workshop on the Economics of Peer-to-Peer Systems*, 2004.

[6] C. Dellarocas. Immunizing online reputation reporting systems against unfair ratings and discriminatory behavior. In *Proceedings of the 2nd ACM conference on Electronic commerce*, pages 150–157, Minneapolis, Minnesota, United States, 2000. ACM.

[7] W. Elmenreich. Fusion of continuous-valued sensor measurements using confidence-weighted averaging. *Journal of Vibration and Control*, 13(9-10):1303–1312, 2007.

[8] K. Faceli, A. de Carvalho, and S. Rezende. Combining intelligent techniques for sensor fusion. *Applied Intelligence*, 20:199–213, 2004.

[9] S. L. Feld and B. Grofman. The accuracy of group majority decisions in groups with added members. *Public Choice*, 42(3):273–285, 1984.

[10] J. Frolik, M. Abdelrahman, and P. Kandasamy. A confidence-based approach to the self-validation, fusion and reconstruction of quasi-redundant sensor data. *IEEE Transactions on Instrumentation and Measurement*, 50(6):1761–1769, 2001.

[11] D. L. Hall and J. Llinas. An introduction to multisensor data fusion. *Proceedings of the IEEE*, 85(1):6–23, 1997.

[12] M. A. Hossain, P. K. Atrey, and A. El Saddik. Learning multisensor confidence using a reward-and-punishment mechanism. *IEEE Transactions on Instrumentation and Measurement*, 58(5):1525–1534, 2009.

[13] A. Jøsang. Artificial reasoning with subjective logic. In *Proceedings of the Second Australian Workshop on Commonsense Reasoning*, volume 48. Perth, 1997.

[14] S. Lakshmivarahan. *Learning Algorithms Theory and Applications*. Springer-Verlag, 1981.

[15] N. Littlestone and M. K. Warmuth. The weighted majority algorithm. *Inf. Comput.*, 108(2):212–261, 1994.

[16] K. Marzullo. Tolerating failures of continuous-valued sensors. *ACM Transactions on Computer Systems (TOCS)*, 8(4):284–304, 1990.

[17] J. Mundinger and J.-Y. Le Boudec. Analysis of a reputation system for mobile Ad-Hoc networks with liars. *Perform. Eval.*, 65(3-4):212–226, 2008.

[18] K. S. Narendra and M. A. L. Thathachar. *Learning Automata: An Introduction*. Prentice-Hall, New Jersey, 1989.

[19] D. T. H. Ng, B. J. Oommen, and H. E. R. Adaptive learning mechanisms for ordering actions using random races. *IEEE Transactions on Systems, Man and Cybernetics*, 23(5):1450–1465, 1993.

[20] A. S. Poznyak and K. Najim. *Learning Automata and Stochastic Optimization*. Springer-Verlag, Berlin, 1997.

[21] Y. Sawaragi and N. Baba. A note on the learning behavior of variable-structure stochastic automata. *IEEE Transactions on Systems, Man and Cybernetics*, 3:644–647, 1973.

[22] S. Sen and N. Sajja. Robustness of reputation-based trust: boolean case. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1*, pages 288–293, Bologna, Italy, 2002. ACM.

[23] M. A. L. Thathachar and P. S. Sastry. Varieties of learning automata: An overview. *IEEE Transactions on Systems, Man, and Cybernetics-Part B: Cybernetics*, 32(6):711–722, December 2002.

[24] M. A. L. Thathachar and P. S. Sastry. *Networks of Learning Automata: Techniques for Online Stochastic Optimization*. Kluwer Academic, Boston, 2003.

[25] M. L. Tsetlin. *Automaton Theory and the Modeling of Biological Systems*. Academic Press, New York, 1973.

[26] A. Yazidi, O.-C. Granmo, and B. J. Oommen. Service selection in stochastic environments: a learning-automaton based solution. *Applied Intelligence*, 36(3):617–637, 2012.

[27] A. Yazidi, B. J. Oommen, and M. Goodwin. On solving the problem of identifying unreliable sensors without a knowledge of the ground truth: The case of stochastic environments. *Unabridged journal version of this paper, 2015. To be submitted for publication.*

[28] B. Yu and M. P. Singh. Detecting deception in reputation management. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 73–80, Melbourne, Australia, 2003. ACM.

[29] G. Zacharia and P. Maes. Trust management through reputation mechanisms. *Applied Artificial Intelligence*, 14(9):881–907, 2000.