

Causality-Aided Falsification

Takumi Akazaki *
The University of Tokyo, Japan
JSPS Research Fellow
akazaki@ms.k.u-tokyo.ac.jp

Yoshihiro Kumazawa
The University of Tokyo, Japan
kumazawa@is.s.u-tokyo.ac.jp

Ichiro Hasuo †
National Institute of Informatics, Tokyo, Japan
hasuo@nii.ac.jp

Falsification is drawing attention in quality assurance of heterogeneous systems whose complexities are beyond most verification techniques' scalability. In this paper we introduce the idea of *causality aid* in falsification: by providing a falsification solver—that relies on stochastic optimization of a certain cost function—with suitable causal information expressed by a *Bayesian network*, search for a falsifying input value can be efficient. Our experiment results show the idea's viability.

1 Introduction

Falsification In computer science, *verification* refers to the task of giving a mathematical proof to the claim that the behavior of a system \mathcal{M} satisfies a desired property φ (called a *specification*), under any circumstances (such as choices of input to the system \mathcal{M}). A mathematical proof thus obtained gives a level of confidence that is fundamentally different from empirical guarantees given by *testing*.

Extensive research efforts have yielded a number of effective verification techniques and they have seen successful real-world applications. At the same time, however, it is also recognized that *large-scale heterogeneous systems* are still beyond the scalability of most of these verification techniques. Notable among such are *cyber-physical systems (CPSs)* that exhibit not only discrete digital dynamics but also continuous physical dynamics. Imagine a car today: it contains not only dozens of chips (like ECUs) but also continuous dynamics (wheels, suspensions, internal combustion, etc.).

It is in this CPS context that the idea of *falsification* is found its use [17].

The falsification problem

- **Given:** a model \mathcal{M} (a function from an input signal to an output signal), and a *specification* φ (a temporal formula)
- **Answer:** a *critical path*, that is, an input signal σ_{in} such that the corresponding output $\mathcal{M}(\sigma_{\text{in}})$ does not satisfy φ

Two benefits of falsification are particularly appealing. For one, a system model \mathcal{M} can be totally a *black box*: once \mathcal{M} as a function $\sigma_{\text{in}} \mapsto \mathcal{M}(\sigma_{\text{in}})$ is given as an oracle, we can check if an educated guess σ_{in} is a solution or not—without knowing \mathcal{M} 's internal working. This is an advantage given that many CPSs do have black-box components: they can come from external suppliers, or they can be physical dynamics too complex to mathematically model (typically one uses *look-up tables* to describe such).

*Supported by Grants-in-Aid for JSPS Fellows No. 15J09877.

†Supported by JST ERATO HASUO Metamathematics for Systems Design Project (No. JPMJER1603), and JSPS Grant-in-Aid No. 15KT0012.

Algorithm 1 Falsification by optimization, with a cost function f_φ

Input: σ_0 ▷ The initial guess
1: $v_0 := f_\varphi(\mathcal{M}(\sigma_0));$ ▷ $v_i = f_\varphi(\mathcal{M}(\sigma_i))$ is the score of the input v_i
2: **for** $i = 1 \dots N$ **do** ▷ N is the greatest number of iteration
3: $\sigma_i := \operatorname{argmin}_\sigma \left(f_\varphi(\mathcal{M}(\sigma)) \mid \begin{array}{l} \text{under the previous observations} \\ (\sigma_0, v_0), (\sigma_1, v_1), \dots, (\sigma_{i-1}, v_{i-1}) \end{array} \right);$
4: $v_i := f_\varphi(\mathcal{M}(\sigma_i));$
5: **if** $v_i < 0$ **then return** $\sigma_i;$
6: **end if** ▷ Falsification succeeded, because we assume $f_\varphi(\tau) < 0$ implies $\tau \not\models \varphi$
7: $i := i + 1;$
8: **end for**

Another appealing feature of falsification is its affinity with *machine learning (ML)* and *optimization* techniques. In automatic verification techniques the greatest challenge is *state-space explosion*: the size of the input space for σ_{in} grows exponentially with respect to its dimension, often to the extent that exhaustive search in it is no longer possible. Recent surges in ML and optimization algorithms can offer potent countermeasures against this *curse of dimensionality*: specifically, after observing output $\mathcal{M}(\sigma_1), \dots, \mathcal{M}(\sigma_n)$ for input $\sigma_1, \dots, \sigma_n$, those algorithms can “learn” from these previous attempts and suggest an input signal σ_{n+1} with which $\mathcal{M}(\sigma_{n+1}) \not\models \varphi$ is likely.

