

GiBy: A Giant-Step Baby-Step Classifier For Anomaly Detection In Industrial Control Systems

Sarad Venugopalan¹ and Sridhar Adepu²

¹School of Computer Science, University of Bristol

²Department of Computer Science, Swansea University

Abstract—The continuous monitoring of the interactions between cyber-physical components of any industrial control system (ICS) is required to secure automation of the system controls, and to guarantee plant processes are fail-safe and remain in an acceptably safe state. Safety is achieved by managing actuation (where electric signals are used to trigger physical movement), dependent on corresponding sensor readings; used as ground truth in decision making. Timely detection of anomalies (attacks, faults and unascertained states) in ICSs is crucial for the safe running of a plant, the safety of its personnel, and for the safe provision of any services provided. We propose an anomaly detection method that involves accurate linearization of the non-linear forms arising from sensor-actuator(s) relationships, primarily because solving linear models is easier and well understood. Further, the time complexity of the anomaly detection scenario/problem at hand is lowered using dimensionality reduction of the actuator(s) in relationship with a sensor. We accomplish this by using a well-known water treatment testbed as a use case. Our experiments show millisecond time response to detect anomalies and provide explainability; that are not simultaneously achieved by other state of the art AI/ML models with eXplainable AI (XAI) used for the same purpose. Further, we pin-point the sensor(s) and its actuation state for which anomaly was detected.

Index Terms—Anomaly Detection, Cyber-Attacks, Safety, Security, Cyber-Physical System, Industrial Control System.

I. INTRODUCTION

Physical/mechanical automation components are more recently becoming digitized and analogue systems are also being developed with digital interfaces to enable integration with digital systems. Industries such as manufacturing (including assembly, quality check and distribution) and critical infrastructure (CI) - including the water sector (water and wastewater treatment, and distribution), energy sector (generation, transmission and distribution), and transportation sector, play an important role in global society. The provision of high-quality CI services in particular should be noted as significant for public health, well-being and safety. The integration and embedding of computational processing units expose the ICS used in industries and CI to an expanded attack surface, making it vulnerable to cyber-attacks [1]. For example, malware was found in the manufacturing sector focused on a steel plant [2]. This attack led to an unregulated furnace and the resulting dangerous infrastructure could not be shut down as normal, leading to physical damage to the plant [3]. Recent power sector-focused malware such as FrostyGoop [4] were able to infiltrate CI and disrupt services. Disruption and damage have been recorded in the water sector due to attacks

on CPS at the Maroochy Shire plant in Australia [5], multiple plants in Texas, USA [6], Mayo in Ireland [7], and a mitigated attack at the Oldsmar plant in USA [8]. The spate of real-world attacks on manufacturing and CI operational technology (OT), and their consequences reiterate the value in timely detection of these attacks and subsequent actions to bring these systems into a safe and non-compromised state.

In this work, we use the Secure Water Treatment (SWaT) [9] testbed as a use case in ICS anomaly detection. We use system-specific information to improve anomaly detection, as opposed to other models that employ specialized solutions on a subset of available information [10]–[12], or prefer to use a generic data-driven approach [13]. Specifically, we rely on *i.*) the real-time data logs to retrieve state information from plant sensors and actuators, *ii.*) the PLC control logic sheet to determine sensor boundary conditions and corresponding actuation steps (if any), and *iii.*) infer sensor-actuator(s) relationships from the plant system architecture.

Going over the complete actuation state space is a case of solving problems in exponential time (see Section II-B). We employ a dimensionality reduction technique to reduce the time complexity by considering only the nearest-neighbour actuators related to a given sensor. This is followed by an accurate linearization of a reduced number of actuator(s)-sensor relationships. We note that the water treatment plant (see Fig. 3), has limited nonlinear sensor-actuator(s) relationships. This is attributed to the (mostly) serial nature of numerous industrial processes, including those involving assembly lines in a factory setting and a water treatment plant. Once the sensor-actuator(s) relationships are linearized, it is trivial to determine the bounds required to detect anomalies in the ICS. The advantage of this approach is that it typically allows detection and explanation within a millisecond (see Section V). It saves valuable time with fast detection and explanation, helping with faster diagnostics, mitigation, or recovery. The novelty being that our solution is fit for use in near real-time and resource-constrained decision-making control systems (see Section VI-F). We make the following main contributions.

- i.*) We present our anomaly detector in Section III.
- ii.*) An extended solution is proposed in Section IV to detect rare event anomalies.
- iii.*) We compare our solution results with AI/ML models for anomaly detection and explainable AI (see Section V).
- iv.*) We discuss how explainability [14] is incorporated into our solution (see Section VI-E).

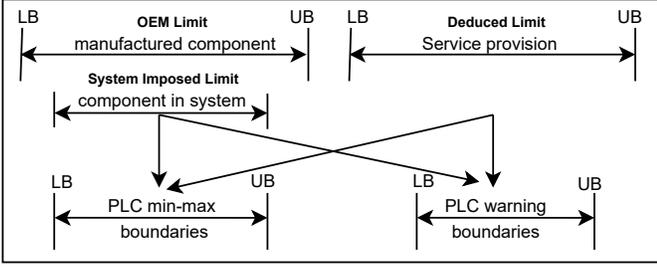


Fig. 1. Safe state determination using sensors. Boundaries are fed into a programmable logic controller. Lower bound (LB) and upper bound (UB).

The paper is organised as follows. Section II provides the problem context and discusses the threat model. Section III presents the detector solution. The solution is extended in Section IV. Experiments are carried out and results presented in Section V. The work is discussed in Section VI. The related work is provided in Section VII. The conclusions are drawn, and future work is proposed in Section VIII.

II. PROBLEM CONTEXT AND THREAT MODEL

A. Sensor readings and boundary conditions

In context, boundaries refer to safety limits for plant operations. These limits (see Fig. 1) apply to (i) transport components' operational characteristics (such as pump speed and pipe pressure), and (ii) service provisions (such as the chemicals required to maintain the hydrogen ion concentration (pH) and oxidation-reduction potential (ORP) for water, in the required safe range for consumption). The distinction between transport characteristics and service provisioning is made based on requirement. For example, pipes and pumps may be used to safely transport a variety of fluids, whereas service provisions for water treatment implies that provisioning used is specific for water quality and consumption safety. A service is most often provided using a conglomeration of different transport components and service provisions.

Transportation components from reputable manufacturers are supplied with detailed specifications, defining the acceptable safe operational range and conditions of use. When used as part of a system, a new appropriate system-imposed limit may apply depending on the system design and other operational bottlenecks. For example, a pipe may be Original Equipment Manufacturer (OEM) safety rated for a pressure of 0-5 Bars. However, downstream constraints may only allow a pressure of 1-3 Bars. In such cases, a system-imposed limit is put in place. With respect to service provisioning boundaries, they are deduced from extensive knowledge and tests to determine the safe range for consumption. For example, a water treatment process must ensure that adding acid or base chemicals do not significantly result in the pH of water moving away from 7 (neutral).

The OEM/system imposed limits for transport components and deduced limits for service provisioning parts are fed into one or more controllers, such as a PLC (see Fig. 1). The min-to-max operational range values on the PLC assists with taking action and (possibly) alerting an operator, when

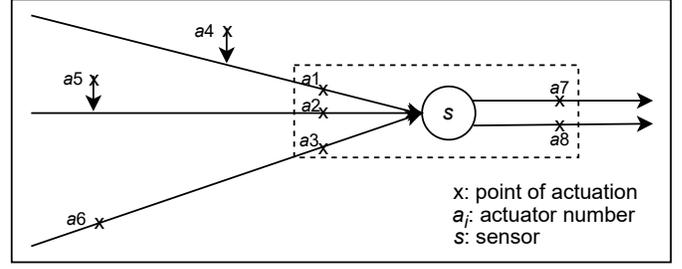


Fig. 2. A sensor and dependent series-parallel actuators in a physical process. Its nearest neighbor actuator(s) are shown inside the dotted rectangle.

the safety threshold is breached. Additionally, the PLCs may also include logic to act, even before a breach, by providing a warning buffer. This may be achieved by setting tighter boundary conditions. Traditionally, at least two layers of safety are provided, the lower boundary is the warning boundary which alerts operators to take action, the upper boundary is the min-max boundary, and this should never be passed. In some systems if the min-max boundary is passed the personnel must evacuate the site and services are put on hold until safe conditions are assured.

B. Reducing time complexity: core idea

In this section, we explain the core idea on which we base our solution. The set of boundaries [LB, UB] in anomaly detection, used as safe range for operations may not always be trivial to optimally determine. Taking into consideration all actuators together causes a time complexity blow up, resulting in a longer time to solve the problem optimally. For example, let each actuator have k actuation sequences (such as open/transition/close; here $|k| = 3$), and let there be n such actuators, in relation with a given sensor s . Hence, there are a total of k^n actuation states, in relation with sensor s . In Fig. 2, the readings from a sensor s may change due to actuation in plant processes, at points marked X. The actuators at the input side are named $a_1 - a_6$, and the output side are $a_7 - a_8$. Now, if we can reduce the state space for sensor relations to its immediate neighbour actuators ($a_1 - a_3$ and $a_7 - a_8$; see dotted rectangle in Fig. 2), we reduce our solution seek time. Therefore, we employ a dimensionality reduction of the actuator(s) in relation with sensor s .

