

# Compositional Probabilistic Analysis of Temporal Properties Over Stochastic Detectors

Ivan Ruchkin, *Member, IEEE*, Oleg Sokolsky, *Member, IEEE*, James Weimer, *Member, IEEE*, Tushar Hedao, and Insup Lee, *Fellow, IEEE*

**Abstract**—Runtime monitoring is a vital part of safety-critical systems. However, early stage assurance of monitoring quality is currently limited: it relies either on complex models that might be inaccurate in unknown ways or on data that would only be available once the system has been built. To address this issue, we propose a compositional framework for modeling and analysis of noisy monitoring systems. Our novel 3-value detector model uses probability spaces to represent atomic (noncompositional) detectors, and it composes them into a temporal logic-based monitor. The error rates of these monitors are estimated by our analysis engine, which combines symbolic probability algebra, independence inference, and estimation from labeled detection data. Our evaluation on an autonomous underwater vehicle found that our framework produces accurate estimates of error rates while using only detector traces, without any monitor traces. Furthermore, when data are scarce, our approach shows higher accuracy than noncompositional data-driven estimates from monitor traces. Thus, this article enables accurate evaluation of logical monitors in early design stages before deploying them.

**Index Terms**—Cyber-physical systems (CPSs), detection algorithms, formal languages, parameter estimation, probabilistic logic.

## I. INTRODUCTION

**M**ANY cyber-physical systems (CPS) operate in challenging safety-critical contexts, such as roadways, urban airspaces, and coastal waters. Not only is the control of the CPS expected to be safe but also safety monitoring is necessary to trigger fail-safes and notify human operators. This monitoring typically detects whether the CPS has experienced a fault/failure or finds itself in an environment where it may be unsafe. Such environments include plants with hardware malfunctions, severe weather conditions, and unmapped physical locations.

In the last decade, the runtime verification research has developed formal techniques and tools to increase confidence in the system's safety [1], [2]. Logically and temporally related

events can be monitored based on specifications in linear temporal logic (LTL) [3] and its descendants. When observations violate these specifications, the conditions are possibly unsafe and the system should execute a fail-safe fallback maneuver (e.g., stopping or landing).

Runtime safety monitors rely on inputs from noisy and potentially unreliable perception mechanisms. For example, the algorithms processing visual sensor data may randomly fail to detect features, or instead report nonexistent artifacts [4]. Furthermore, some perception systems are based on statistical tests, which at times may not have sufficient data for a conclusive answer. Such detectors feed stochastic errors into monitors, which become difficult to analyze only using non-deterministic logical constructs: they do not distinguish the more and less likely scenarios. Thus, we require *probabilistic guarantees for logical monitors*. These guarantees often take the form of false-positive and false-negative rates.

Unfortunately, the existing work falls short of providing such guarantees given perception uncertainties, especially in early design stages when data are scarce. One way to address this problem is model driven: probability constructs are embedded into a logical property (e.g., in probabilistic computation tree logic (PCTL) [5]), and the property is model checked on a probabilistic dynamics model. This method requires a detailed closed-loop model of the whole system—not only of the perception mechanism but also of the controller and the environment. Such models are rarely available in practice, and even when available they tend to be inaccurate, hard to validate, and lack scalability.

Another path to probabilistic monitoring guarantees is directly through the testing/execution data: safety monitors are treated as black-box binary random variables, thus turning the problem into the observation of the Bernoulli trials and estimating their error probabilities using standard statistical techniques [6]. In this case, the binary variable representation unnecessarily abstracts away our knowledge about the perception uncertainties and the logical structure of the monitor. This method requires substantial and specific data that may be difficult to collect; e.g., the system needs to be repeatedly put in an unsafe state to calculate the false-negative rate of a safety monitor. Usually, such data are unavailable in early design stages.

In this article, we present an approach to *estimating the error rates of logic-based runtime monitors*. We combine the strengths of model-based and data-driven approaches by taking advantage of logical specifications and using only perception

Manuscript received April 17, 2020; revised June 17, 2020; accepted July 6, 2020. Date of publication October 2, 2020; date of current version October 27, 2020. This work was supported in part by the Air Force Research Laboratory and the Defense Advanced Research Projects Agency as part of the Assured Autonomy Program under Contract FA8750-18-C-0090; and in part by the Office of Naval Research under Contract N68335-19-C-0200. This article was presented in part at the International Conference on Embedded Software 2020 and appears as part of the ESWEEK-TCAD special issue. (Corresponding author: Ivan Ruchkin.)

The authors are with the Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104 USA (e-mail: ruchkin@cis.upenn.edu).

Digital Object Identifier 10.1109/TCAD.2020.3012643

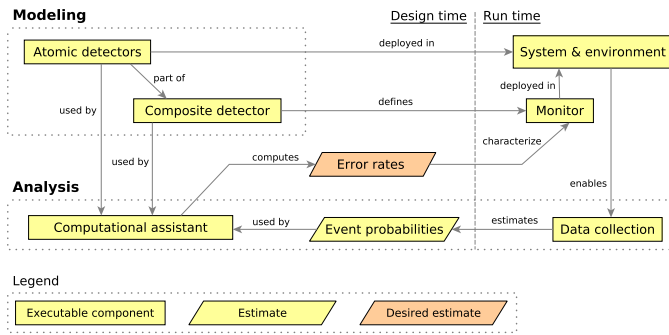


Fig. 1. Our modeling and analysis framework.

(nonmonitor) data—without any closed-loop system models. Thus, we can evaluate monitors before they are built.

Shown in Fig. 1, our approach consists of two parts: 1) *modeling* and 2) *analysis*. Inspired by statistical hypothesis tests [7], the modeling part focuses on binary ground-truth conditions (e.g., whether a vehicle faces an obstacle), which are true or false at each time. At the heart of our model are *stochastic detectors* of those ground truths with three outcomes: 1) true; 2) false; and 3) unknown. Logical monitors are modeled by composing, these detectors using LTL operators; the composed detection outcomes are determined by the 3-value<sup>1</sup> LTL [9], whereas the ground truths are composed by classic binary LTL.

The analysis part, performed offline by a heuristic computational assistant, treats a monitor’s error rate as a probability in an outcome space composed from the detector spaces, in accordance with the monitor’s logical definition. This probability is represented by a symbolic formula that can be decomposed into smaller probabilities of events associated with individual detectors. These probabilities (shown at the bottom of Fig. 1) are easier to estimate because they require only labeled data from the detectors, but not monitors.

We evaluated our approach on a simulated underwater vehicle that detects and follows an underwater pipeline. One safety property we monitor is that after losing sight of the pipeline, the vehicle rediscovers it within a given period of time; otherwise, the vehicle declares that it is lost, surfaces, and needs to be evacuated by a ship crew. We implemented a monitor for this property and applied our approach to estimate its false-negative rate (along with the false-positive rate of a different monitor, which checks whether the vehicle reliably follows the pipeline). This evaluation showed that our approach provides estimates on par with purely data-driven estimation, exceeds them when little data are available, reuses labeled detector data due to its compositional nature, and executes within seconds. However, the approach is dependent on accurate independence assumptions between detectors.

This article makes three contributions.

- 1) A probabilistic model of 3-outcome detectors (Section IV) based on a logic for these detectors (Section III).
- 2) An analysis of error rates based on probabilistic calculations and using labeled detection data (Section V).

<sup>1</sup>We do *not* mean the 3-value inconclusiveness due to incomplete traces [8].

- 3) An application of the above model and analysis to two monitors in a simulated underwater vehicle (Section VI).

This article is accompanied by a supplement [10] with additional derivations and illustrations. The source code and data can be found at <https://github.com/bisc/prob-comp-asst>.

The next section introduces a running example.

## II. MOTIVATING EXAMPLE

The DARPA assured autonomy program considers an unmanned underwater vehicle (UUV) that follows a continuous underwater pipeline and scans it for defects. To discover and track the pipeline, the vehicle uses two side-looking sonars pointing to the sides and downward from the vehicle. Thus, going above the pipeline in parallel to it provides the necessary data. The quality of the pipeline is investigated manually offline, after the UUV mission is complete, by processing the collected scan data via synthetic aperture imaging. The UUV is also equipped with a forward-facing sonar for obstacle avoidance.

The UUV software is comprised of a perception subsystem, a controller, and a safety monitor. The perception subsystem continuously processes the sensor data from both side-scan sonars. Since the data are noisy, the pipeline detector (represented as a predicate  $PI$  here) produces a potentially inaccurate 3-valued output at each moment.

- 1) Confident that the pipeline is visible (value T).
- 2) Confident that the pipeline is not visible (value F).
- 3) Not confident in either outcome (value U).

This output of  $PI$  is fed into the controller and the safety monitor. The former actuates the propeller and fins to follow the pipeline as well as possible. The goal of the latter is to determine when the UUV has definitively lost track of the pipeline and needs to be evacuated. The monitor is required to raise an alarm when the pipeline is not visible (i.e., F) for  $d \in [5, 30]$  s. To define this monitor in classic LTL, we map T to the Boolean truth value, and both F and U to the falsehood value. Thus, we write the following LTL formula, from which a monitor<sup>2</sup> can be automatically generated [2]:

$$\neg PI \rightarrow \diamond_{[1,d]} PI. \quad (1)$$

Our goal is to determine two probabilistic properties of the above monitor at design time: 1) the false-positive rate (FPR: “detected but not present”) and 2) the false-negative rate (FNR: “not detected but present”). High FPR is undesirable because it takes the effort of an entire vessel to evacuate the UUV. High FNR is a safety issue because if the UUV is lost for longer than  $d$  seconds, it may find itself in an area where the surfacing may be unsafe due to potential collisions (e.g., near a busy port). The next section describes our logical model for monitors.