One can say that falsification is after all *adaptive testing*: most falsification solvers rely on stochastic guess; hence their failure do not prove “ $\mathcal{M}(\sigma_{\text{in}}) \models \varphi$ for every σ_{in} .” However in many real-world scenarios falsification is as good as one gets, because of systems’ complexity and their black-box components within. Existing stochastic optimization-based solvers (such as S-TaLiRo [8] and BREACH [13]) have shown striking performance, too, scaling up to various Simulink diagrams from automotive applications. Moreover, falsification has special appeal to real-world engineers: while it takes certain familiarity to come to appreciate correctness proofs, counterexamples discovered by falsification easily convince engineers that there are issues to be resolved.

Search of Cost Functions A technical cornerstone that set off the study of falsification is *robust semantics* of temporal formulas [14, 15]. With CPS application in mind we assume that input and output of our system model \mathcal{M} are given by (time-variant) signals. For them it is standard to specify properties using some *temporal logic*, such as metric interval temporal logic (MITL) [7] and signal temporal logic (STL) [20]. In robust semantics [14, 15] a signal σ and a formula φ are assigned a *continuous* truth value $\llbracket \sigma, \varphi \rrbracket \in \mathbb{R}$ that designates how robustly the formula is satisfied. This departure from the conventional Boolean semantics (where $\llbracket \sigma, \varphi \rrbracket \in \{\text{t}, \text{ff}\}$) allows one to adopt a hill climbing-style optimization algorithm to look for a falsifying input signal.

Algorithm 1 is a high-level description of falsification by optimization. Here a *cost function* f_φ carries a signal (output of the system \mathcal{M}) to a real; we assume that its value is linked with satisfaction of φ , that is specifically, $f_\varphi(\tau) < 0$ implies $\tau \not\models \varphi$. We assume that the value of f_φ for a given input can be effectively computed; we assume the same for the function \mathcal{M} . Still in Line 3 the true solution may not be available since the global structure of \mathcal{M} is unknown—this reflects our black-box view on \mathcal{M} . Therefore in Line 3 we make a guess based on the previous trials.

The robust semantics of temporal formulas in [14, 15] is a prototype of such a cost function (Algorithm 1). Subsequently in the study of falsification, search of better cost functions has been an important topic. For example, sometimes *time robustness* [14]—as opposed to *space robustness* in the original work [15]—yields smoother hills to climb down, aiding optimization. Combination of space and time robustness is pursued in [5], where they enrich logics with *averaged* modalities to systematically enhance expressivity. Additional bias is put on cost functions in [12] so that search for falsifying input

covers a greater number of discrete modes of a system \mathcal{M} . After all, the question here is how to enrich cost functions, extracting additional information from a system \mathcal{M} and/or a specification φ .

Contribution: Causality Aid in Falsification In this paper we build on the observations in [4] and propose to aid falsification using *causal information*. We lay out the idea using a simple example.

Example 1 (incremental counter) Consider the pseudocode shown on the right. We think of: $i_0, i_1, \dots, i_N \in [-1, 1]$ as the values of a time-variant input signal i at time $t = 0, 1, \dots, N$, respectively; Lines 1–10 as a system \mathcal{M} that takes such input and returns the value of `cnt` as output; and the assertion $\text{cnt} \leq N$ as a specification φ , that is $\square_{[N,N]}(\text{cnt} \leq N)$ in temporal logic. It is clear that, to falsify φ , all the input values i_0, i_1, \dots, i_N must lie in $(-0.2, 0.2)$; otherwise the counter is reset and never reaches $\text{cnt} = N + 1$.