Let there be normal and attack plant operation datasets for the sensor and actuator states logged during plant operations, under normal and attack conditions, respectively. A normal training dataset comprises of all seen (normal) state relations embedded into it. We implicitly rely on all actuator relations $f(a_1, \dots, a_8)$ w.r.t. the sensor, but explicitly consider the actuation sequences for $a_1 - a_3$, for input side linearization. The excluded actuators $a_4 - a_6$ are indirectly responsible for the change in measurement readings at sensor s . However, these

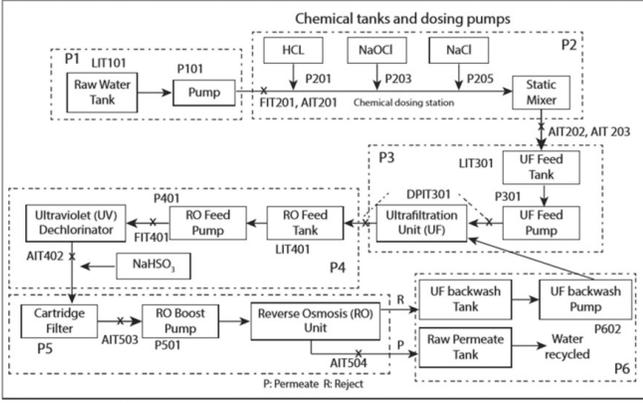


Fig. 3. High level system architecture of SWaT water treatment testbed [9].

relations are captured at the nearest neighbour (nn-) actuators¹ $a_1 - a_3$, and mapped to sensor readings via $f(a_1, \dots, a_8) = s_{reading}$, provided in the normal dataset. Note, only the input actuator state space is reduced. The sensor reading remains a function of all related actuators. Sensor reading is readily available from the normal dataset and its measured range, w.r.t. a_1, \dots, a_8 actuation remains unchanged. In short, we refined our search space to the sensor's nn-actuator(s). However, the set of normal operations seen in the normal dataset might be a tiny subset of the entire normal state space. In this context, everything not seen as part of the normal plant operation dataset and derivations are treated as anomalous (similar to the concept used in one-class classification [15], [16]), and our detector flags it as a warning. The detector warnings provide an additional layer of safety along with the PLC boundaries discussed in section II-A, by generating a suitable alert when action is required.

C. SWaT Water Treatment Testbed and datasets

The SWaT water treatment (see Fig. 3) is a six-stage process (P_1 to P_6) and emulates a typical treatment plant. The water for treatment is the input to P_1 and treated water is the main output of P_6 . Different (and redundant) sensors are strategically placed to make measurements across the testbed. Actuators are used for control. The sensors and actuators used in $P_1 - P_6$ appear in Table I. Process P_1 is designed to get the raw water into the system. P_2 is focused on chemical dosing and P_3 with ultra-filtration. P_4 includes UV-dechlorination and P_5 carries out reverse osmosis (RO). P_6 uses a backwash tank for cleaning and stores treated water in another tank (that is recycled). The device names consist of two parts where the first 3 characters are shortened for devices type and the last 3 numbers refer to the process stage and number of devices. For example, the device FIT-101 refers to the flow meter in the first process where it is the first flow meter.

¹Henceforth, it is referred to as nn-actuators. It is common to place the actuator immediately preceding the related sensor. If not, we yet consider that directly related actuator as a nearest neighbour — to improve detector resolution, when that information is available. Resolution is further enhanced by picking non-neighbour PLC imposed actuation, affecting the sensor reading; as part of the design (see Section VI-D).

This description of each device can be found as follows: FIT-XXX: Flow meter, LIT-XXX: Level Transmitter, AIT-XXX: Analyser, DPIT-XXX: Differential pressure indicating transmitter, PIT-XXX: Pressure meter, MV-XXX: Motorized valve, P-XXX: Pump and UV-XXX: UV Dechlorinator [9]. The dataset was collected for eleven days, of which the first seven days is under normal working conditions. The final four days included injecting/ swapping with anomalous data, involving 41 CPS attacks on the SWaT water testbed. The plant was run from an empty to a fully operational state. The attacks were conducted by altering the OT network traffic, spoofing the sensor values, and issuing bogus SCADA commands.

Process	Sensor	Actuator
P1	LIT-101, FIT-101	MV-101, P101
P2	AIT-201, AIT-202, AIT-203, FIT-201	MV-201, P-201, P-202, P-203, P-204, P-205, P-206
P3	DPIT-301, FIT-301, LIT-301	MV-301, MV-302, MV-303, MV-304, P-301, P-302
P4	AIT-401, AIT-402, FIT-401, LIT-401	P-401, P-402, P-403, P-404, UV-401
P5	AIT-501, AIT-502, AIT-503, AIT-504, FIT-501, FIT-502, FIT-503, FIT-504, PIT-501, PIT-502, PIT-503	P-501, P-502
P6	FIT-601	P-601, P-602, P-603

TABLE I
THE SENSORS AND ACTUATORS IN EACH PROCESS STAGE [12].

The SWaT datasets provided the state view of the fifty-one sensors and actuators in the plant at one second resolution, as time series data. A state may include actuators in either open/close, on/off or open/transition/close state, depending on the type of the actuator used. The normal dataset consisted of 495,000 records of plant state under normal operations. The attack dataset had 449,919 records. Note that the attack dataset is a mix of normal and attack records. The attack records accounted for where less than 6% of total data [12]. If the plant was under attack during the record generation (within the dataset) then this record is tagged as *Attack*, and otherwise as *Normal*.

D. Threat Model

An attacker is assumed to have followed an attack vector [17] to gain access into plant operational technology. This adversary is assumed to be able to inject control commands and modify sensor readings and actuator status in OT. This leads to four attack types [9] — single stage single point (SSSP) where exactly one sensor/actuator data is under attack, single stage multi-point (SSMP) where multiple sensor/actuator data in exactly one process is under attack, multistage single point (MSSP) where exactly one sensor/actuator data in multiple processes are under attack, and multistage multi-point (MSMP) where multiple sensor/actuator data in multiple processes are under attack. It is assumed there are no attacks on the plant when the normal dataset is collected, which would be used as the training dataset.

III. GIANT-STEP BABY-STEP SOLUTION

A. Design rationale and considerations

The sensor and actuator values in the training dataset are viewed as ground truth data. Our core requirements [18]

are (R1) near real-time detection, (R2) anomaly cause identification, and (R3) minimising false alarms. An automated response is not a core consideration because governance considerations involve human intervention and human-in-the-loop solutions [19], [20] to be qualified as compliant to safety standards i.e. a fully automation solution might not be viable. While the ability to automate is provided, the level of automation is left to the discretion of the implementer. One-class classification is used due to the typical imbalance between normal and attack dataset [21]; the latter contributing only a tiny percentage of attack data. We train using the normal dataset and employ unsupervised learning [22].

B. Solution outline

To meet the requirements in Section III-A, we use binary classification to solve a decision problem [23]. The proposed anomaly detector has two main steps — giant and baby step. The giant-step is used to determine $[LB, UB]$ boundary for measured sensor values (available in normal dataset), in relation with every nn-actuation state. For example, consider the more detailed SWaT process stages as illustrated in Fig. 4. In this system, water flows into the tank when motorized valve (actuator) MV101 transitions from close to open. The tank level indicator (sensor) LIT101 measures the water tanks instantaneous water level. When pump (actuator) P101 is turned on, water is pumped out from the tank, and the flow transmission (sensor) FIT201 registers non-zero flow values. The two nn-actuators for sensor LIT101 are MV101 and P101. MV101 has three actuation states (close-1/transition-0/open-2) and P101 has two actuation states (on-2/off-1), in combination this allows for six possible states across the two actuators connected to the sensor LIT101. The $[LB_i, UB_i]_{i=1, \dots, 6}$ bounds for LIT101 is computed by our giant-step solution for all actuator combinations of MV101 and P101: 11, 12, 01, 02, 21, and 22.

The baby-step computes the $[LB, UB]$ boundary based upon the *rate of change* in sensor reading(s) associated with its nn-actuator(s). Again, consider the level indicator sensor LIT101 in Fig. 4. Instead of directly measuring the instantaneous level value of water in the tank, here we measure the rate of change of water level between consecutive time intervals. As seen earlier, we find $[LB_i, UB_i]_{i=1, \dots, 6}$ bounds for LIT101 in relation to actuator combinations of MV101 and P101. These measurements form the basis of the min-max bounds for the rate of change of water level; for each of the six actuation states w.r.t. the level sensor. The combination of the proposed giant-step and baby-step methods is not only able to detect sensor values that are out of bound for dependent actuator states but also detect whether the rate of change of sensor value is also within bounds. When the readings from the sensor under observation are beyond determined thresholds/bounds, the detector issues a warning event to a suitable log; to be processed and considered as a trigger for the generation of an alert to system operators.

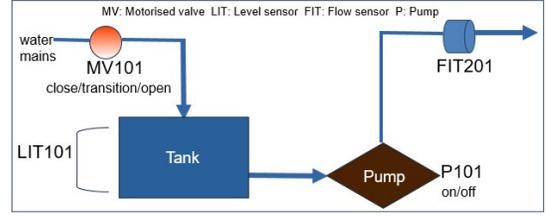


Fig. 4. Mid-granular level view of P_1 and early P_2 SWaT process stages.

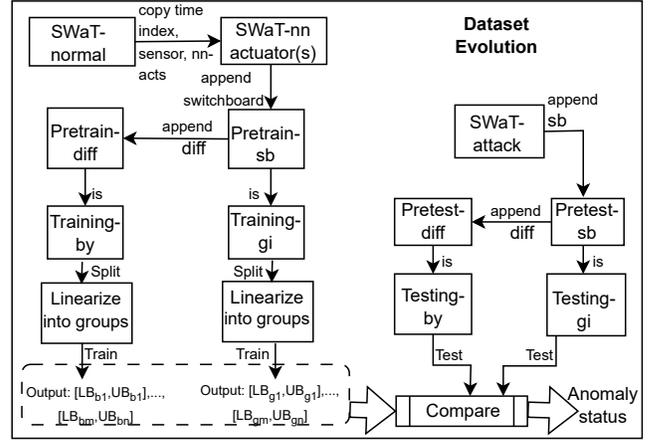


Fig. 5. Evolution of training and testing dataset. Pretest, testing and comparison are carried out in real-time, once the output from training is available.

C. A switchboard to map actuation state, building linearized state groups, and enable explanation

The training dataset is built on the normal SWaT dataset and the testing dataset from the SWaT attack dataset (see Fig. 5). Dataset prefixed with *example* are hypothetical and are used to drive the core idea (see TABLE II and III).