## III. THREE-VALUE TEMPORAL LOGIC FOR DETECTORS

Our model of monitors has two aspects: 1) logical and 2) stochastic. The logical aspect (this section)

<sup>2</sup>A monitor raises an alarm when the formula it monitors evaluates to F.

specifies monitors as logical formulas over detector outputs. The stochastic aspect (next section) represents these monitors as probabilistic compositions of detectors. Both aspects focus on the outputs of detectors, abstracting away heterogeneous details of the exact functionality of the detectors.

A detector  $D$  is—abstractly for now—a random process that at each moment tries to determine the presence or absence of some hypothesis  $H_1$ . We model the output of  $D$  at a certain moment as a pair of random variables ( $DO, GT$ ).

- 1) *Detection outcome*  $DO \in \{T, F, U\}$  is the 3-value output of the detector: T (confident in  $H_1$ ), F (confident in the absence of  $H_1$ , i.e., the presence of null hypothesis  $H_0$ ), and U (not confident in either  $H_1$  or  $H_0$ ).
- 2) *Ground truth*  $GT \in \{T, F\}$  is the actual presence of  $H_1$ , which can either be true (T) or false (F).

Sampled at consecutive times, a detector produces a pair of finite equal-length sequences ( $s_{do}, s_{gt}$ ), where  $s_{do}$  is a sequence (or, trace) of  $DO$  values and  $s_{gt}$  is a sequence of  $GT$  values. For example,  $PI$  is a detector for  $H_1$ : “pipeline present.” For it,  $s_{gt}(t)$  is whether the pipeline is indeed present at time  $t$ , and  $s_{do}(t)$  is the 3-value detection output.

Monitors are specified in  $LTL_{3d}$ , a 3-value LTL for detectors with bounded temporal modalities. This logic is our modification of the 3-value LTL with unbounded modalities [9], which, in turn, builds on Kleene’s strong logic of indeterminacy  $K_3$  [11].  $LTL_{3d}$  adds bounds on modalities and two extra negations to handle uncertain values.  $LTL_{3d}$  considers fully known traces: the value U is orthogonal to the inconclusiveness from not knowing a complete trace [8].

Formulas in  $LTL_{3d}$  are constructed from atomic detectors  $D^1 \dots D^k$ , the truth constant T, three operators of negation ( $\neg_s, \neg_w, \neg_{se}$ ), conjunction  $\wedge$ , and bounded until  $U_{[m,n]}$  (where  $m, n \in \mathbb{N}_0$  and  $m < n$  for the rest of this section)

$$\varphi ::= \varphi U_{[m,n]} \varphi \mid \varphi \wedge \varphi \mid \neg_s \varphi \mid \neg_w \varphi \mid \neg_{se} \varphi \mid D \mid T$$

where  $D \in \{D^1 \dots D^k\}$ . These operators conventionally define other syntax: disjunction  $\vee$ , three implications  $\rightarrow_s, \rightarrow_w,$  and  $\rightarrow_{se}$ , and modalities eventually  $\diamond$  and always  $\square$ ; their definitions are available in the supplement [10].

Any  $LTL_{3d}$  formula  $\varphi$  can be evaluated in two ways. The first evaluation is over the values of  $DO$ , calculating the 3-value detection results. The second evaluation is over the values of  $GT$ , characterizing the binary ground-truth “oracle” (and coinciding with the standard LTL semantics). Both evaluations happen on a finite state trace  $s = s_1 s_2 \dots s_L$ , where  $L \in \mathbb{N}$  is the trace length. We assume we always have a trace of sufficient length to evaluate a given formula:  $L \geq n$ . A tail of trace  $s$  by starting from position  $t < L$  is denoted as  $s(t..) = s_t s_{t+1} s_{t+2} \dots s_L$ .

We formulate the 3-value evaluation first, denoted  $\llbracket \varphi \rrbracket_s$ . Here, the state trace  $s$  is a collection of the  $DO$  traces  $s_{do}^1 \dots s_{do}^k$  of the respective detectors

$$\begin{aligned} \llbracket T \rrbracket_s &= T \\ \llbracket D^j \rrbracket_s &= s_{do}^j(1), \text{ where } j \in [1, k] \\ \llbracket \neg_s \varphi \rrbracket_s &= \begin{cases} T & \text{if } \llbracket \varphi \rrbracket_s = F \\ F & \text{if } \llbracket \varphi \rrbracket_s = T \\ U & \text{if } \llbracket \varphi \rrbracket_s = U \end{cases} \end{aligned}$$

$$\llbracket \neg_w \varphi \rrbracket_s = \begin{cases} T & \text{if } \llbracket \varphi \rrbracket_s = F \text{ or } \llbracket \varphi \rrbracket_s = U \\ F & \text{if } \llbracket \varphi \rrbracket_s = T \\ U & \text{never} \end{cases}$$

$$\llbracket \neg_{se} \varphi \rrbracket_s = \begin{cases} T & \text{if } \llbracket \varphi \rrbracket_s = F \\ F & \text{if } \llbracket \varphi \rrbracket_s = T \text{ or } \llbracket \varphi \rrbracket_s = U \\ U & \text{never} \end{cases}$$

$$\llbracket \varphi_1 \wedge \varphi_2 \rrbracket_s = \begin{cases} T & \text{if } \llbracket \varphi_1 \rrbracket_s = T \text{ and } \llbracket \varphi_2 \rrbracket_s = T \\ F & \text{if } \llbracket \varphi_1 \rrbracket_s = F \text{ or } \llbracket \varphi_2 \rrbracket_s = F \\ U & \text{otherwise} \end{cases}$$

$$\llbracket \varphi_1 U_{[m,n]} \varphi_2 \rrbracket_s = \begin{cases} T & \text{if } \exists i \in [m, n] \cdot \llbracket \varphi_2 \rrbracket_{s(i..)} = T \\ & \text{and } \forall j \in [m, i) \cdot \llbracket \varphi_1 \rrbracket_{s(j..)} = T \\ F & \text{if either } \forall i \in [m, n] \cdot \llbracket \varphi_2 \rrbracket_{s(i..)} = F \\ & \text{or } \exists i \in [m, n] \cdot \llbracket \varphi_1 \rrbracket_{s(i..)} = F \text{ and} \\ & \quad \forall j \in [m, i) \cdot \llbracket \varphi_2 \rrbracket_{s(j..)} = F \\ U & \text{otherwise.} \end{cases}$$

*Definition 1 (3-Value Satisfaction):* Formula  $\varphi$  is 3-value satisfied on state trace  $s$ , denoted  $s \models_3 \varphi$ , if  $\llbracket \varphi \rrbracket_s = T$ .

A few remarks are in order. Similarly to  $K_3$ , we interpret U as a value of insufficient information, leading to the semantics that take advantage of the known information, e.g.,  $U \wedge F \equiv F$ . We adopt this principle because CPS often need to make decisions with limited information rather than waiting for the complete picture to be known. Similarly to the 3-value LTL [9], our Until operator handles uncertainty in the values of its operands—not in whether the available trace values are sufficient to evaluate the Until.

Unlike the existing 3-value logics, our  $LTL_{3d}$  has several negations. Operator  $\neg_s$  is equivalent to the negation in  $K_3$ : it flips the certain outcomes (T, F) while preserving the uncertain value. The two other negations, *weak*  $\neg_w$  and *strong exclusive*  $\neg_{se}$ , are never uncertain, with  $\neg_w$  erring on the side of  $H_1$  and  $\neg_{se}$  erring on the side of  $H_0$ .

The 2-value evaluation ( $\llbracket \varphi \rrbracket_s$ ) is based on the  $s_{gt}$  traces  $s_{gt}^1 \dots s_{gt}^k$  of the respective detectors. The semantics are the same as for  $\llbracket \varphi \rrbracket_s$  except not using the U values: atomic detectors can only evaluate to T or F because here the state trace  $s$  is based on the ground-truth traces

$$\llbracket D^j \rrbracket_s = s_{gt}^j(1), \text{ where } j \in [1, k].$$

*Definition 2 (2-Value Satisfaction):* Formula  $\varphi$  is 2-value satisfied on state trace  $s$ , denoted  $s \models_2 \varphi$ , if  $\llbracket \varphi \rrbracket_s = T$ .

*Remark 1:* We use bold operators (e.g.,  $\wedge, \neg_s,$  and  $U$ ) where *both* 3-value and 2-value interpretations are possible; nonbold operators (e.g.,  $\wedge, \neg$ ) have the usual 2-value interpretation.

Relations  $\models_3$  and  $\models_2$  are used to define a monitor’s output and ground truth, respectively.

*Definition 3 (Monitor):* A monitor  $M$ , defined by formula  $\varphi$  in  $LTL_{3d}$  over detectors  $D^1 \dots D^k$ , is a function that maps the detector traces  $(s_{do}^1, s_{gt}^1) \dots (s_{do}^k, s_{gt}^k)$  to a pair of monitor traces  $(m_{do}, m_{gt})$

$$\begin{aligned} m_{do}(t) &= s_{gt}^1(t..) \dots s_{gt}^k(t..) \models_3 \varphi \\ m_{gt}(t) &= s_{gt}^1(t..) \dots s_{gt}^k(t..) \models_2 \varphi. \end{aligned}$$

Notice that detectors produce a 3-value  $s_{do}$  and a 2-value  $s_{gt}$ , whereas monitors produce two 2-value traces:  $m_{do}$  and  $m_{gt}$ . In

all cases, we interpret the value T an “alarm,” and F and U as “no alarm.”

To sum up, the ground-truth values are substituted into logical formulas in the 2-value case, and the detector output values are substituted in the 3-value case.

With the logical aspect defined, we rewrite the motivating property (1) in  $LTL_{3d}$ . We use  $\neg_w$  to be conservative: the unknown presence of pipeline (U) is interpreted as its loss

$$\varphi_{pr} := \neg_w \mathbf{PI} \rightarrow_w \diamond_{[1,d]} \mathbf{PI}. \quad (2)$$

We define a pipe recovery monitor  $M_{pr}$  to alarm only when *certain* about *violations* of  $\varphi_{pr}$ . Thus, we use  $\neg_s$  to indicate this requirement of certainty

$$M_{pr} = \neg_s \varphi_{pr} = \neg_s (\neg_w \mathbf{PI} \rightarrow_w \diamond_{[1,d]} \mathbf{PI}). \quad (3)$$

We implemented  $M_{pr}$  over the  $s_{do}$  trace of detector  $\mathbf{PI}$ , and it raises an alarm when T occurs in its  $m_{do}$ . The GT trace  $m_{gt}$  is calculated from  $s_{gt}$  of  $\mathbf{PI}$ . Monitor code generation relies on the standard approaches [2] and is out of this article’s scope.

A *monitoring error* occurs if, for some  $t$ ,  $m_{do}(t) \neq m_{gt}(t)$ . We are interested in two types of errors: if  $m_{do}(t) = T$  and  $m_{gt}(t) = F$ , it is a *false positive*; if  $m_{do}(t) = F$  and  $m_{gt}(t) = T$ , it is a *false negative*. Since detectors and monitors are random variables, we talk about probabilities, or *rates*, of either type of errors.

Having monitors with high error rates can be a safety-critical concern, as illustrated on  $M_{pr}$  in Section II. Typically, these rates are evaluated by collecting, labeling, and analyzing their output data. The data are collected over multiple executions, each providing one  $m_{do}$ , labeling which yields one  $m_{gt}$ . Then, the error rate estimate is the count of errors (e.g., false positives) divided by the count of cases where this type of error (e.g., a false positive) was possible (e.g., all true negatives).

We believe that this estimation can be done in *early design stages* and using *only* the data from the atomic detectors (without executing the monitors). The central problem of this article is, hence, to estimate the error rates of a monitor of a  $LTL_{3d}$  formula. This estimation relies on labeled traces of atomic detectors (paired  $s_{do}$  and  $s_{gt}$ ).

#### IV. PROBABILISTIC COMPOSITION OF DETECTORS

This section models detectors as probability spaces behind the random variables ( $DO$ ,  $GT$ ), described in the previous section. Then, we show how compositions of these spaces correspond to formulas in  $LTL_{3d}$ , in order to analyze these formulas probabilistically in Section V.

##### A. Stochastic Detectors

We start by defining a probability space of one detector.

*Definition 4 (Atomic Stochastic Detector):* An *atomic stochastic detector* ( $\mathbf{D}$ ) is a pair of random variables ( $DO$ ,  $GT$ ) over a probability space  $(\Omega, \mathcal{F}, \Pr)$ .

- 1)  $\Omega$  is an elementary outcome space, defined as the set of all six possible pairs of values for ( $DO$ ,  $GT$ ): (T, T), (T, F), (F, T), (F, F), (U, T), (U, F).
- 2)  $\mathcal{F} = 2^\Omega$  is a sigma-algebra of events over signature  $\Omega$ , defined as the powerset of  $\Omega$  (including the empty set).

We use the following *marginal events* as shortcuts for the individual values of either random variable<sup>3</sup>:

$$\begin{aligned} \text{gtt}(\mathbf{D}) &= (*, T) = \{(T, T), (F, T), (U, T)\} \\ \text{dot}(\mathbf{D}) &= (T, *) = \{(T, T), (T, F)\} \\ \text{gtf}(\mathbf{D}) &= (*, F) = \{(T, F), (F, F), (U, F)\} \\ \text{dof}(\mathbf{D}) &= (F, *) = \{(F, T), (F, F)\} \\ \text{dou}(\mathbf{D}) &= (U, *) = \{(U, T), (U, F)\}. \end{aligned}$$

- 3)  $\Pr$  is a single-detector discrete probability measure over  $\mathcal{F}$ , subject to the usual Kolmogorov axioms.

Our concept of a stochastic detector encompasses a wide range of algorithms with ternary outcomes over time series, including anomaly detectors, machine learning classifiers, and sensor fusion algorithms. We have conveniently abstracted away the inputs of these algorithms, making this model broadly applicable. As mentioned in Section III, our convention is that a detector raises an alarm iff event dot occurs, thus introducing the asymmetry inherent in statistical tests (which reject  $H_0$  only with sufficient evidence).

Now, we handle time and multiple detectors. When detector  $\mathbf{D}$  produces readings over time points  $t_1 \dots t_k$  (e.g., a UUV gets consecutive readings of the pipeline), we model it with a sample of  $k$  “instances” of the original detector:  $\mathbf{D}^1 \dots \mathbf{D}^k$ , i.e., a stochastic process. These copies have their own respective variables  $DO^1 \dots DO^k$  and  $GT^1 \dots GT^k$ . These instances are not necessarily independent identically distributed; in general, their individual probability spaces are coupled to form a larger, *multidetector probability space*, which we use to describe compositions.

*Definition 5 (Multidetector Probability Space):* For detectors  $\mathbf{D}^1 \dots \mathbf{D}^k$ , a *multidetector probability space*  $(\Omega, \mathcal{F}, \Pr)$  is defined as follows.

- 1)  $\Omega = \Omega^1 \times \dots \times \Omega^k$  is a  $k$ -dimensional space of elementary outcomes; each dimension produces an elementary outcome pair for the respective detector.
- 2)  $\mathcal{F} = 2^\Omega$  is an event sigma-algebra, which is a powerset of the outcome space  $\Omega$  (including the empty set). Thus, an event is a set of  $k$ -sized vectors.
- 3)  $\Pr$  is a global probability measure over  $\mathcal{F}$ , extended from  $\Pr^1 \dots \Pr^k$  in a manner consistent with them (this requirement is formalized in the supplement [10]).

For a given set of atomic detectors, there are infinite multidetector spaces because individual probabilities  $\Pr^1 \dots \Pr^k$  can be coupled in many ways to form the joint  $\Pr$ . Analysis of possible couplings is one of the tasks we handle in Section V.

Now, consider  $k$  different<sup>4</sup> detectors for  $n$  time points, leading to a  $k \times n$  multidetector space. Although we can handle the general joint distribution from Definition 5, for tractability it is common to impose domain-specific constraints on  $\Pr$ . A common type of them is independence constraints, denoted  $A \perp\!\!\!\perp B$  for random variables  $A$  and  $B$ ; a set of independence constraints is denoted  $\mathcal{I}$ . Detectors of a random walk would have independent sequential actions:  $A^i \perp\!\!\!\perp A^{i+1}$ . This means

<sup>3</sup>By definition, events dof, dot, and dou partition  $\Omega$ . The same applies to gtt and gtf. Any such partitioning events constitute a well-formed detector.

<sup>4</sup>The detectors may be of the same phenomena or different ones. We address the general case where all ground truths are different.

that *any combination* of marginal events produced by these variables is independent (as expanded in the supplement [10]).

One of the common measurement models is that the error is *stateless* between measurements, meaning that previous measurements do not provide any extra information<sup>5</sup> given the current ground truths. In our terms, this implies the independence of any sequence of outcomes given their respective ground truths—an assumption we will use in Section V

$$\forall i, j : \mathbb{N}, 1 \leq i < j \leq k \cdot \\ DO^i \perp\!\!\!\perp DO^{i+1} \perp\!\!\!\perp \dots \perp\!\!\!\perp DO^j \mid GT^i \dots GT^j. \quad (4)$$

*Remark 2:* For convenience, we will perform symbolic calculations over marginal events introduced in Definition 4. These events can be connected with operators from classic binary logic:  $\neg$ ,  $\wedge$ , and  $\vee$ . We use nonbold fonts for them to interpret as usual (e.g.,  $\wedge$  means that both events occurred). These operators correspond to set operations over events in  $\mathcal{F}$ : complement, intersection, and union, respectively.

The above probabilistic model expresses any error rate, denoted  $\text{er}$ , straightforwardly as a conditional probability over marginal events. This article focuses on two specific error rates.

- 1) *False-Positive Rate:*  $\text{fpr}(\mathbf{D}) = \Pr(\text{dot}(\mathbf{D}) \mid \text{gtf}(\mathbf{D}))$ .
- 2) *False-Negative Rate:*  $\text{fnr}(\mathbf{D}) = \Pr(\neg \text{dot}(\mathbf{D}) \mid \text{gtt}(\mathbf{D}))$ .

Our model allows error rates to change between detectors  $\mathbf{D}^1 \dots \mathbf{D}^n$ . However, in this article, we assume the error rates to be constant across time moments because here they represent the inherent qualities of detectors. Mathematically, this means that our model averages out temporal variation in an error rate.

### B. Composition of Stochastic Detectors

Now, we will show how to interpret formulas in  $\text{LTL}_{3d}$  as compositions of detector spaces. Intuitively, a composite detector defined by a formula is a result of applying this formula's evaluation to the atomic detectors. The composite outputs are distributed over a multidetector probability space.

*Definition 6 (Composite Detector):* A composite detector  $\mathbf{D}'$  for formula  $\varphi$  over detectors  $\mathbf{D}^1 \dots \mathbf{D}^k$  with state trace  $s$  is a pair of random variables ( $DO'$ ,  $GT'$ ) over a multidetector probability space for  $\mathbf{D}^1 \dots \mathbf{D}^k$ , with these values at time  $t$

$$DO'(t) = \llbracket \varphi(\mathbf{D}^1 \dots \mathbf{D}^k) \rrbracket_{s(t..)} \\ GT'(t) = \llbracket \varphi(\mathbf{D}^1 \dots \mathbf{D}^k) \rrbracket_{s(t..)}.$$

Our insight here is that logical formulas correspond to composite detectors in multidetector spaces. For example,  $\varphi_{pr}$  and  $\mathbf{M}_{pr}$  are composite detectors over  $\mathbf{PI}, \mathbf{PI}^1 \dots \mathbf{PI}^d$  defined in a  $(d+1)$ -dimensional event space, according to (2) and (3). These spaces can be analyzed probabilistically, yielding estimates of error rates.