1) *Switchboard*: A switchboard is used to linearize actuation states. It is also part of a bijective map between the data output from training and the testing (sensor) value. A switchboard state (*sb*) is a numerical string. It uniquely identifies the nn-actuator(s) state for the sensor, at time index

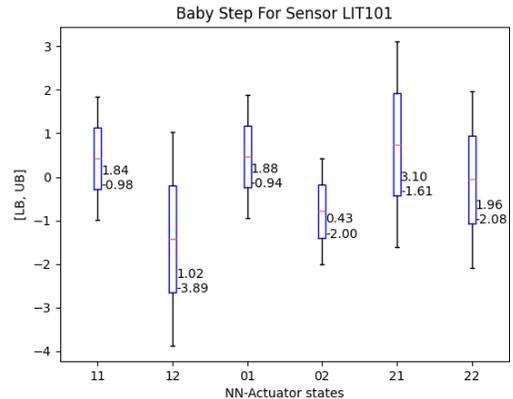


Fig. 6. Plotted six $[LB, UB]$ baby-step training bounds for sensor LIT101. SWaT normal dataset is used for training.

t , and is expressed as a concatenation of the nn-actuator state. Consider two nn-actuators, one with 3 possible states (off = 1, transition = 0, on = 2), and the other with 2 possible states (off = 1, on = 2). This results in six switchboard states — 11, 12, 01, 02, 21, and 22. First, we parse through the normal training dataset and copy the time index, sensor reading, nn-actuators (see first four columns in example TABLE II). Next, we append a switchboard state (for nn-actuator(s)), at each time index t . This is called the pretrain-sb dataset (see additional column labeled *Switchboard* in TABLE II). Next, for the baby-step method, we compute the difference (diff) between consecutive sensor values and append it to a new column. We call this *Pretrain-diff* dataset (see additional column labeled *LIT101Diff* in TABLE II). This is also the *Training-by* dataset. For the giant-step method, the direct sensor values are used. Hence, *Training-gi* and *Pretrain-sb* datasets are identical (see Fig. 5).

2) *Linearized state groups*: Further, we split the *Training-by* dataset into linearized groups (LSG_{sb}), for all $sb \in Training-by$ dataset. The first group contains only dataset rows associated with the $sb = 11$ switchboard state for this example, in the exact order it appeared in the *Training-by* dataset. The remaining groups are determined in a similar manner. We are now able to study the behaviour within each linearized group. The output of passing the training dataset through 1) *Switchboard*, and 2) *Linearized state groups*, are a set of $[LB, UB]$ state bounds, to be used in anomaly detection (see Fig. 5). A similar linearization split is carried out for *Training-gi* dataset, and another set of $[LB, UB]$ state bounds are determined. This completes the training steps. For testing, the test sensor value $testval$ (with corresponding sb state appended to it) is read from the testing dataset.

3) *Explainability*: The switchboard provides bijective mapping from training dataset to testing dataset, using a unique numerical string to represent nn-actuation state. The $[LB, UB]$ bounds are determined from the linearized groups. The bijective nature implies the existence of an inverse mapping, from the testing set to training set (see Fig. 5). I.e., provided with the nn-actuator(s) and the test sensor value, we can determine its corresponding training $[LB, UB]$ bounds for the giant and baby-step. When the sensor value is out of bound, the explainability includes sensor label/name, actuation state, the time index t , and the bound it breached (LB or UB). For the explainability pseudocode refer to lines 6-9 in Algorithm 2.

D. An example of operations

In TABLE II, the sensor considered is LIT101. Its nn-actuators are MV101 and P101 (see Fig. 4). The switchboard state at t is an ordered concatenation of its nn-actuator state. LIT101Diff is the difference between consecutive LIT101 readings. Assume first four columns in TABLE II were the only entries in the normal dataset. The remaining two columns are derived using information from the earlier columns. As part of the training, the linearized groups LSG_{sb} are computed for each of existing switchboard states on the training dataset. Let btr denote the baby-step training. The btr step computes the min and max boundaries, $[LB, UB]_{sb=11}^{btr} = [0.0011, 0.1570]$

and $[LB, UB]_{sb=01}^{btr} = [0.0785, 0.4711]$ (see TABLE II). These trained $[LB, UB]$ bounds are used to determine the anomaly status.

Index	LIT101	MV101	P101	Switchboard	LIT101Diff
1	121.2518	1	1	11	—
2	121.4088	1	1	11	0.1570
3	121.4099	1	1	11	0.0011
4	121.6050	0	1	01	0.1951
5	121.6835	0	1	01	0.0785
6	122.1546	0	1	01	0.4711

TABLE II
AN EXAMPLE NORMAL (TRAINING) DATASET FOR SENSOR LIT101 AND ITS NN-ACTUATORS. SWITCHBOARD AND DIFF VALUES ARE DERIVED.

Index	LIT101	MV101	P101	Switchboard	LIT101Diff
1	123.2151	1	1	11	—
2	121.6835	1	1	11	-1.5316

TABLE III
AN EXAMPLE ATTACK (TESTING) DATASET SHOWING ROW ENTRIES FOR SENSOR LIT101 AND ITS NN-ACTUATORS.

An example attack (testing) dataset (see first four columns) is shown in TABLE III. As seen earlier, the remaining two columns are derived. The testing switchboard state $sb = 11$ is mapped to matching $[LB, UB]_{sb}^{btr}$, seen in training. The corresponding testing sensor value is checked to determine if it satisfies the condition $LB \leq testsensorval \leq UB$. For time index 1 in TABLE III, LIT101Diff is undefined. This is because computing a diff, i.e., current value - previous value, requires two consecutive (time indexed) LIT101 sensor values. For testing index 2, a diff exists. With respect to corresponding (training) $sb = 11$ actuation state, its baby-step bounds are $[0.0011, 0.1570]$. The verification check therefore is $0.0011 \leq LIT101Diff \leq 0.1570$? Which in our example ($LIT101Diff = -1.5316$) returns the *False* condition, resulting in the baby-step method issuing an out of bounds warning event. The 6 baby-step bounds for LIT101 determined from training using the SWaT normal dataset are shown in Fig. 6. Had the giant-step been used, we would train and test directly on the sensor values in column labeled LIT101 instead of column labeled LIT101Diff, in TABLE II and TABLE III, respectively.

For other sensors such as FIT201 (see Fig. 4), it is of interest to find the giant and baby-step training bounds. When water is pumped using P101, the flow rate measured by FIT201 settles in a steady range of (mostly) adjacent values. Training directly on the flow rate, and finding bounds, allows us to detect anomalous flows. The baby-step is also of significance w.r.t. this sensor. We are also interested in the bounds on rate of flow change, w.r.t. its nn-neighbor(s) actuation. Later, in section IV we will see that the baby and giant-step may be combined with an extended detection algorithm, to provide up to four different detection mechanisms.

E. Giant-Step Baby-Step Anomaly Detector

The **giant-step baby-step** (GiBy) anomaly detector is based on the switchboard implementation, linearization step, and

determination of training bounds detailed in Section III-C. The algorithm is run for each sensor with a nn-actuator(s) dependency. We note that the GiBy algorithm is able to take as input, nn-actuator(s) with different actuation sequences. For example, the first actuator a_1 might have $|k_1| = 2$, a_2 with $|k_2| = 3$ and a_3 actuator with $|k_3| = 2$. There is also no limitation on the number of input nn-actuators imposed by the algorithm. Algorithm 1 shows GiBy-core training for one sensor. The giant-step training in lines 30-33 involves determining the switchboard states in lines 3-7, linearized state groups (LSG) in lines 15-23 and the bounds in lines 24-29. The baby-step training in lines 34-38 additionally involves computing the *diff* in lines 8-14, before it proceeds with the linearization step and bound determination. Algorithm 2 shows GiBy-core testing for one sensor. The giant-step test in lines 10-13 are used to determine the switchboard state and map the test sensor actuation state to its training [LB, UB] bounds (see lines 1-9). The baby-step in lines 14-19 also carries out similar steps, except it is carried out for the test sensor diff value. Note the definition BoundsCheck in lines 1-9 also provides an explanation when an anomaly is detected. The underlined variables are replaced with its actual value when printed.

IV. EXTENDING DETECTION CAPABILITIES

Extended capabilities are built to provide rare event anomaly detection for a time-series window of sensor readings tending towards the boundaries or around the centre of the probability distribution (see Fig. 8). Stealth attacks on ICS were studied in Urbina et al [24]. The goal of an adversary is to keep detection statistic below the selected threshold. Some detection solutions using a residual-error threshold [13] are vulnerable to stealth attacks. We used stateful detection to catch anomalies based on the probability distribution of state values in a time-window. The extended detector is able to flag a subset of time-series sensor readings clustered around (within) the boundaries or towards the centre of the distribution (see Fig. 8); that are rare events not captured during training. Note that it may be used with giant or baby-step. While it improves stealth attack detection (see Section V), this extended detector may not flag attack record readouts exceptionally close to normal operations; typically, when normal operations range, and its statistics are known to an (insider) attacker.

The anomaly probability is deduced using a proposed empirical method. The empirical method is preferred when a relatively large number of entries are available in the training dataset (this is case for the SWaT used within Section II-C). The examples in TABLE IV and TABLE V are hypothetical, and drives the core idea. The extended detector relies on a probability score deduced from the sensor reading.