Given a fully defined  $\text{Pr}$  in some multidetector space, a logical formula determines the probabilities of events of the corresponding composite detector. In our experience, however, a fully defined probability measure on a multidetector space

<sup>5</sup>A slightly stronger assumption, common in time-series analysis [12], is written  $DO^i = GT^i + \epsilon_t$ , where  $\epsilon_t$  is a white noise series.

is rarely available or needed: it would require a lot of data and/or prior knowledge to determine. Instead, this space can be manipulated symbolically (i.e., without assigning numeric values) to express the desired probabilities through the known or easy-to-measure ones (e.g., probabilities of the atomic events), without committing to a fully defined  $\text{Pr}$ . These symbolic manipulations are convenient to carry out in the classic binary propositional logic, so we calculate using probability formulas over event predicates.

Let us provide some examples to illustrate event predicates in compositions, say conjunction. Its marginal events are derived from the logical semantics to be as follows<sup>6</sup>:

$$\begin{aligned} \text{gtt}(\mathbf{D}^a \wedge \mathbf{D}^b) &= \text{gtt}(\mathbf{D}^a) \wedge \text{gtt}(\mathbf{D}^b) \\ \text{gtf}(\mathbf{D}^a \wedge \mathbf{D}^b) &= \text{gtf}(\mathbf{D}^a) \vee \text{gtf}(\mathbf{D}^b) \\ \text{dot}(\mathbf{D}^a \wedge \mathbf{D}^b) &= \text{dot}(\mathbf{D}^a) \wedge \text{dot}(\mathbf{D}^b) \\ \text{dof}(\mathbf{D}^a \wedge \mathbf{D}^b) &= \text{dof}(\mathbf{D}^a) \vee \text{dof}(\mathbf{D}^b) \\ \text{dou}(\mathbf{D}^a \wedge \mathbf{D}^b) &= \text{dou}(\mathbf{D}^a) \wedge \text{dou}(\mathbf{D}^b) \vee \\ &\quad \text{dou}(\mathbf{D}^a) \wedge \text{dot}(\mathbf{D}^b) \vee \\ &\quad \text{dou}(\mathbf{D}^b) \wedge \text{dot}(\mathbf{D}^a). \end{aligned}$$

Such rewritings for other compositions can be found in the supplement [10].

Another rewriting to note: in our finite-trace semantics, the temporal modalities “always” and “eventually” for integers  $m$  and  $n > m$  can be expressed with  $\wedge$  and  $\vee$  (which, in turn, can be analyzed based on their marginal events)

$$\begin{aligned} \square_{[m,n]} \mathbf{D} &::= \mathbf{D}^m \wedge \dots \wedge \mathbf{D}^n \\ \diamond_{[m,n]} \mathbf{D} &::= \mathbf{D}^m \vee \dots \vee \mathbf{D}^n. \end{aligned}$$

In summary, this section showed that an error rate of an  $\text{LTL}_{3d}$  formula can be represented as a probability of events in a multidetector space. We take advantage of this fact to calculate estimates of error rates in the next section.

## V. ANALYSIS OF ERROR RATES

The goal of our analysis is to compute a numeric estimate of the false-positive or false-negative rate of a monitor, such as  $\text{fnr}(\mathbf{M}_{pr})$ , which is treated as a probability in an event space  $\mathcal{F}$  over some detectors  $\mathbf{D}^1 \dots \mathbf{D}^k$ . This is challenging because, even though it is precisely defined, a computable formula for that probability is *not known a priori*—and there is generally an infinite number of them. Finding a computable formula is complicated by knowing probability values for some set of events and having traces for another set of events.

We factor this problem into three analyses; each takes different inputs and produces a numeric estimate  $\hat{\text{er}}$  of a given error rate  $\text{er}$ .

- 1) *Exact Compositional Calculation (ECC):* To compute  $\hat{\text{er}}$ , this analysis requires a monitor formula  $\mathbf{M}$ ,

<sup>6</sup>The form of  $\text{dou}(\mathbf{D}^a \wedge \mathbf{D}^b)$  follows from events  $\text{dot}(\mathbf{D}_a \wedge \mathbf{D}_b)$  and  $\text{dof}(\mathbf{D}^a \wedge \mathbf{D}^b)$ , being a complement of their union. Specifying the marginal events fully defines a detector, as shown in the supplement [10].

user-specified probability values  $\text{Pr}$  of selected events in  $\mathcal{F}$ , independence assumptions  $\mathcal{I}$ , and transformation rules  $\mathcal{R}$

$$\hat{\text{er}}(\mathbf{M}) = \text{ECC}_{\text{er}}(\mathbf{M}, \text{Pr}, \mathcal{I}, \mathcal{R}).$$

- 2) *Noisy Compositional Calculation (NCC)*: This analysis requires a monitor formula  $\mathbf{M}$ , labeled traces for atomic detectors (for simplicity represented by one trace pair:  $s_{gt}$  and  $s_{do}$ ), independence assumptions  $\mathcal{I}$ , a set of symbolic “preferred” probabilities  $\hat{\mathbb{P}}$  of user-specified events in  $\mathcal{F}$ , and transformation rules  $\mathcal{R}$  (which we elaborate in the next section). The preferred probabilities is a set of probability symbols ( $\text{Pr}$ ) over the events of  $\mathcal{F}$ . The user chooses  $\hat{\mathbb{P}}$  as a heuristic<sup>7</sup> for probabilities to express the error rate through. For example, we use  $\hat{\mathbb{P}} = \{\text{fpr}(\mathbf{PI}), \text{fnr}(\mathbf{PI})\}$  because the composite error rates can often be compactly expressed using the atomic error rates. The goal is to find a rate formula with a minimal number of nonpreferred probabilities and estimate these probabilities from the labeled detector traces

$$\hat{\text{er}}(\mathbf{M}) = \text{NCC}_{\text{er}}(\mathbf{M}, s_{gt}, s_{do}, \mathcal{I}, \hat{\mathbb{P}}, \mathcal{R}).$$

- 3) *Black-Box Calculation (BBC)*: This analysis uses only labeled monitor traces, represented with a pair  $(m_{gt}, m_{do})$ , produced by monitor  $\mathbf{M}$  whose rate is being estimated. This analysis does not need the  $\mathbf{M}$  formula: it estimates the error rate directly from the traces

$$\hat{\text{er}}(\mathbf{M}) = \text{BBC}_{\text{er}}(m_{gt}, m_{do}).$$

To perform BBC, we use the standard approach using count-based frequency estimation, exemplified in the second last paragraph of Section III, which is an optimal estimator for event probabilities [7]. Generally, to estimate an error rate  $\rho = \text{Pr}(A|B)$ , BBC is computed as follows:

$$\begin{aligned} & \text{BBC}_{\text{Pr}(A|B)}(m_{gt}, m_{do}) \\ &= \frac{|\{(m_{gt}(t), m_{do}(t)) \text{ s.t. } A(t) \wedge B(t)\}|}{|\{(m_{gt}(t), m_{do}(t)) \text{ s.t. } B(t)\}|}. \end{aligned} \quad (5)$$

This article proposes a novel semiautomated heuristic analysis to compute ECC and NCC. To support the analysis, we implemented a semiautomated computational assistant based on the Mathematica’s symbol manipulation (see its source code at <https://github.com/bisc/prob-comp-asst>).

### A. Transformation Rules

Our computational assistant symbolically transforms formulas and computes numeric estimates. The assistant starts with the formula of the desired error rate ( $\text{er}$ ) of a monitor and transforms it toward a computable formula. Here, “computable” means that each probability in this formula is either known or can be estimated from  $(s_{gt}, s_{do})$ .

This assistant manipulates formulas using transformation rules. Each transformation rule has this form

$$[G] A \rightarrow C, \text{ where}$$

<sup>7</sup>The reasons for choosing a specific set  $\hat{\mathbb{P}}$  may vary. Some probabilities have strong priors or more data for accurate estimation; others can limit the formula search or help find a small formula.

- 1)  $A$  is the antecedent to replace, e.g., it can be  $\text{Pr}(X, Y)$  where  $X$  and  $Y$  are placeholders for any events;
- 2)  $C$  is the consequent to replace the antecedent, e.g., it can be  $\text{Pr}(X)\text{Pr}(Y)$ , where  $X$  and  $Y$  are from  $A$ ;
- 3)  $G$  is the guard condition, such as  $X \perp\!\!\!\perp Y$ .

All transformation rules can be partitioned into four lists, the combination of which is denoted  $\mathcal{R}$ . The four lists are as follows (their rules are fully specified in the supplement [10]).

- 1)  $\mathcal{R}_{\text{log}}$  are logical tautologies of  $\text{LTL}_{3d}$ , e.g.,  $\neg_s \neg_s \mathbf{D} = \mathbf{D}$ .
- 2)  $\mathcal{R}_{\text{ev}}$  are rules substituting the logical operators of  $\text{LTL}_{3d}$  for their definitions in terms of detector events.
- 3)  $\mathcal{R}_{\text{prob}}$  are rules to algebraically manipulate probability expressions. This includes the rules to turn probabilities of event conjunctions into products of event probabilities, if these events are independent.
- 4)  $\mathcal{R}_{\text{indep}}$  are rules to infer conditional independence. We use the standard semigraphoid axioms: contraction, decomposition, weak union, and intersection [13].

Each of the four lists can be *applied* to a formula. The application of each list proceeds one rule at a time in the order of the list. A rule is applied to every matching subformula and until no subformulas match the antecedent and satisfy the guard. If/once a rule can no longer be applied, the next rule is read. If no rule is applicable, the list application terminates.

All rule lists use only sound transformations: assuming  $G$  holds, we know that  $A = C$ . We ensure that by using only the standard rules of probability, the Boolean algebra, conditional independence, as well as the easy-to-show tautologies of  $\text{LTL}_{3d}$ . Therefore, applying  $\mathcal{R}$  never produces a formula not equivalent to the original one,  $\text{er}(\mathbf{M})$ , under the assumptions  $\mathcal{I}$ .

The outcome of rule application significantly depends on the rule lists chosen by the user. While we provide large sets of rules, it is up to the user to pick a subset and order that would be effective for a given scenario. To do so, the user needs to understand the rules and their effects on probability expressions.

The presence and order of rules in  $\mathcal{R}$ , and particularly  $\mathcal{R}_{\text{prob}}$ , determine the completeness, termination, and algorithmic efficiency of ECC and NCC. Some rule lists may lack the necessary transformations (e.g., adding and removing negations to events) to produce a computable formula. Application of other rule lists may not terminate; e.g., replacing  $\text{Pr}(X)$  with  $\text{Pr}(X \wedge Y)$  requires a rule that introduces a variable  $Y$ , which could be applied indefinitely. Such rules can lead to a set of formulas that grows exponentially (to consider all combinations of events) and is possibly unbounded (if normal forms for event expressions are not enforced).

We expect a large number of formula-finding problems to be solvable with rule lists. Unfortunately, we do not have a precise characterization of this class, leaving it for future work. For the monitors and error rates in this article, we found sufficient and efficient rule lists, given in full detail in the supplement [10].

### B. Steps of ECC and NCC

The ECC and NCC analyses proceed in five steps, illustrated in Fig. 2. Their abstract description below is exemplified in Section V-C.

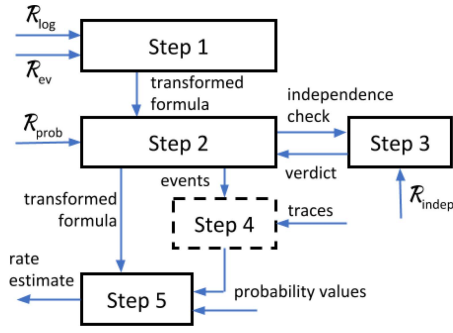


Fig. 2. ECC/NCC analysis steps.

*Step 1:* Rewrite the formula in terms of marginal events dot, dof, dou, gtt, and gtf—without any operators of  $LTL_{3d}$ .

*Inputs:*

- 1)  $er(\mathbf{M})$ —the goal error rate  $er$  of a monitor  $\mathbf{M}$ ;
- 2)  $\mathcal{R}_{log}$ —transformation rules for the logic operators;
- 3)  $\mathcal{R}_{ev}$ —transformation rules for events of detectors.

*Output:*  $f(\Pr(\mathcal{F}))$ —a formula  $f$  over probabilities of the Boolean combinations of marginal events of atomic detectors.

*Substeps:*

- 1.1 Apply rules  $\mathcal{R}_{log}$  to  $er(\mathbf{M})$  to simplify the composite detector formula. While not required, this substep can shorten the subsequent steps, e.g., by removing negation operators when valid (examples in the supplement [10]).
- 1.2 Apply rules  $\mathcal{R}_{ev}$  to the result of the previous step. This substep replaces the 3-value logical constructs with an equivalent encoding using replacing them with binary events dot, dof, dou, gtt, and gtf. The expression becomes more verbose but also manipulable with the classical Boolean logic.

*Step 2:* Algebraically reduce the formula to the preferred or known probabilities.

*Inputs:*

- 1)  $f(\Pr(\mathcal{F}))$ —a formula  $f$  over probabilities of Boolean combinations of marginal events of atomic detectors;
- 2)  $\mathcal{R}_{prob}$ —transformation rules for probabilities;
- 3) (ECC only)  $\Pr(\mathcal{F})$ —known probabilities of some given events in  $\mathcal{F}$ . The rest are unknown;
- 4) (NCC only):  $\hat{\Pr}(\mathcal{F})$ —a set of preferred probabilities.

*Output:*  $f(\Pr(\mathcal{F}))$ —a formula  $f$  over  $\Pr$  of events in  $\mathcal{F}$ .

*Substeps:*

- 2.1 Apply rules from  $\mathcal{R}_{prob}$ , with two caveats.
  - a) Do not apply rules to known/preferred probabilities.
  - b) If a rule has an independence guard, send this independence statement to step 3. If step 3 returns  $\top$ , apply the rule; otherwise, skip to the next rule.
- 2.2 After all  $\mathcal{R}_{prob}$  have been applied, proceed to this step.
  - a) For ECC, go to step 5;
  - b) For NCC, send the events under probabilities in the current formula  $f$  to step 4 (then go to step 5).

*Step 3:* Check of whether an independence statement hold under the independence assumptions.

*Inputs:*

- 1) an independence statement of form  $X \perp\!\!\!\perp Y$  or  $X \perp\!\!\!\perp Y|Z$ , where  $X$ ,  $Y$ , and  $Z$  are event expressions;
- 2)  $\mathcal{I}$ —a set of independence assumptions;

3)  $\mathcal{R}_{indep}$ —a list of transformation rules for independence.  
*Output:*  $\top$  (if the statement can be assumed to be true) or  $\text{F}$  (if not enough information to conclude independence).

*Substeps:*

- 3.1 Apply rules  $\mathcal{R}_{indep}$  to  $\mathcal{I}$ . This produces a larger set of independence statements assumed to be true;
- 3.2 Return  $\top$  if the given statement is in the produced set. Otherwise, return  $\text{F}$ .

*Step 4 (NCC Only):* Produce numeric probability estimates for the set of events requested in substep 2.2.

*Inputs:*

- 1) events  $\mathbb{E}$  in  $\mathcal{F}$ . requested in substep 2.2;
- 2) a labeled detector trace  $(s_{gt}, s_{do})$ .

*Output:* Probability estimates  $\hat{\Pr}$  for the requested events.

*Substeps:*

- 4.1 for each event in  $\mathbb{E}$ , calculate a count-based estimate analogously to (5) but over trace  $(s_{gt}, s_{do})$ ;
- 4.2 send all probability estimates  $\hat{\Pr}(\mathbb{E})$  to step 5.

*Step 5:* Calculate a numeric estimate  $\hat{er}$  of error rate  $er$ .

*Inputs:*

- 1) Formula  $f(\Pr(\mathcal{F}))$  from step 2.2;
- 2) (ECC only) known probabilities  $\Pr(\mathcal{F})$ ;
- 3) (NCC only) estimated probabilities  $\hat{\Pr}(\mathcal{F})$ .

*Output:* A numeric estimate of error rate  $\hat{er}(\mathbf{M})$ .

*Substeps:*

- 5.1 Replace each probability in  $f(\Pr(\mathcal{F}))$  with its value;
  - a) ECC uses known probability values  $\Pr(\mathcal{F})$ ;
  - b) NCC uses the estimates  $\hat{\Pr}(\mathcal{F})$  from step 4.
- 5.2 Calculate the numeric expression and return the result.

### C. Example of ECC/NCC Application

As an illustration, we walk through<sup>8</sup> the NCC calculation of  $\text{fnr}(\mathbf{M}_{pr})$ . As inputs we have the definition of  $\mathbf{M}_{pr}$  in (2) with some known value of  $d$ , a labeled trace  $(s_{gt}, s_{do})$  for  $\mathbf{PI}$ , and  $\hat{\Pr} = \{\text{fpr}(\mathbf{PI}), \text{fnr}(\mathbf{PI})\}$ , rules  $\mathcal{R}$ , and two independence assumptions: (4) and that only the current ground truth affects the detector output

$$\forall i : \mathbb{N}, 1 \leq i \leq k.$$

$$DO^i \perp\!\!\!\perp GT^1 \dots GT^{i-1}, GT^{i+1} \dots GT^k \mid GT^i. \quad (6)$$

The initial formula  $er(\mathbf{M})$  is as follows:

$$\text{fnr}(\neg_s(\neg_w \mathbf{PI} \rightarrow_w \diamond_{[1,d]} \mathbf{PI})). \quad (7)$$

Step 1.1 replaces  $\rightarrow_w$  with  $\vee$  and  $\neg_w$ , rewrites  $\diamond_{[1,d]}$  using  $\vee$ , and distributes  $\neg_s$  over the  $\vee$  operators in the parentheses. The outcome is a different detector under  $\text{fnr}$

$$\text{fnr}(\neg_w \mathbf{PI} \wedge \neg_s \mathbf{PI}^1 \wedge \dots \wedge \neg_s \mathbf{PI}^d). \quad (8)$$

Step 1.2 rewrites (8) as a conditional probability over the marginal events of the detector under  $\text{fnr}$ . Specifically, the rewriting uses events dot and gtf according to the definitions of operators  $\vee$ ,  $\neg_w$ , and  $\neg_s$

$$\begin{aligned} & (\text{dof}(\mathbf{PI}^0) \vee \text{dou}(\mathbf{PI}^0)) \wedge \text{dof}(\mathbf{PI}^1) \wedge \dots \wedge \text{dof}(\mathbf{PI}^d) \\ & \text{gtf}(\mathbf{PI}^0) \wedge \text{gtf}(\mathbf{PI}^1) \wedge \dots \wedge \text{gtf}(\mathbf{PI}^d). \end{aligned}$$

<sup>8</sup>The full derivation for this error rate formula is in the supplement [10].