1) *Empirical anomaly score*: Consider the hypothetical example in Fig. 7. The range of sensor readings is shown along the x-axis, and the frequency range of these readings is shown up the y-axis. The distribution of sensor values for a given sensor and a specified actuation state is shown. As an example, we begin to calculate the anomaly probability for the sensor reading value of 2 (shown by the red line). We compute four probabilities w.r.t. the sensor — Equation 2 and

Algorithm 1: Giant-step Baby-step Training

Input: Set1: Normal SWaT dataset (NSD). Set2: Sensor label s and time index i . Set3: Sensors-actuators relationship graph RG .
Output: The state bounds [LB, UB] for the sensor.

```

1 Def NearestNeighbors( $s, RG$ ):
2    $\lfloor$  return nearestActsList:= FindInOutSensorEdges( $s$ )
3 Def SwitchBoardState( $i, s, actstate, RG$ ):
4   nnActs:= NearestNeighbors( $s, RG$ )
5   sbFlag:= isMember( $actstate, nnActs$ )
6   if sbFlag==True: return sb:= Concat(nnActs)
7   else: return -1
8 Def DatasetDiff( $NSD, s$ ):
9   for index  $i = 1$  to end //  $i \in NSD$ 
10  do
11    sensorval:= FindSensorVal( $NSD, i, s$ )
12    diffSenVal:= sensorval[i]-sensorval[i-1]
13    NSD[i]:= AppendToRow( $NSD[i], diffSenVal$ )
14  return NSD
15 Def LinearizeStates( $s, NSD, RG$ ):
16  for index  $i = 0$  to end //  $i \in NSD$ 
17  do
18    nnActs:= NearestNeighbors( $s, RG$ )
19    senval:= FindSensorVal( $NSD, i, s$ )
20    actstate:= FindAllActState( $i, s, NSD$ )
21    sb:=SwitchBoardState( $i, s, actstate, RG$ )
22    LSG[sb]:= AppendRow( $i, s, senval, nnActs, sb$ )
23  return LSG// Linearized State Group
24 Def DetermineBounds( $LSG$ ):
25  for Each linearized state group LSG do
26    LB:= min( $senval_i$ )  $\forall i \in LSG$ 
27    UB:= max( $senval_i$ )  $\forall i \in LSG$ 
28    LSGBoundList:= AppendRow( $s, sb, LB, UB$ )
29  return LSGBoundList
30 Def GiantStepTrain( $s, NSD, RG$ ):
31  LSG:=LinearizeStates( $s, NSD, RG$ )
32  GiBoundList:= DetermineBounds( $LSG$ )
33  return GiBoundList
34 Def BabyStepTrain( $s, NSD, RG$ ):
35  NSDdiff:= DatasetDiff( $NSD, s$ )
36  LSG:=LinearizeStates( $s, NSDdiff, RG$ )
37  ByBoundList:= DetermineBounds( $LSG$ )
38  return ByBoundList

```

3 are the left and right anomaly probabilities, respectively. The empirical anomaly probability is shown in Equation 5, and finally it's not-anomaly probability is $1 - Pr_{anom}$. Let sen be the sensor reading of interest (for this example $sen = 2$). It is the third reading from the left. Therefore, sensor index $n = 3$. Let sen_i^f be the frequency at the i^{th} sensor index. Let l denote left and r for right. We find

$$T = \sum_{i=1}^{end} sen_i^f \quad (1)$$

$$P(sen_l^n) = \left(\sum_{i=1}^{n-1} sen_i^f \right) / T \quad (2)$$

$$P(sen_r^n) = \left(\sum_{i=n+1}^{end} sen_i^f \right) / T \quad (3)$$

For the example with $n = 3$, we have $T = 36$, $P(sen_l^n) = 3/36 = 0.083$ and $P(sen_r^n) = 31/36 = 0.861$. The center of the probability distribution is $p = 0.5$. The anomaly probability

Algorithm 2: Giant-step Baby-step Testing

Input: Set1: Attack SWaT dataset (ASD). Set2: Sensor label s and time index i . Set3: Sensors-actuators relationship graph RG and training bounds lists.

Output: The state bounds [LB, UB] for the sensor.

```

1 Def BoundsCheck ( $i, s, senval, actstate, RG, BoundList$ ):
2    $sb :=$  SwitchBoardState( $i, s, actstate, RG$ )
3    $Bounds :=$  retrieveTrainingBounds( $s, sb, BoundList$ )
4    $LB :=$  Bounds['Low']
5    $UB :=$  Bounds['High']
6   if  $LB \leq senval \leq UB == False$  or  $sb == -1$  then
7     print("Explanation: Anomaly DETECTED for sensor  $s$  at
8       time index  $i$  for actuation state  $actstate$  because sensor
9       value  $senval$  not in  $[LB, UB]$ . Otherwise, actuation
10      state is invalid or not seen in training when  $sb = -1$ .")
11   else
12     print("No anomaly was detected.")
13
14 Def GiantStepTest ( $i, s, ASD, RG, GiBoundList$ ):
15    $actstate :=$  FindAllActState( $i, s, ASD$ )
16    $senval :=$  FindSensorVal( $ASD, i, s$ )
17   BoundsCheck( $i, s, senval, actstate, RG, GiBoundList$ )
18
19 Def BabyStepTest ( $i, s, ASD, RG, ByBoundList$ ):
20    $actstate :=$  FindAllActState( $i, s, ASD$ )
21    $senvalNow :=$  FindSensorVal( $ASD, i, s$ )
22    $senvalPrev :=$  FindSensorVal( $ASD, i - 1, s$ )
23    $senval :=$   $senvalNow - senvalPrev$ 
24   BoundsCheck( $i, s, senval, actstate, RG, ByBoundList$ )

```

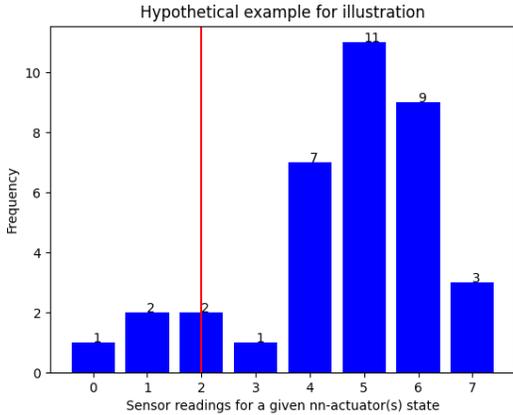


Fig. 7. An example sensor readings distribution for a given nn-actuation state. Some readings may repeat and is represented by its y-axis frequency.

for sen is computed using Equation 5. The edge cases are shown in Algorithm 3.

2) *Sliding window in training and testing:* The goal is to detect a series of consecutive sensor reading probabilities whose values cluster around the probability boundaries $anomprob = 1$, and $anomprob = 0$ (see example Fig. 8). For training, we multiply a series of consecutive not-anomaly probabilities as part of a sliding window, for sensor readings in a time-series dataset. The min and max values $[min, max]$ of the product are computed. For testing, the same sliding window multiplication is carried out on the test dataset. For the multiplied test probability $testprob$, when the condition $min \leq testprob \leq max$ is violated, an anomaly is raised and the explanation provided. The set of consecutive test sensor

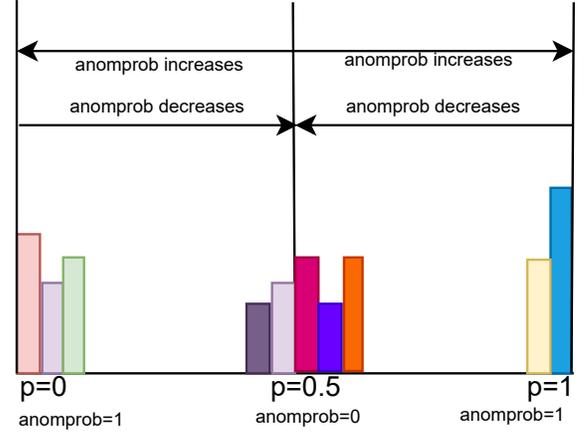


Fig. 8. A hypothetical illustration of the distribution probabilities for sensor readings given its nn-actuation state, centred at $p = 0.5$. The anomaly probability is 0 at the centre and tends towards 1, as it approaches either of the boundaries. The x-axis is range of probability distribution and y-axis, its frequency.

values for which anomalies are raised are seen as rare events, not captured during training of normal dataset. It may also be used to detect multiple rare events (if any), over a given time interval. TABLE IV shows a hypothetical training dataset for sensor LIT101 and switchboard status 11. Column labeled 1 - PrAnom, is the empirically determined not-anomaly probability. PrAnom is derived using Equation 5. Column labeled SWProduct is the sliding window (sw) product. The sliding window length (sw_{len}) was chosen to be 3, for ease of illustration. We compute

Index	LIT101	LIT101Diff	sb	1-PrAnom	SWProduct
1	186.2518	-	11	-	-
2	186.4088	0.1570	11	0.8	-
3	186.4099	0.0011	11	0.2	-
4	186.6050	0.1951	11	0.6	0.096
5	186.6835	0.0785	11	0.6	0.072
6	187.1546	0.4711	11	0.2	0.072

TABLE IV
AN EXAMPLE TRAINING DATASET. NN-ACTUATORS ARE NOT SHOWN. SB IS SWITCHBOARD, $SW_{len} = 3$. COLUMNS LABELED 1-PRANOM AND SWPRODUCT ARE DERIVED.

$$SWProduct(Index) = \prod_{i=Index-sw_{len}+1}^{Index} 1 - PrAnom_i \quad (4)$$

For the example in TABLE IV, the min and max bounds for SWProduct are [0.072, 0.096] for $sb = 11$.

Index	LIT101	LIT101Diff	sb	1-PrAnom	SWProduct
1	111.2324	-	11	-	-
2	111.2824	0.05	11	0.4	-
3	111.3324	0.05	11	0.4	-
4	111.3824	0.05	11	0.4	0.064

TABLE V
TEST FOR EXTENDED ANOMALY. TESTVAL IS LIT101DIFF, $SW_{Len} = 3$.