Hence, step 1.2 produces this probability formula

$$\Pr\left(\neg\left(\left(\text{dof}(\mathbf{PI}^0) \vee \text{dou}(\mathbf{PI}^0)\right) \wedge \text{dof}(\mathbf{PI}^1) \wedge \dots \wedge \text{dof}(\mathbf{PI}^d)\right) \mid \text{gtf}(\mathbf{PI}^0) \wedge \dots \wedge \text{gtf}(\mathbf{PI}^d)\right). \quad (9)$$

Then step 2.1 switches the negation to  $1 - \Pr$

$$1 - \Pr\left(\left(\text{dof}(\mathbf{PI}^0) \vee \text{dou}(\mathbf{PI}^0)\right) \wedge \text{dof}(\mathbf{PI}^1) \wedge \dots \wedge \text{dof}(\mathbf{PI}^d) \mid \text{gtf}(\mathbf{PI}^0) \wedge \dots \wedge \text{gtf}(\mathbf{PI}^d)\right). \quad (10)$$

Now, step 2.1 matches a rule with an independence guard and sends the events under  $\Pr$  to step 3. step 3.2 finds that the assumed (4) matches the provided events and returns  $\mathbf{T}$ . Step 2.1 now replaces the above probability with a product of probabilities of each detector's events. Furthermore, step 2.1 reduces the conditioning of  $\text{dof}^i$  to  $\text{gtf}^i$  using a rule that checks (6). A few algebraic transformations in step 2.1 yield the *final formula* for  $\text{fnr}(\mathbf{M}_{pr})$

$$1 - (1 - \text{fpr}(\mathbf{PI}))(1 - \text{fpr}(\mathbf{PI}) - \Pr(\text{dou}(\mathbf{PI}) \mid \text{gtf}(\mathbf{PI})))^d. \quad (11)$$

Step 2.2 leads to step 4, which estimates the probabilities  $\text{fpr}(\mathbf{PI})$  and  $\Pr(\text{dou}(\mathbf{PI}) \mid \text{gtf}(\mathbf{PI}))$  from the trace  $(s_{gt}, s_{do})$  of  $\mathbf{PI}$ . Step 4.1 counts the events of interest [e.g., false positives for  $\text{fpr}(\mathbf{PI})$ ] and divides it by the counts of possibilities for false positives [i.e., ground-truth negatives for  $\text{fpr}(\mathbf{PI})$ ].

Finally, step 5.1 substitutes the estimates of  $\text{fpr}(\mathbf{PI})$  and  $\Pr(\text{dou}(\mathbf{PI}) \mid \text{gtf}(\mathbf{PI}))$  into (11). Step 5.2 yields a numeric estimate of  $\text{fnr}(\mathbf{M}_{pr})$ . We investigate the accuracy of such estimates in the next section.

To summarize, this section introduced three analyses of error rates. ECC consists of two stages: 1) symbolic manipulation (steps 1–3) and 2) numeric computation (step 5). NCC consists of three stages: 1) symbolic manipulation (steps 1–3); 2) estimation from data (step 4); and 3) numeric computation (step 5). BBC consists only of one step—estimation from data.

ECC/NCC can take advantage of their modularity: 1) prior knowledge about detectors can be used for accurate of; 2) that the symbolic manipulation can be performed once per monitor and reused should the data or numerical inputs change; and 3) labeled monitor traces are not required.

## VI. EVALUATION: UUV CASE STUDY

Our validation goals are fourfold; we evaluate: 1) the *accuracy* of predictions of monitor error rates by our approach (NCC) compared to the purely data-driven approach (BBC); 2) the *dependency* of this accuracy on the amount of provided data; 3) the *sensitivity* of our approach to the independence assumptions; and 4) the *computational costs* of our approach. Note that in 1), NCC is *not* intended to beat the accuracy of BBC because it is used on different inputs in different circumstances; instead, NCC aims to achieve a comparable accuracy without any traces  $(m_{gt}, m_{do})$  from monitors.

At a high level, our evaluation goes as follows. First, we specify and implement two monitors: 1)  $\mathbf{M}_{pr}$  in (3) and 2)  $\mathbf{M}_{rf}$  defined below in (13). Second, we collect multiple

labeled traces (outputs and ground truth) of each monitor and detector. Third, we produce three estimates of an error rate of each monitor: ECC uses the error rates of the atomic detector known-by-construction, NCC uses the labeled traces of the atomic detector, and BBC uses the labeled traces of monitors.

To evaluate 1), we compare the residuals of NCC and BBC relative to ECC, which in this case plays the role of the true monitor error rate. To evaluate 2), we compare these residuals at different amounts of data provided to NCC and BBC. To evaluate 3), we compare NCC and BBC on traces with sensor noise that invalidate our independence assumption in (4). Finally, to evaluate 4), we measure the time it takes to perform the five steps of NCC.

### A. Data Collection

We based our evaluation on the DARPA assured autonomy case study described in Section II. Data were collected from a ROS Gazebo UUV simulator ([https://github.com/uuvsimulator/uuv\\_simulator](https://github.com/uuvsimulator/uuv_simulator)) customized with underwater vehicle dynamics and pipeline generation. We executed 73 pipeline-following missions, totaling 7.7 h of simulation time (the data can be found at <https://github.com/bisc/prob-comp-asst>). Each mission started with the pipeline in view and terminated 30 s after the pipeline was no longer visible (either due to successful completion or getting lost). We kept all the system design parameters (including the controller) fixed across all missions, with the only variables being the random seed that determines the pipeline shape<sup>9</sup> and the initial position relative to the first pipeline segment.

During each mission, the pipeline detector  $\mathbf{PI}$  provided a trace of binary values, each indicating the presence of a pipeline in the UUV's view at that moment, at the rate of 1 Hz.  $\Pr(\text{dou})$  for  $\mathbf{PI}$  is set to 0 for simplicity. We used two (mutually exclusive) types of noise in  $\mathbf{PI}$ : 1) *stateless* and 2) *stateful*. The former has fixed  $\text{fpr}$  and  $\text{fnr}$  equal to 0.1 by construction, without any dependency on the past noise (only on the current GT). The latter is statefully dependent on the past five detection results, varying with the noise weight  $w \in [0, 1]$  from 0.1 at best to 0.5 at worst (a coin toss equivalent), simulating the propensity of noise to persist over time<sup>10</sup>

$$\begin{aligned} \text{fpr}(t) &= 0.1 + 0.4w \cdot \frac{\sum_{i:1..5} \mathbb{1}^{\left\langle DO^{t-i} = \mathbf{T} \right\rangle}}{5} \\ \text{fnr}(t) &= 0.1 + 0.4w \cdot \frac{\sum_{i:1..5} \mathbb{1}^{\left\langle DO^{t-i} = \mathbf{F} \right\rangle}}{5}. \end{aligned} \quad (12)$$

We investigated two monitors, chosen for their different syntactic structure.

- 1) The pipeline recovery monitor  $\mathbf{M}_{pr}$ , which has been described in Section II and formalized in (3)

$$\mathbf{M}_{pr} := \neg_s(\neg_w \mathbf{PI} \rightarrow_w \diamond_{[1,d]} \mathbf{PI}).$$

<sup>9</sup>The pipeline is always a noncircular and nonself-intersecting sequence of linear segments, and no segment connection angle is less than 90 degrees.

<sup>10</sup> $\mathbb{1}^{\langle a \rangle}$  is an indicator function; it returns 1 if  $a = \mathbf{T}$  and 0 otherwise.



- 2) A monitor  $M_{rf}$  for the “reliable following” detector  $\varphi_{rf} := \square_{[0,d]}PI$ , requiring that the pipeline is confidently perceived for  $d$  samples:

$$M_{rf} := \neg_s \square_{[0,d]}PI. \quad (13)$$

By feeding the ground-truth and noisy detector samples into each monitor, we obtained the  $m_{gt}$  and  $m_{do}$  traces respectively. Thus, for each of the two composite detectors, each mission provided four binary-valued traces of equal length: a ground truth/noisy pair of traces for the pipeline detection, and a ground truth/noisy pair of traces for the monitor.

All experiments were run on a Lenovo X1 laptop with Ubuntu 18.04 LTS, Intel Core i7-6600U CPU at 2.60 GHz, 16-GB DDR3 RAM, and an internal SATA SSD hard drive.

### B. Variables and Sampling

We estimate one rate for each monitor: FNR for  $M_{pr}$  using (11), and FPR for  $M_{rf}$  using a formula derived in the supplement [10]. We chose these specific rates because they have a larger range of values for different  $d$ , from 0.2 to over 0.95. The other two rates— $fpr(M_{pr})$  and  $fnr(M_{rf})$ —have low values, barely exceeding 0.01 at their highest.

Thus, our results have four *independent variables*: 1) a chosen monitor; 2) a mission number; 3) a deadline;<sup>11</sup> and 4) a noise weight  $w$ . Our primary *dependent variables* are ECC, NCC, and BBC. Given a monitor, ECC depends only on  $d$ , whereas NCC and BBC depend on  $d$  and a subset of missions from which the data are used. In all comparisons below, we always provided the same subset of missions to both NCC and BBC.

Our secondary dependent variables are the errors between the estimates and the amounts of information available in the traces. We use the root-mean-squared error (RMSE) as a standard measure between real-valued point estimates. For estimating FNR, the amount of information is determined by the number of gtt events of the monitor: BBC considers the FNR a probability parameter of a binomial random variable, and each gtt is a sample of this variable. Likewise, for estimating FPR, the information is measured by the count of gtf events.

The evaluation 2 samples across all nontrivial subsets of our trace dataset (sizes from 1 to 72) with replacement. We pick a fixed number (40) of randomly (uniformly) chosen subsets for each size. In the analyses related to RMSE, the averaging is done over fixed-sized bins for the amounts of information (i.e., the counts of monitor gtt or gtf events). The only exception to this sampling strategy is the first accuracy analysis, where we use our full dataset for an estimate.

### C. Results

1) *Evaluation 1—Accuracy*: We evaluated the accuracy of calculations for different values of deadline  $d$  using the full dataset with stateless noise ( $w = 0$ ). For both monitors, the differences between NCC, BBC, and ECC were small, with all three closely tracing each other (see the supplement [10]

<sup>11</sup>The “deadline” is the value of time bound  $d$  in both monitor definitions.




Fig. 3. Log-RMSE of FNR predictions ( $d = 10$ ) for  $M_{pr}$ , relative to the amount of information.




Fig. 4. NCC and BBC for FNR of  $M_{pr}$  ( $d = 10$ ) relative to weight  $w$ . Baseline ECC is shown for  $w = 0$ , for context.

for the figures). For FNR of  $M_{pr}$ , the error of BBC increases slightly for  $d > 20$ . The RMSE of NCC (relative to ECC) across all values of  $d \in [2, 30]$  is 0.015, and 0.013 for BBC (also relative to ECC). While NCC has a larger error, it stays consistent across all  $d$  values, whereas BBC accuracy varies because BBC needs data for each value of  $d$ . For larger  $d > 20$ , fewer datapoints trigger the monitor’s ground-truth alarm, making it harder for BBC to produce accurate estimates. For the FPR of  $M_{rf}$ , the differences were negligible: the average absolute error is on the order of 0.001 for both NCC and BBC.

2) *Evaluation 2—Data Dependency*: Here, we analyze how estimate error changes relative to the amount of information (keeping  $w = 0$ ,  $d = 10$  for illustration). Fig. 3 shows the RMSE of NCC and BBC for  $M_{pr}$  produced by sampling different data subsets. Each point in the plot is produced by averaging the squared error for a 25-wide bin,<sup>12</sup> and the minimum number of samples per bin is 37. This figure indicates that NCC is more accurate than BBC for small amounts of data, while the BBC gets increasingly more accurate for larger amounts. Notice a logarithmic scale: the small-data errors are orders of magnitude larger than the large-data errors. The sudden drop of BBC is a sampling artifact: BBC converges to the full-dataset estimate of FNR, which in this case happens to be close to ECC. No such phenomenon was observed for  $M_{rf}$ : NCC is more accurate for all data amounts, as shown in the

<sup>12</sup>Binned over the  $x$ -axis; each bin’s result is plotted at its lower bound.




Fig. 5. Dependency of computation times of two formula manipulations and data estimation on the formula size.

supplement [10], confirming that the BBC error is higher for larger deadlines.

3) *Evaluation 3—Sensitivity to Independence Assumptions:* In the previous two evaluations, the noise profile ( $w = 0$ ) matched our independence assumption (4). Here, though, we intentionally set  $w \in (0, 1]$  to disrupt this assumption and observe the effect on the NCC estimates. Fig. 4 shows that for full-data estimates NCC and BBC increase significantly for larger  $w$ . We do not have the ground-truth values for monitor error rates for  $w > 0$ : a true Markov model for their calculation needs the true detector event probabilities, which are simulated and not known exactly. Therefore, ECC is only shown to illustrate the increase in the FNR. NCC diverges from BBC by over 0.1 for  $w \geq 0.9$ , overestimating the FNR. The supplement [10] shows an analogous picture for  $M_{rf}$ .

4) *Evaluation 4—Computational Costs:* To measure the computation times of our approach, we consider its two parts: 1) *formula manipulation* (steps 1–3 and 5)<sup>13</sup> and 2) *estimation from data* (step 4). The former is performed in ECC and NCC and has indistinguishable performance in both cases. The latter is performed in NCC (for detector probabilities  $\hat{P}_r$ ) and BBC (for monitor rate estimates  $\hat{e}_r$ ).

The computation times for formula manipulation, shown in Fig. 5, increase with monitor formula size, linearly determined by parameter  $d$ . While small formulas take fractions of a second, the magnitude of growth varies with the formula and the rule list. The time to analyze  $\text{fpr}(M_{rf})$  never takes longer than a second, exhibiting a slow and linear growth. On the other hand, the analysis of  $\text{fnr}(M_{pr})$  exhibits a polynomial growth in complexity and takes close to 10 s for  $d = 30$ . We witnessed an exponential complexity growth for some rule lists that introduce many probability terms (e.g., by transforming disjunctions into sums of probabilities). The estimation of rates for monitors and detectors alike does not depend on the formula size and, for our dataset, takes approximately 1 s on average for one detector/monitor.

#### D. Interpretation of Results

NCC and BBC estimates benefit differently from getting more data. How useful extra data are depends on the probability of the GT events used in the respective calculation.

<sup>13</sup>Step 5 is virtually instantaneous and was timed together with the symbolic manipulation steps (1–3), which are computationally intensive.

For estimating  $\text{fnr}(M_{pr})$ , NCC uses events  $\text{gtf}(PI)$  [to estimate  $\text{fpr}(PI)$ ], which are 1.58 times more likely than the  $\text{gtt}(M_{pr})$  events (for  $d = 10$ ) used by BBC. Whereas for estimating  $\text{fpr}(M_{rf})$ , NCC uses events  $\text{gtt}(PI)$  [to estimate  $\text{fnr}(PI)$ ], which are 2.02 times more likely than the  $\text{gtf}(M_{rf})$  events (for  $d = 10$ ) used by BBC. Thus, NCC for  $M_{rf}$  gets a larger information advantage over BBC per trace, than in case of  $M_{pr}$ .

The absence of relevant ground-truth events puts the BBC at a significant disadvantage. We notice that larger  $d$  for  $M_{pr}$ , but not  $M_{rf}$ . This is explained by the reduced chance of pipe-loss events ( $\text{gtt}(M_{pr})$ ) as the definition of a pipe loss is relaxed by increasing  $d$ . On the other hand, satisfaction events for reliable-following ( $\text{gtf}(M_{pr})$ ) do not drop as drastically, since, for most of each trace, the UAV follows the pipeline successfully. As a result, NCC performs better than the BBC for  $d \geq 15$ . We conclude that composite detectors with infrequent violations would benefit from our approach for FNR estimation, whereas those with infrequent satisfactions should use our approach for FPR estimation.

The evaluation also highlights two benefits of compositions. First, the events for atomic detectors are typically more prevalent than those for monitors, making it easier to estimate and compose the error rates of detectors—as opposed to observing rare monitored events. Second, NCC can reuse the results of its steps. If the monitor formula changes, it can redo the formula manipulation and reuse the estimates  $\hat{P}_r$ . If the detectors or available data change, NCC can redo steps 4 and 5 and reuse the outcomes of symbolic manipulation for any value of  $d$ . In contrast, whenever the formula (e.g., the value of  $d$ ) changes, BBC needs new monitor traces and redoes the whole analysis. These benefits have the potential to reduce the data collection burden and speed up design space exploration: poor-quality monitors can be eliminated from the design or used to justify hardware/software upgrades.

#### E. Limitations

Our approach is indeed sensitive to the independence assumptions: the more they are violated, the more we deviate from the data-driven estimates. However, for  $M_{pr}$  and  $M_{rf}$ , assuming independence of events that are in fact dependent leads to upper-bounding the true error rate. This happens because independent events lead to smaller probabilities in (10) compared to a stateful dependency. Hence, in this case, our estimates are conservative upper bounds.

Therefore, one guideline is to carefully pick independence assumptions by consulting the domain and system experts (by asking them about stateful/stateless events, causal relations, etc.). For example, we could handle  $w > 0$  in the noise model in (12) by making the Markovian assumptions on detector  $PI$  such as  $DO^t \perp\!\!\!\perp DO^{t-j} \mid GT^t, DO^{t-1} \dots DO^{t-5}$  for  $j > 5$ , instead of (4). This would change the FNR calculation, requiring estimation of, e.g.,  $\Pr(\text{dof}(PI^t) \mid \text{gtf}(PI^t), \text{dof}(PI^{t-1}) \dots \text{dof}(PI^{t-5}))$ . Our analysis procedure would remain the same.

Another way to handle independence assumptions is to validate them on data, e.g., by informally comparing probabilities

or doing chi-squared tests. Any way of validating these assumptions would fit into our framework.

The produced error rate estimates are not invariant to some changes in the system. Examples include changing the sampling rate and retraining the perception. To update the estimates for a new perception, data collection and step 4 of our analysis need to be repeated. However, the  $PI$  rate estimates we produced would be invariant to deploying a higher-performance controller, which may increase  $\Pr(\text{gtt}(PI))$  but not change  $\text{fpr}(PI)$ . We plan to investigate reusing the part of the data and calculations for limited changes in the system.

The relative errors of NCC and BBC will vary for other monitors and systems. Simple monitors of low-variance systems may not benefit from our approach. For complex systems though, our approach could avoid prohibitively difficult testing of rare monitored conditions. Either way, we expect our qualitative observations to still hold.

## VII. RELATED WORK

The existing work related to this article can be grouped into two large categories: 1) detection and monitoring of uncertain events and 2) probabilistic and logical reasoning.

### A. Detection, Monitoring, and Error

A combination of detection algorithms is a well-explored technique in many areas, giving rise to a variety of relations between the objects of composition. In perception, sensor fusion reconciles multiple uncertainties about an observed phenomenon. Popular approaches, such as the Kalman filtering and the Bayesian methods, link the inputs via process and measurement models [14], whereas we do not rely on such models and therefore only require specific knowledge about certain probabilities. In machine learning, ensemble methods, such as bagging, boosting, and stacking [15] aim to improve the learning performance given multiple algorithms (e.g., for anomaly detection [16]). They aggregate the outputs of learners in dynamic ways, such as voting on them or averaging them. This article focuses on prespecified logical/temporal relations between the outputs of black-box detection algorithms.

Logically specified properties are often used for safety monitoring and runtime verification in CPS. For STL, robustness [17], [18] takes a worst-case view on error, measuring the maximum deviation to violate the property. We focus on the average-case (probabilistic) error measures, such as false-positive and false-negative rates [6]. Typically estimated directly from samples, such measures do not necessarily transfer to the deployment environment. We provide a more rigorous way to estimate error rates based on probabilistic computations.

Monitoring of probabilistic logical properties addresses a different problem, but uses relevant techniques: it relies on the information from the observed data to test whether a probability constraint holds. Approaches, such as ProMo [19], PTPSC [20], and BaProMon [21] perform online checks of probability constraints via statistical hypothesis tests. Similarly, statistical approaches can estimate monitor error rates, although at design time: each rate is modeled as a

Bernoulli random variable, with its parameter estimated from monitor traces. As we show in Section VI, the downside of this approach is the need to observe and label potentially rare events for each monitor under analysis, necessitating extensive data collection. In contrast, our approach uses atomic, more observable inputs.