The example test dataset in TABLE V shows the test sensor readings and its diff. The empirical not-anomaly probability

(1-PrAnom) for LIT101Diff is 0.4, since the diffs are the same. The fourth test index with LIT101Diff has sliding window product, $SWProduct = 0.064$. The verification check therefore is $0.072 \leq SWProduct \leq 0.096$? Which in our example ($SWProduct = 0.064$) returns the *False* condition, and anomaly is flagged by the detector. For the experiments in Section V, the SWaT normal dataset is used for training through use of its sensor values across its linearized group states (see Fig. 5) — for sliding window sizes 5, 10, 25, 50 and 100. For example, a sliding window of size 100 implies that 100 consecutive sensor readings for the given state at one second resolution were trained, and the $SWProduct$ min, max values were the output of that training. This was tested against the SWaT attack dataset. Different window sizes are used to capture anomalies across different time ranges and provide different detection sensitivity depending on for how long the attack lasted. The extended detection training methods are provided in Algorithm 3. The FindMinMaxProduct function at line 37 is used to find the min and max sliding window product bounds for different window lengths and each switchboard state in LSG (derived from training dataset).

$$Pr_{anom}(sen) = \begin{cases} (0.5 - \min(P(sen_l^n), P(sen_r^n))) \cdot 2, \\ \text{when } P(sen_l^n) \neq 0.5, P(sen_r^n) \neq 0.5. \end{cases} \quad (5)$$

For testing the extended algorithm, a similar sequence of steps is used but with the testing dataset. The sensor values used here are the attack dataset readings for giant-step and its diff for baby-step. To find the test sensor value anomaly probability, a binary search is carried out on the corresponding *sortedSenList* (this is the *LSG* values sorted in ascending order on frequency, see lines 1-3 in Algorithm 3) derived from extended training. If *testval* equals the sorted sensor list value, then *PrLeft*, *PrRight* and *PrAnom* (lines 6-27) are computed for this *sortedSenListVal*. For nearest $sortedSenListVal1 < testval < sortedSenListVal2$; on *sortedSenListVal1*, compute its $PrLeft = 1 - PrRight$ and on *sortedSenListVal2*, compute its $PrRight = 1 - PrLeft$. I.e., we include the frequency of the tested value in the probability computation. Now that we have the *PrLeft* and *PrRight* probabilities for *testval*, $1 - PrAnom$ is computed as seen in line 28. The products are then computed for the sliding window lengths and checked against the min, max extended bounds determined in training. An anomaly flag is raised when it is out of bounds. Note that sorting and static probabilities were computed once, stored during training and reused later.

V. EXPERIMENTAL RESULTS AND ANALYSIS

The experiments are carried out on a Windows 11 machine, with an Intel(R) Core(TM) i7-1195G7 @ 2.90GHz processor, with one of four cores used in experimentation. The GiBy-core detector (Section III-E) and extended algorithm (Section IV) are used. The SWaT normal dataset is used for training and attack dataset is for testing. Table VIII shows the attacks that are detected by GiBy. Attacks 5, 9, 12, 15 and 18 in TABLE VIII had no physical impact on the system. Hence, it is not displayed. For Attack 3, the 1 mm rate of change

Algorithm 3: Extended Detection Training Functions

```

1 Function SortAscendingByValueOnFrequency (LSG, sb):
2   sortedSenList:= SortAscendingByValueOnFrequency(LSG,
3     sb)// Sort the Linearized group for
4     switchboard state on frequency of
5     repeated sensor values as in Fig. 7.
6   return sortedSenList
7
8 Function Total (sortedSenList):
9   T:=  $\sum_{i=1}^{end} sortedSenList_i^f$  return T
10
11 Function FindPrLeft (index, sortedSenList, T):
12   n:= index
13   PrLeft:=  $(\sum_{i=1}^{n-1} sortedSenList_i^f) / T$ 
14   return PrLeft
15
16 Function FindPrRight (index, sortedSenList, T):
17   n:= index
18   end:= lastIndex(sortedSenList)
19   PrRight:=  $(\sum_{i=n+1}^{end} sortedSenList_i^f) / T$ 
20   return PrRight
21
22 Function FindAnomPr (PrLeft, PrRight):
23   if PrLeft == 0 or PrRight == 0 then
24     PrAnom:= 1 // flag anomaly
25   else if PrLeft ≠ 0.5 and PrRight ≠ 0.5 then
26     PrAnom:= |0.5-min(PrLeft, PrRight)| · 2
27   else
28     PrAnom:= 0.5
29   if PrLeft + PrRight < 0.5 then
30     PrAnom:= |PrLeft+PrRight|
31   if PrRight == 0 and PrLeft ≠ 1 then
32     PrAnom:= PrLeft
33   else if PrLeft == 0 and PrRight ≠ 1 then
34     PrAnom:= PrRight
35   return index, 1-PrAnom // Not anomaly probability
36
37 Function FindSWProduct (index, swlen, LSG, sb):
38   sortedSenList:= SortAscendingByValueOnFrequency(LSG, sb)
39   T:= Total(sortedSenList)
40   PrLeft:= FindPrLeft(index, sortedSenList, T)
41   PrRight:= FindPrRight(index, sortedSenList, T)
42   PrAnom:= FindAnomPr(PrLeft, PrRight)
43   swproduct:=  $\prod_{i=Index-swlen+1}^{Index} (1 - PrAnom_i)$ 
44   return swproduct
45
46 Function FindMinMaxProduct (index, swlen, LSG, sb):
47   // swlen 5, 10, 25, 50, 100 are used
48   end:= lastIndex(LSG, sb)
49   minswprod:= inf // A large positive number
50   maxswprod:= -inf // A large negative number
51   for i=1 to end do
52     swprod:= FindSWProduct(index, swlen, LSG, sb)
53     minswprod:= min(minswprod, swprod)
54     maxswprod:= max(maxswprod, swprod)
55   return minswprod, maxswprod // Do for each sb state

```

is within the range of normal behaviour of the tested state. Hence, this rate of change is not flagged for one of two attack states. The extended algorithm can detect the anomaly when the sliding window (multiplied) probability exceeds the seen threshold in training. Attack 13, 14, and 17 are not detected because the changing MV-304 and MV-303 states yet showed normal DPIT-301 readings in attack dataset. Attack 24 on P-203 is not detected because despite turning on this pump, the pH sensor readings were in normal range seen in training. The same is also true with Attack 29 on P-203.

	nn-act(s), sen	diff (if any) & Linearize	Core bounds	Extended bounds	Total time
Giant-step	219s	82s	16s	407s	724s
Baby-step	162s	65s	15s	269s	511s

TABLE VI
AVERAGE TIME TAKEN TO TRAIN ONE SENSOR IN NORMAL DATASET.

	nn-act(s), sen	diff (if any)	Total pre-test time	Total test time/#records	Test time/sensor
Giant-step	148s	—	148s	323s/449919	0.00071s
Baby-step	145s	47s	192s	320s/449919	0.00071s

TABLE VII
AVERAGE TIME TAKEN TO TEST ONE SENSOR IN ATTACK DATASET.

The time taken to train and test the normal and attack SWaT datasets are shown in TABLE VI and TABLE VII, respectively. The total training time per sensor was seen to be 9 to 12 minutes. For testing, the total time to copy the nn-actuator(s) and sensor to a new spreadsheet, and computing diff (if any) in the attack dataset are determined in pre-testing steps. The total pre-test time shown is for a total of 449919 attack dataset records. The test step (includes detecting anomaly and providing explanation) is typically seen to be under 1 millisecond per sensor. The time taken for training and pre-testing are linear to the number of records seen in the normal and attack datasets, respectively.

For ECOD and Deep-SVDD in Mathuros et al. [12], the training time is shorter in comparison with GiBy. This is because they train approximately 24 minutes of data, i.e., $24 \cdot 60 = 1440$ records for each of the six process stages of SWaT; as opposed to the entire normal dataset of 495000 records trained by GiBy. ECOD detection time was reported to range from 0.72s to 4.62s, whereas for Deep-SVDD, it was 15.88s to 16.45s. Their Integrated Gradient (IG) XAI used with Deep-SVDD was the fastest with a time of 5.33s. However, the XAI provided there use feature ranking to explain what feature likely caused the anomaly. With GiBy, the explanation is precise and pinpoints to the sensor and nn-actuation state (see line 6 in Algorithm 2).

were also seen in the training/normal dataset. However, they are considered as attacks because the sensor readings changed even though it remained below the detection threshold. For example, in attack 29, P-203 is ON and HCl flowed out of the chemical tank and was mixed with water. However, the pH sensor AIT-203 readings remained at levels seen in normal dataset training, though it resulted in wastage of chemical (HCl in this case). It is possible to detect other edge case attack types such as attack 3 and 6; where the sensor reading is either a fixed value or changes by a constant value but remained below the detection threshold. Such a detection may be carried out by observing the number of times the sensor reading pattern repeats and then assigning probabilities for that event. Also, for any detector training there is an entropy loss; where the training dataset input is compressed into a smaller trained output (for example see Fig. 5). For these reasons, we avoid trying to detect specific attacks and designed GiBy to be a generalized one-class classifier. However, the implementer is free to plug-in their edge-case detection models onto GiBy classifier. Without any plug-in, GiBy in TABLE VIII has 33 Y's and 11 N's in the Detected column. This translates to a 75% detection of the attacks in the table. Since attack 4 and attack 29 on P-201 has non-existent training data, a correct representation of detected attacks is 33/42, which is 78.5%. We emphasize that the detection accuracy of 78.5% is specific to attacks in TABLE VIII. Accuracy might be alternatively reported as *i.*) number of detected attack records in SWaT attack dataset divided by the number of total attack records, or worse as *ii.*) number of total correctly detected labels (Attack/Normal) in attack dataset divided by total number of records in attack dataset. Hence, the interpretation of what accuracy measures and how it is reported may be inconsistent. Next, we use an example attack on SWaT to show the fallibility of anomaly detectors against subtle sensor reading manipulation.

a) Undetectable attacks: Despite best efforts, a powerful

attacker who is able to compromise the PLC may introduce a small constant drift δ to the sensor reading to avoid detection. Consider the following attack scenario where an attacker has successfully compromised the PLC. Let an attacker manipulate the level sensor LIT-101 reading with the goal to underflow the tank and dry pump using P-101 (see Fig. 4). Further, let the lower limit for LIT-101 be 10 mm; below which the PLC turns OFF the pump P-101; to prevent dry pumping. For the tank filling state MV101 = OPEN and P101 = OFF, the attacker would spoof the LIT101 value on the PLC to $Levelval_i = Levelval_i + (\delta \cdot i)$; at each time step $i = \{1, 2, \dots, tmax\} \in tsteps$. This causes the water level to appear to be higher than it actually is. For the tank draining state MV101 = CLOSE and P101 = ON, the attacker would spoof the LIT101 value on the PLC to $Levelval_j = Levelval_j + \delta \cdot (i + j)$; again, causing the tank to appear fuller than it is. By adjusting the value of δ and the number of time-steps i , it results in tank under-flow and P-101 dry pumping. For example, consider LIT-101 value is 12 mm. Let $\delta = 0.01$ and $tsteps = 200$. When tank is in filling state for 150 time-steps, $Levelval_{150} = Levelval_{150} + (0.01 \cdot 150)$. Our accumulated drift is now 1.5 mm. Next, let the tank be in draining state for the next 50-time steps. Then,

Attack no.	Attack point	Attack event	Detected	nn-actuator/sensor	Remark
1	MV101	Close MV101	Y	MV101, LIT101	Detected with GiBy-step
2	P102	Turn on P102	Y	P101, P102, P103	Detected with GiBy-step
3	LIT101	Increase by 1 mm each second	N	MV101, P101, LIT101	Detected because state not seen in training
4	MV104	Open MV104	N	MV101, P101, LIT101	One of two attack states detected with Baby-step extended algorithm
5	LIT101	Increase by 1 mm each second	Y	MV101, P101, LIT101	MV104 reading not in SVDD dataset
6	AIT203	Set value of AIT203 as 0	Y	P101, MV101, AIT203	Detected with GiBy-step
7	LIT101	Water level increased above 1000 mm set point	Y	MV101, LIT101	Detected with GiBy-step
8	DPTF301	Set value of DPTF301 as -0.48kg	Y	P101, DPTF301	Detected with GiBy-step
10	P101	Set value of P101 as ON	N	P101, EV401, P101	Detected with Baby-step
11	P101	Set value of P101 as 0	Y	P101, EV401, P101	Detected with Baby-step
12	MV301	Close MV301	N	MV301, MV304, DPTF301	DPTF301 reads normal. No other sensor relates directly with state
14	MV301	Do not use MV301 opens	N	MV301, MV304, DPTF301	DPTF301 reads normal. No other sensor relates directly with state
16	LIT101	Increase water level	Y	MV101, LIT101	Detected with Baby-step extended algorithm
17	MV101	Close MV101	N	MV101, MV104, DPTF301	DPTF301 reads normal. No other sensor relates directly with state
19	AIT203	Set value of AIT203 as 16.25kPa	Y	P101, P102, AIT203	Detected with GiBy-step
20	AIT203	Set value of AIT203 as 25.75kPa	Y	P101, P102, AIT203	Detected with GiBy-step
21	MV101, LIT101	Set value of LIT101 as continuously. Value of LIT101 as 700 mm	Y	MV101, P101, LIT101	Detected with Baby-step
22	UV401, P101	Stop UV401. Value of AIT502 set as 150. Force P101 to remain ON	Y	UV401, P101, AIT502	Detected with Baby-step
23	P101, LIT101, MV101	Value of DPTF301 set to 20.4 bar. Increase MV101 opens. Keep MV101 opens.	Y	MV101, P101, DPTF301	Detected with Baby-step
24	P101, P102	Turn off P101 and P102	N, Y	P101, P102, AIT203	Attack not detected as pH sensor values in normal range
25	LIT101, P101	Set value of LIT101 as 1000. P101 in large on	Y	P101, LIT101	Attack detected as OMP sensor values were out of trained bounds
26	P101, LIT101	P101 in water tank continuously. Set value of LIT101 as 801 mm	Y, N	P101, P102, LIT101	Attack detected as both pumps not simultaneously ON in training
27	P101, LIT101	Value of LIT101 set as 0 mm and P101 set as 0 mm	Y	P101, P102, LIT101	Not detected as LIT101 sensor value changes were in trained bounds
28	P101	Close P101	N	—	Turning off P101 will stop water inflow to LIT101 in the expected outcome
29	P101, P102, P103	Turn on P101, Turn on P103, Turn on P102	N, N, Y	P101, P102, AIT203	P101 is never turned on in training. No training data available
30	LIT101, P101, MV101	Turn P101 on continuously. Set value of LIT101 as 700 mm. P101 started itself because LIT101value	Y, Y	MV101, P101, LIT101	Both P101 and P102 not simultaneously ON in training data. This state is detected because line 6
31	LIT101	Set LIT101 to 500 mm	Y	P101, P102, LIT101	Detected with Baby-step
32	LIT101	Set LIT101 to above 1000	Y	MV101, LIT101	Detected with Baby-step
33	LIT101	Set LIT101 to above 1000	Y	MV101, P101, LIT101	Detected with Baby-step
34	P101	Turn P101 off	N	P101, P102, LIT101	P101 is OFF and P102 is ON in training dataset
35	P101, P102	Turn P101 off	N	—	P101 is OFF and P102 is OFF, and water stop flowing is expected behavior
36	LIT101	Set LIT101	Y	MV101, P101, LIT101	Detected with Baby-step
37	P101, P102	Close P101. Set value of P102	Y	P101, P102	Detected with Baby-step
38	AIT203	Set value of AIT203 as 260	Y, Y	UV401, AIT203	Attack on AIT203 detected with Baby-step
39	P101	Set value of P101 as 0.5	Y, Y	P101, P102, AIT203	Attack on P101 detected with Baby-step
40	AIT203	Set value of AIT203 as 260	Y, Y	P101, P102, AIT203	Attack on AIT203 detected with Baby-step
41	P101	Set value of P101 as 0	Y	UV401, P101	Detected with Baby-step
42	LIT101	Decrease value by 1 mm each second	Y	P101, P102, LIT101	Detected by Baby-step extended algorithm

TABLE VIII
THE ATTACKS DETECTED USING GIBY.

1) *Analysis:* The attacks in TABLE VIII — 3, 13, 14, 17, 24, 26, 29 were either undetected or partially detected. This is because the range of sensor values in the attack dataset

$Levelval_{50} = Levelval_{50} + 0.01 \cdot (150 + 50)$. This causes additional drift of 0.5, and a total drift of $1.5 + 0.5 = 2$ mm. As the tank continues to drain, and LIT-101 sensor reading arrives at 11.9 mm, in reality the tank level is at 9.9 mm. The PLC control logic considers the tank level to be above the set lower-bound, but tank underflow occurred, and P-101 is dry pumping. To make detection further intractable, the drift δ may be further reduced and the time-steps increased. Further, this attack may be used on one or more sensors simultaneously and might be used to create synchronized failures when additional system information is available. Such subtle attacks tailored to remain below the detection threshold are difficult to detect (refer to Section VIII for possible remedial steps).

VI. DISCUSSION

A. Detector limitations

Our core detection solution (see Section III) is limited to its nn-actuators, for any given sensor. We mapped a smaller state-space of actuator states to the sensor reading. As a result, we forsook the determination of a larger number of tighter individual bounds that takes exponential time for the convenience of a practically smaller solution seek space. However, our sensor measurements are taken directly from the normal operations dataset. As a result, the overall determined [LB, UB] bounds remained unchanged, and hence does not change the overall safety limits determined. For example, consider eight dependent actuators influencing a sensor reading. When the actuation sequence for each actuator is $k = \{\text{on}, \text{off}\}$, then $|k| = 2$. When $|k| = 2$ and number of actuators is 8, there are a total of $2^8 = 256$ actuation states. This requires determination of $[LB_i, UB_i]_{\forall i=1, \dots, 256}$ bounds. Additionally, let us assume that only 3 out of the 8 dependent actuators are nearest neighbours, w.r.t. sensor s . Using GiBy, we would have only considered nearest neighbour ($2^3 = 8$) actuation states, and determined $[LB_j, UB_j]_{\forall j=1, \dots, 8}$. As a result, there was no change in the overall state sensor bounds (safety limits) determined, although some detection resolution was lost by considering only the nearest neighbours.

Another limitation expected from the anomaly detector is a sensitivity to out of bound measurements, resulting in warnings being issued when no attack is present. This limitation could occur even in the system normal (initial) operating phase. The normal dataset is expected to cover its everyday boundaries of operation. However, in some cases the normal dataset may not cover all boundaries of normal operations, and these normal boundaries may be surpassed during future plant operations. Any logged warning observed, that is deemed as normal, will have to be reviewed and trained, to avoid those warnings from re-appearing. However, this will stabilize once sufficient normal (training) data indicative of its range of operations is available. This is a training dataset limitation that affects our model. A final limitation is that it may not detect all invalid actuation states. It is limited to the resolution of the actuator states seen for the nn-actuators.

B. Detection in an early prototype system design

In Fig. 3, we note that hydrochloric acid (HCl) is added into the first tank. Water is pumped using P101 and combined

with the HCl pumped by P201. The pH is then measured using sensor AIT201. However, if the HCl being transported is mislabelled as NaOCl, each is poured into an incorrect tank. This means the tank holding HCl now holds NaOCl and vice versa. AIT201 no longer measure the pH of water mixed with HCl, and it instead measures the mix with NaOCl. Now, if the dosing rates are different, the chemicals will be administered at an incorrect rate. For this scenario, it is not possible to detect the pH of water after the addition of HCl, without the assistance of a redundant pH sensor downstream. This design shortcoming is rectified in the actual SWaT testbed. The normal dataset provided is indicative that NaCl was used in the first tank, HCl in second and NaOCl in the third tank. This is because the sensor measurements from the dataset, corresponded with conductivity (AIT201), pH (AIT202) and oxidation reduction potential (AIT203). An inadequate design may lead to detection failures, for no fault of the detector modelled. We emphasize that good detection capabilities are reliant on meticulous system design.

C. GiBy Detection capabilities

In the threat model (Section II-D), the attacks are classified into single-stage single/multiple component attacks and multi-stage single/multiple component attacks. The component considered is a sensor or an actuator. The goal of this classification is to build a quantifiable measure (score) on the quality of the detector, w.r.t. its ability to detect these types of attacks. Without loss of generality, the classification may be merged into single and multiple sensor/actuator attack(s) on a system.