### B. Probability, Logic, and Inference

We distinguish two groups of combinations of formal logics and probability theory [22]: those with probabilistic entailment (i.e., logics of probability [23]), and those with probability operators (i.e., probabilistic logics [24]). The former type of logics (e.g., Adams' logic [25]) uses the classic logical syntax and performs *uncertainty propagation*, from the uncertainties of premises to the uncertainty of the conclusion. These logics address a problem different from ours: they manipulate uncertain knowledge about certain statements (e.g., “the system is safe” has a 95% chance of being true), whereas this article manipulates certain statements about uncertain phenomena (e.g., “the chance of false positive is 5%” is true).

The latter type of logics use probability operators over events; for example, the PCTL [5] can describe probabilistic transition systems. A typical use of such logics is to model-check the desired property on a known model or develop a set of premises that would prove or disprove the property deductively. Similarly to these logics, we reason about probabilities of related events, but do not use explicit probability operators in our syntax. Furthermore, our approach also does not require models of systems dynamics or trustworthy premises that logically entail the desired statement.

A notable recent probabilistic logic is the chance-constrained temporal Logic (C2TL) [26], which turns constraints on probabilities of deterministic predicates into optimization problems for the controller synthesis. Similar to our approach, C2TL probabilistically handles perception uncertainty of logically related events. The key difference is that we solve a probability estimation—not control—problem and thus rely on error rates of atomic detectors instead of chance constraints and a plant model.

Logic-free probabilistic inference often uses graphical models, such as the Bayesian nets and the Markov chains [13]. Usually approximate, this inference computes a probability query given a network fitted to the data. Instead of fitting a model, we discharge part of the problem by performing exact symbolic transformations, which allows us to handle arbitrary logical relationships. In the future, we hope to use graphical models to perform (in)dependence calculations for our approach.

## VIII. DISCUSSION

The reader may notice that the values of FPR and FNR described in the evaluation are unacceptably high for a safety-critical monitor: most of the time, the  $M_{pr}$  would not detect that the UAV has lost the pipeline. This finding came unexpectedly to the authors, and it illustrates the benefit of applying our approach in early design stages. Two responses would reduce the monitor's error rates: 1) significant improvements

to the pipe detector (i.e., reducing its FPR and FNR below 0.1) or 2) redefining the properties (e.g., the loss of pipeline occurs if it was not visible for 2 out of 3 s). For instance, if the pipe detector is improved such that  $\text{fpr}(PI) = 0.01$ , then  $\text{fnr}(M_{pr})$  is lowered from 0.651 to 0.096 for  $d = 10$ . Similarly, setting  $\text{fnr}(PI) = 0.01$  reduces  $\text{fpr}(M_{pr})$  from 0.686 to 0.105 for  $d = 10$ . This reduction indicates that our estimation can inform impactful design choices.

We note a dual relationship between the modalities in detector formulas and the FNR/FPR computations. The detectors with the box modality lead to straightforward computations of FNR (given standard independence assumptions) and relatively large values for it, whereas the computations for FPR rely on complex formulas and produce small values. For diamond modalities, the situation is reversed: simple computations for large FPRs, and complex computations for small FNRs. Formally, this is due to the conditioning in the error rate definition—either on conjunction (simple) or a disjunction (complex).

#### A. Future Work

Several important research directions are opened by this article. One such direction is enhancing the framework outputs. First, one could improve accuracy for limited data by moving from point estimates to interval estimates with uncertainty. Also, the framework could explore the design space by optimizing error rates over parameters in probability formulas.

Another direction is transitioning the analysis to run time. One could continuously maintain the estimates of the analysis inputs and update the error rate estimates as needed. Monitoring independence assumptions between detectors would be particularly useful to compensate for design-time simplifications. Furthermore, our analysis could predict the probability of future events (e.g., crashes) that are logically connected to the observed events.

Finally, one can envision a CPS design environment that iterates through sets of independence assumptions that would be sufficient to calculate the desired rate given the known inputs. This environment would suggest alternative independence setups to the engineers along with the additional required inputs (and their uncertainties), making the design choices more transparent and automated.

#### B. Conclusion

This article proposed a modeling and analysis framework for logical compositions of detectors. The modeling part represented atomic detectors as compact probability spaces and composed them with 3-value temporal-logic formulas. The analysis part estimated the error rates of composite detectors by manipulating symbolic probability expressions, via independence inference and estimation from data. Our evaluation on a pipeline detection case study showed that our framework's estimates were accurate within fractions of a percent and comparable to purely data-based (noncompositional) estimates on monitor traces.

#### ACKNOWLEDGMENT

The authors thank Taylor Carpenter and Hyonyoung Choi for their help with the UUV simulator. They also thank the anonymous reviewers for their insightful comments.

#### REFERENCES

- [1] K. Havelund and G. Reger, "Runtime verification logics a language design perspective," in *Models, Algorithms, Logics and Tools: Essays Dedicated to Kim Guldstrand Larsen on the Occasion of His 60th Birthday*. Cham, Switzerland: Springer, 2017, pp. 310–338.
- [2] E. Bartocci *et al.*, *Specification-Based Monitoring of Cyber-Physical Systems: A Survey on Theory, Tools and Applications*. Cham, Switzerland: Springer, Jan. 2018, pp. 135–175.
- [3] A. Pnueli, "The temporal logic of programs," in *Proc. 18th Annu. Symp. Found. Comput. Sci.*, Oct. 1977, pp. 46–57.
- [4] C. Szegedy *et al.*, "Intriguing properties of neural networks," in *Proc. Int. Conf. Learn. Represent.*, 2014, pp. 1–6.
- [5] H. Hansson and B. Jonsson, "A logic for reasoning about time and reliability," *Formal Aspects Comput.*, vol. 6, pp. 512–535, Sep. 1994.
- [6] J. L. Fleiss, B. Levin, and M. C. Paik, *Statistical Methods for Rates & Proportions*, 3rd ed. Hoboken, NJ, USA: Wiley Intersci., 2003.
- [7] H. Cramer, *Mathematical Methods of Statistics*. Princeton, NJ, USA: Princeton Univ. Press, 1999.
- [8] A. Bauer, M. Leucker, and C. Schallhart, "Monitoring of real-time properties," in *FSTTCS*. Berlin, Germany: Springer, 2006.
- [9] B. Konikowska, "A three-valued linear temporal logic for reasoning about concurrency," ICS PAC, Warsaw, Poland, Rep., Nov. 1998.
- [10] I. Ruchkin, O. Sokolsky, J. Weimer, T. Hedao, and I. Lee, "Supplementary materials for compositional probabilistic analysis of temporal properties over stochastic detectors," Univ. Pennsylvania, Dept. Comput. Inf. Sci., Rep., Jul. 2020, doi: [10.13140/RG.2.2.27777.07522](https://doi.org/10.13140/RG.2.2.27777.07522).
- [11] S. C. Kleene, *Introduction to Metamathematics*. Amsterdam, The Netherlands: North-Holland, 1952.
- [12] P. J. Brockwell, *Time Series: Theory and Methods*. New York, NY, USA: Springer, 2009.
- [13] D. Koller, N. Friedman, and F. Bach, *Probabilistic Graphical Models: Principles and Techniques*. Cambridge, MA, USA: MIT Press, Jul. 2009.
- [14] Y. Bar-Shalom, P. K. Willett, and X. Tian, *Tracking and Data Fusion: A Handbook of Algorithms*. Mansfield, CT, USA: YBS, Apr. 2011.
- [15] O. Sagi and L. Rokach, "Ensemble learning: A survey," *WIREs Data Min. Knowl. Discov.*, vol. 8, no. 4, 2018, Art. no. e1249.
- [16] C. C. Aggarwal, *Outlier Analysis*. New York, NY, USA: Springer, Jan. 2013.
- [17] J. V. Deshmukh, A. Donze, S. Ghosh, X. Jin, G. Juniwal, and S. A. Seshia, "Robust online monitoring of signal temporal logic," *Formal Methods Syst. Des.*, vol. 51, no. 1, pp. 5–30, Aug. 2017.
- [18] H. Abbas, Y. V. Pant, and R. Mangharam, "Temporal logic robustness for general signal classes," in *Proc. Int. Conf. Hybrid Syst. Comput. Control*, Apr. 2019, pp. 45–56.
- [19] L. Grunske and P. Zhang, "Monitoring probabilistic properties," in *Proc. ESEC/FSE*, 2009, pp. 183–192.
- [20] P. Zhang, W. Li, D. Wan, and L. Grunske, "Monitoring of probabilistic timed property sequence charts," *Softw. Pract. Exp.*, vol. 41, no. 7, pp. 841–866, 2011.
- [21] Y. Zhu, M. Xu, P. Zhang, W. Li, and H. Leung, "Bayesian probabilistic monitor: A new and efficient probabilistic monitoring approach based on Bayesian statistics," in *Proc. 13th Int. Conf. Qual. Softw.*, Jul. 2013, pp. 45–54.
- [22] R. Haenni, J.-W. Romeijn, G. Wheeler, and J. Williamson, *Probabilistic Logics and Probabilistic Networks*. Dorchester, The Netherlands: Springer, 2011.
- [23] A. Sernadas, J. Rasga, and C. Sernadas, "On probability and logic," *Portugaliae Mathematica*, vol. 74, no. 4, pp. 267–313, Feb. 2018.
- [24] G. De Bona, F. G. Cozman, and M. Finger, "Towards classifying propositional probabilistic logics," *J. Appl. Logic*, vol. 12, no. 3, pp. 349–368, 2014.
- [25] E. Adams, *A Primer of Probability Logic*. Stanford, CA, USA: CSLI, 1998.
- [26] S. Jha, V. Raman, D. Sadigh, and S. A. Seshia, "Safe autonomy under perception uncertainty using chance-constrained temporal logic," *J. Autom. Reason.*, vol. 60, no. 1, pp. 43–62, Jan. 2018.