Our core detector (see Section III) can detect single sensor or actuator manipulations, when it breaks the safety bounds of the relation between sensor and its actuation state. When both sensor and corresponding actuator(s) are spoofed, it is no longer able to individually detect that attack, but it depends on whether any of the sensor readings downstream², as a result, went out of bounds. For example, in Fig. 3, consider an attack where P101 is spoofed to ON and FIT201 reading to *normalFlowRate*, despite the pump being OFF and no water flowing. Assume the controller logic is to turn on the first HCl dosing pump (P201), when the FIT201 flow rate y is between $0 < x < y < z$ and to OFF otherwise. The attack causes the dosing pump to release HCl even when no water is pumped³. However, due to pH sensor AIT201, this attack is detected. Further, if we assume that AIT201 is also spoofed to show normal range of measurements — a redundant downstream sensor AIT501 incorporated in $P5$ is able to detect the abnormality in pH. However, if AIT501 is also spoofed, it may not be able to detect the attack from the measurements made. As a side effect, changing the pH might possibly change its conductivity sensor reading, permitting the attack to be detected. It would be difficult to conclude without further experimentation.

In some scenarios, it may be sufficient to detect the anomaly at a later stage, and in others, it requires to be detected in

²Downstream bounds are also captured by GiBy, since it is computed for every sensor with nn-neighbour actuator(s).

³The success of this attack depends on whether the attacker can compromise the controller PLC and manipulate its measurements.

the same stage. For example, if pump P101 is running above its safe rpm, it might cause the adjacent (following) pipe to burst, affecting the plant and personnel. Where such a concern exists, the design has to include a pressure sensor. When there is a risk the pressure sensor reading is spoofed, a physical pressure safety (release) valve will have to be incorporated into the design. An anomaly detector’s ability to detect single/multiple attacks is highly dependent on intrinsic design relationships and system dependencies. This tight coupling makes generalized interpretations on detection capability (by only considering the detection model) less meaningful.

D. Non-neighbor PLC imposed actuation

The P_2 stage in Fig. 3 involves chemical dosing. Dosing tanks are required to be in the order NaCl, HCl and NaOCl (see Section VI-B for reason). AIT201 in the SWaT dataset measures conductivity. According to the PLC control logic sheet, the conductivity sensor AIT503 (downstream sensor in P_5) read by the PLC instructs P201 to be turned off, when this sensor value hits a preset threshold. Turning off the salt pump, in-turn regulates the conductivity sensor reading at AIT503. Hence, P201 is additionally considered as a nn-actuator (though not a neighbour) for the sensor AIT503; to capture this relationship. It is added to train the sensor (using GiBy). In general, the PLC logic sheet is examined to include such ‘additional’ nn-actuators, when there is a sensor dependency. Not including this relationship does not change the overall detection bounds captured in training. However, it means that some detection resolution is lost because we are no longer directly capturing the relationship between AIT503 and pump P201.

E. Explainability

Anomaly detection is typically followed by a mitigation response, when required. However, extra processing is required to identify why the detector raised a warning; to inform the plant operator of the detector reasoning. This is to assist with the next steps such as diagnostics, emergency manual shutdown, or plant recovery. For example, the anomaly detector may flag time index $t = 34323^{th}$ second from start of operation as anomalous. This time index has 51 sensor and actuator readings (see TABLE. I). Without further explanation, it is not possible to deduce *i.*) what sensor(s) readings detected the anomalies, *ii.*) what change in actuation states led it to an anomalous state, and *iii.*) what bound it breached.

The more sophisticated AI/ML models using neural network and deep learning are inherently black box models. They are used to detect anomalies but requires additional support from explainable AI (XAI) models to explain the decision made by the AI detector. However, both the ML and XAI models have individual accuracy limitations and is compounded when used in combination. It also takes longer to train and test these models (see Section V). On the contrary, explainability with the proposed GiBy detector is straightforward. Note that we employed a switchboard in Section III-C. Traceability is offered via the one-to-one and onto map, using the switchboard — also in the inverse direction, from sensor reading (under

testing) to its [LB, UB] bounds (determined in the training phase). Our training and testing are carried out per sensor, for all concerned sensors and nn-actuator(s). When an anomaly is detected, this immediately tells us what sensor readings are causing anomalies. Since the sensor is known, we can look at its current and previous actuation states, from the state log. GiBy raises an anomaly warning when *i.*) the sensor reading (under testing) for the nn-actuator state is out of bound. Our detector is able to provide an explanation remark, whether it was because the sensor reading was $> UB$ or $< LB$. *ii.*) when the normal training dataset did not have the nn-state actuator state for the sensor, in it. This may happen for two reasons *a.*) the actuation state is invalid or *b.*) that state is yet to be encountered in the plants’ normal operation and hence not yet seen in the normal training dataset. Another advantage of GiBy, is that it does not produce false positives except when the normal dataset is not wide enough to capture rare non-everyday spikes, anomalies or boundary excursions. Its detection boundaries are clear and limited to what it has seen in normal dataset training. Our detector system is designed to offer a level of explainability that is expected to equip the operator to deal with most contingencies. The capability of our system to provide the correct sensor location and response to faulty components/attacks subsequently reduces the plant downtime.

F. Usability in other sectors

Anomaly detection and mitigation in some sectors (such as power), in certain scenarios, require an actuation within 10-15 milliseconds from event occurrence, to ensure safety. The GiBy algorithm presented in Section III is universally transferable and GiBy-core testing may be parallelized. This is because it produces a series of independent [LB, UB] boundaries. It may be pre-programmed into the control logic of an Intelligent Electronic Device (IED), typically used in a power substation. It may be used to trip the circuit breaker in near real-time (when sensor readings are out of safety bounds), or to inform an operator regarding the anomalous behaviour. The computational cost of GiBy-core tests (see Section III-E) are to see whether the sensor value is in between two bounds [LB, UB]; taking up to two comparisons. The main costs of the extended algorithm tests (see Section IV) are the binary search and sliding window multiplications in each time step. Its time complexity may be reduced to that of binary search on the pre-sorted linearized group (that takes logarithmic time), and a maximum of one multiplication and division: for each time step. Instead of repeatedly multiplying every sensor derived value in the sliding window, the newest time step sensor derived value may be multiplied into the existing product, and the sensor derived value at the tail of the window is divided out. As a result, it might also find (outlier detection) applications in resource constrained devices, edge computing devices and legacy systems with lower computational power.

VII. RELATED WORK

While there are solutions to detect packet anomalies on the OT network [25] and memory anomalies on computation

devices [26], [27], amidst other detection vectors, the focus of this work is on process anomaly detection [28], [29]. Several approaches are available for design-centric and data-centric anomaly detection. In a design-centric approach, for example, Adepur and Mathur [30] captured the design using state condition graphs. It was used to capture the conditions on actuation w.r.t. sensor readings, expressed using Boolean conditions, and represented in a graphical format. The work in Yoong et al. [31] used axiomatic design methodology for systems, where it iteratively decomposes a CPS design to sets of dependent components and transformed into invariants. In Merwa et al. [32], an association rule mining technique was used to generate attack and invariant rules. In a data-centric approach, for example, there exist solutions to distinguish outliers, such as ECOD in Li et al. [10]. It computed a univariate empirical cumulative distribution function for each dimension separately. To measure the chance of a data point being anomalous, they computed its tail probability across all dimensions. Another outlier detector called deep Support Vector Data Description (Deep-SVDD) proposed in Ruff et al. [11], trained a neural network while minimizing the volume of a hypersphere that enclosed the network representations of the data. In WaXAI, Mathuros et al. [12] employed ECOD and Deep-SVDD in the context of SWaT anomaly detection. In addition, they also employed explainable AI [14] models, namely, kernel SHAP [33], LIME [34], ALE [35] and IG [36], to provide explanation for the anomaly detected. Further, a 20% improvement in attribution of attack root cause over SHAP was achieved by using a Factorization Machines based approach by Avdalovic et al. [37].

In this work we provided attribution (explainability) linked directly to the sensor. Explainability was provided for each sensor when out of bounds, to top up the recommendations in Fung et al [38]. The GiBy classifier (see Section III) is in comparison [39], [40] marginally less design centric. However, to improve detection resolution, it may use the know-how of direct sensor-actuator relationships, available from the testbed design document or derived from the system architecture by a subject expert, and PLC imposed non nn-actuation (see Section VI-D) logic available in the PLC control logic sheet. Note, we do not break it down to process stages or look for invariants across sub-systems to find association rules. Instead, boundaries are determined per sensor and relationships are generally constrained to nn-actuator(s). We assume that a suitable approach was used to set the PLC min-max boundaries discussed in Section II-A. These hard boundaries are essential for correct system operations even in the absence of anomalies. Recent developments in using large language models (LLMs) for anomaly detection is extensively reviewed in Su et al [41]. Unlike GiBy, none of the log analysis detectors surveyed there positioned themselves as real-time detectors. A generalized note they made on computational efficiency in LLM deployment for forecasting and anomaly detection was its sheer scale and complexity of those models. It demanded substantial computational resources, likely to limit its scalability and usability for real-time applications. Also of concern was the sustainability and energy efficiency of LLMs with future trends likely to persist on environmentally friendly

models through practices such as optimized algorithms and energy-efficient hardware. In this aspect, unlike LLMs and some of the deep learning AI models discussed; GiBy showed millisecond detection speeds on an everyday used microprocessor demonstrating that the algorithm is energy efficient.

VIII. CONCLUSIONS AND FUTURE WORK

We presented a simple yet powerful anomaly detector that trained four set of bounds — the min-max bounds for the giant-step, the rate of change bounds for baby-step, and the extended detection algorithm bounds, for both giant and baby-steps. The detector acted as a one-class classifier. Any tested value outside trained bounds were flagged as an anomaly. The transparent design of GiBy anomaly detector made it straightforward to implement and easy to interpret. The explainability provided pinpoints the sensor and actuation state for which anomaly was detected, and what bounds it breached. The experiments showed that testing for an anomaly and providing an explanation took around $\frac{1}{1000^{th}}$ of a second per sensor. This made it useful for implementation in systems where near real-time decisions are made, or on devices that are resource constrained.

In our experiments, we observed in TABLE VIII that some attacks were well within the range of normal behavior and could not be detected. We emphasize that detecting all attacks in an attack table (TABLE VIII) is not the same as detecting all attacks on a system. Attack tables, such as those provided by SWaT are useful to check whether the detector detects different types of attacks, based on attacker capabilities. However, detecting some attacks are expected to remain difficult for all anomaly detectors, due to entropy loss in training and indistinguishability from normal operational data. Due to the undetectable attacks discussed in Section V-1a, in future work we will explore other types of detection using digital signature verification and encryption to hide the inferences on sensor readings across OT, segregating and air-gapping high risk sensors and PLCs, paying closer attention to attack prevention and fast recovery strategies that do not compromise plant and personnel safety in the meantime.

ACKNOWLEDGMENT

This research is funded by UK Research and Innovation, Knowledge Transfer Partnership project, KTP reference 13504. The authors are grateful for the discussions with Stephen and George, leading to the manuscript.

REFERENCES

- [1] C. Jaikaran. (Last updated - August, 2023) Cybersecurity: Selected Cyberattacks, 2012-2022.
- [2] TRENDMICRO. (January 12, 2015) German Steel Plant Suffers Significant Damage from Targeted Attack.
- [3] R. M. Lee, M. J. Assante, and T. Conway. (Dec 30, 2014) ICS CP/PE (Cyber-to-Physical or Process Effects) case study paper – German Steel Mill Cyber Attack.
- [4] Dean Parsons. (August 9, 2024) What’s the Scoop on FrostyGoop: The Latest ICS Malware and ICS Controls Considerations.
- [5] Marshall Abrams and Joe Weiss. (2008) Malicious Control System Cyber Security Attack Case Study – Maroochy Water Services, Australia.
- [6] James. (October 13, 2024) 11 recent cyber attacks on the water and wastewater sector.

- [7] Trevor Quinn. (December 7, 2023) Anti-Israeli hackers leave 180 Mayo homes without water in cyberattack.
- [8] Andy Greenberg. (February 8, 2021) A Hacker Tried to Poison a Florida City's Water Supply, Officials Say.
- [9] J. Goh, S. Adepu, K. N. Junejo, and A. P. Mathur, "A dataset to support research in the design of secure water treatment systems," in *Critical Information Infrastructures Security*, Cham, 2017, pp. 88–99.
- [10] Z. Li, Y. Zhao, X. Hu, N. Botta, C. Ionescu, and G. Chen, "Ecod: Unsupervised outlier detection using empirical cumulative distribution functions," *IEEE Transactions on Knowledge and Data Engineering*, pp. 1–1, 2022.
- [11] L. Ruff, N. Görmitz, L. Deecke, S. A. Siddiqui, R. A. Vandermeulen, A. Binder, E. Müller, and M. Kloft, "Deep One-Class Classification," in *ICML*, 2018, pp. 4390–4399.
- [12] K. Mathuros, S. Venugopalan, and S. Adepu, "WaXAI: Explainable Anomaly Detection in Industrial Control Systems and Water Systems," in *Proceedings of the 10th ACM Cyber-Physical System Security Workshop*, ser. CPSS '24. New York, NY, USA: ACM, 2024, p. 3–15.
- [13] C. Feng, V. R. Palleti, A. Mathur, and D. Chana, "A systematic framework to generate invariants for anomaly detection in industrial control systems." Network and Distributed System Security (NDSS) Symposium, 2019.
- [14] S. Ali, T. Abuhmed, S. El-Sappagh, K. Muhammad, J. M. Alonso-Moral, R. Confalonieri, R. Guidotti, J. Del Ser, N. Díaz-Rodríguez, and F. Herrera, "Explainable Artificial Intelligence (XAI): What we know and what is left to attain Trustworthy Artificial Intelligence," *Information Fusion*, vol. 99, p. 101805, 2023.
- [15] N. Seliya, A. A. Zadeh, and T. M. Khoshgoftaar, "A literature review on one-class classification and its potential applications in big data," *J. Big Data*, vol. 8, no. 1, p. 122, 2021.
- [16] J. Cao, H. Dai, B. Lei, C. Yin, H. Zeng, and A. Kummert, "Maximum coreentropy criterion-based hierarchical one-class classification," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 8, pp. 3748–3754, 2021.
- [17] CAPEC. (Accessed on 8 February, 2025) CAPEC VIEW: Domains of Attack (Version 3.9) .
- [18] P. Rieger, M. Chilese, R. Mohamed, M. Miettinen, H. Fereidooni, and A.-R. Sadeghi, "ARGUS: Context-Based Detection of Stealthy IoT Infiltration Attacks," in *32nd USENIX Security Symposium (USENIX Security 23)*, Anaheim, CA, Aug. 2023, pp. 4301–4318.
- [19] R. Bentley and A. Sarkar. (Jun 18, 2024) Humans in AI: The necessity for human-in-the-loop (HILT).
- [20] T. Clemmensen, M. T. Moghaddam, and J. Nørbjerg, "Cyber-physical systems with Human-in-the-Loop: A systematic review of socio-technical perspectives," *Journal of Systems and Software*, vol. 226, p. 112348, 2025.
- [21] L. Yuan, S. Yu, Z. Yang, M. Duan, and K. Li, "A data balancing approach based on generative adversarial network," *Future Generation Computer Systems*, vol. 141, pp. 768–776, 2023.
- [22] C. M. Ahmed, G. R. M R, and A. P. Mathur, "Challenges in Machine Learning based approaches for Real-Time Anomaly Detection in Industrial Control Systems," in *Proceedings of the 6th ACM on Cyber-Physical System Security Workshop*, ser. CPSS '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 23–29.
- [23] D. Chakrabarty. (2019) CS 31: Algorithms (Spring 2019): Lecture 19.
- [24] D. I. Urbina, J. A. Giraldo, A. A. Cardenas, N. O. Tippenhauer, J. Valente, M. Faisal, J. Ruths, R. Candell, and H. Sandberg, "Limiting the Impact of Stealthy Attacks on Industrial Control Systems," ser. CCS '16. New York, NY, USA: ACM, 2016, p. 1092–1105.
- [25] A. Howe, D. Peasley, and M. Papa, "Graph autoencoders for detecting anomalous intrusions in OT networks through dynamic link detection," in *2024 IEEE 21st Consumer Communications & Networking Conference (CCNC)*, 2024, pp. 1–6.
- [26] D. L. Reyes, A. Perez-Pons, and R. B. Dean, "Anomaly detection in embedded devices through hardware introspection," in *2023 Silicon Valley Cybersecurity Conference (SVCC)*, 2023, pp. 1–7.
- [27] Kalanit Suzan Segal and Hadar Cochavi Gorelik and Oleg Brodt and Yuval Elbahar and Yuval Elovici and Asaf Shabtai, "Uefi memory forensics: A framework for uefi threat analysis," 2025.
- [28] M. R. G. Raman and A. P. Mathur, "A hybrid physics-based data-driven framework for anomaly detection in industrial control systems," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 52, no. 9, pp. 6003–6014, 2022.
- [29] Y. K. Saheed, S. Misra, and S. Chockalingam, "Autoencoder via DCNN and LSTM models for intrusion detection in industrial control systems of critical infrastructures," in *2023 IEEE/ACM 4th International Workshop on Engineering and Cybersecurity of Critical Systems (EnCyCris)*, 2023, pp. 9–16.
- [30] S. Adepu and A. Mathur, "Distributed Detection of Single-Stage Multipoint Cyber Attacks in a Water Treatment Plant," in *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security*, ser. ASIA CCS '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 449–460.
- [31] C. H. Yoong, V. R. Palleti, R. R. Maiti, A. Silva, and C. M. Poskitt, "Deriving invariant checkers for critical infrastructure using axiomatic design principles," *Cybersecur.*, vol. 4, no. 1, p. 6, 2021.
- [32] M. Mehmood, Z. Baig, and N. Syed, "Securing industrial control systems (ics) through attack modelling and rule-based learning," in *2024 16th International Conference on COMMunication Systems & NETWORKS (COMSNETS)*, 2024, pp. 598–602.
- [33] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17. Red Hook, NY, USA: Curran Associates Inc., 2017, p. 4768–4777.
- [34] M. T. Ribeiro, S. Singh, and C. Guestrin, "“Why Should I Trust You?”: Explaining the Predictions of Any Classifier," *CoRR*, vol. abs/1602.04938, 2016.
- [35] D. W. Apley and J. Zhu, "Visualizing the Effects of Predictor Variables in Black Box Supervised Learning Models," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, vol. 82, no. 4, pp. 1059–1086, 06 2020.
- [36] M. Sundararajan, A. Taly, and Q. Yan, "Axiomatic Attribution for Deep Networks," *CoRR*, vol. abs/1703.01365, 2017.
- [37] A. Avdalovic, J. Khoury, A. Taha, and E. Bou-Harb, "Enhancing Network Security Management in Water Systems using FM-based Attack Attribution," 2025.
- [38] C. Fung, E. Zeng, and L. Bauer, "Attributions for ML-based ICS Anomaly Detection: From Theory to Practice," in *31st Annual Network and Distributed System Security Symposium, NDSS San Diego, USA, February 26 - March 1, 2024*. The Internet Society, 2024.
- [39] M. A. Umer, A. Mathur, K. N. Junejo, and S. Adepu, "Generating invariants using design and data-centric approaches for distributed attack detection," *International Journal of Critical Infrastructure Protection*, vol. 28, p. 100341, 2020.
- [40] Q. Zhu, Y. Ding, J. Jiang, and S.-H. Yang, "Anomaly detection using invariant rules in Industrial Control Systems," *Control Engineering Practice*, vol. 154, p. 106164, 2025.
- [41] J. Su, C. Jiang, X. Jin, Y. Qiao, T. Xiao, H. Ma, R. Wei, Z. Jing, J. Xu, and J. Lin, "Large Language Models for Forecasting and Anomaly Detection: A Systematic Literature Review," 2024.