Input: $i_0, i_1, \dots, i_N \in [-1, 1]$
Output: `cnt`
1: $t := 0; \text{cnt} := 0$
2: **while** $t \leq N$ **do**
3: `flag := (|i| < 0.2);`
4: **if** `flag` **then**
5: `cnt := cnt + 1`
6: **else**
7: `cnt := 0;`
8: **end if**
9: $t := t + 1;$
10: **end while**
11: **assert**($\text{cnt} \leq N$)

Now consider solving the falsification problem here. It turns out that existing falsification solvers have hard time in doing so: besides the apparent hardness of the problem (following the uniform distributions the success probability would be 0.2^N), there is the following “causal” reason for the difficulty.

Assume $i_N \notin (-0.2, 0.2)$, meaning that `cnt` is reset to 0 at the last moment. In this case the earlier input values i_0, i_1, \dots, i_{N-1} have no effect in the final output `cnt`, nor in the robust semantics of the specification $\square_{[N,N]}(\text{cnt} \leq N)$. Therefore there is no incentive for stochastic optimization solvers to choose values i_0, \dots, i_{N-1} from $(-0.2, 0.2)$. More generally, desired bias is imposed on earlier input values i_0, \dots, i_k *only after* later input values i_{k+1}, \dots, i_N have been suitably fixed. Given the system (the above program) as a black box and the specification $\square_{[N,N]}(\text{cnt} \leq N)$ alone, there is no way for optimization solvers to know such causal dependency.

Our enhancement of falsification algorithms consists of leveraging causal information expressed as *Bayesian networks*. See Fig. 1, where we fix $N = 5$ for presentation. The Bayesian network expresses causal dependence of the original specification $\varphi = \varphi_5$ on other specifications $\varphi_0, \dots, \varphi_4$. The newly introduced specifications $\varphi_i = \square_{[i,i]}(\text{cnt} \leq i)$, for each $i = 0, \dots, 4$, express that the counter `cnt` has already been reset by time i . Therefore in order to falsify φ_5 , i.e. to keep incrementing `cnt`, these additional specifications must be falsified, too. The last observation is expressed in the Bayesian network in Fig. 1, specifically in the conditional probabilities $\Pr(\llbracket \varphi_{i+1} \rrbracket = \text{ff} \mid \llbracket \varphi_i \rrbracket = \text{tt}) = 0$.

Now our falsification algorithm looks not only at φ_5 but also at the other predicates $\varphi_0, \dots, \varphi_4$. This way we successfully impose bias on earlier input values i_0, \dots, i_4 to lie in $(-0.2, 0.2)$ —as demonstrated by our experimental results later.

Following the idea illustrated in the last example, our main contribution in this paper is a *causality-aided* falsification algorithm that uses *Bayesian networks of temporal formulas* as input on the specification side. Such a Bayesian network can be derived from an original specification φ alone; they can also be derived through inspection of a system model \mathcal{M} ; or we can use both φ and \mathcal{M} . In order to efficiently leverage the causal information expressed by a Bayesian network, we follow [9, 10] and use variations of *Gaussian process optimization* as our optimization algorithms (Line 3 of Algorithm 1). The feature that they allow to guess both average and variance (see §2.2) turns out to be particularly useful. We imple-

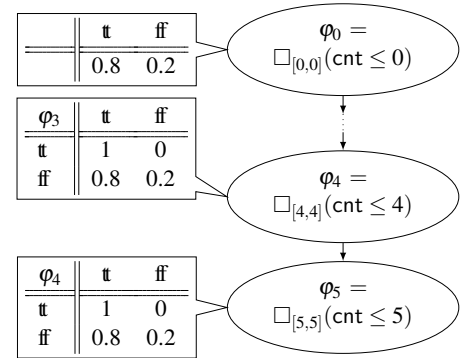


Figure 1: Causality annotation for incremental counter (where $N = 5$)

mented the algorithm; our experimental results, although preliminary, seem to support the effectiveness of our approach.

General methodologies of deriving such Bayesian networks are outside the paper’s focus, although we do have some preliminary ideas and we exploited them for our current examples. One is the use of *probabilistic predicate transformers* that are a classic topic in semantics [18, 19, 22] and are shed fresh light on in the context of *probabilistic programming languages* (see e.g. [23]). This idea follows the earlier observations in [6]; it successfully generates the Bayesian network in Fig. 1 for Example 1. Another idea is parse tree-like decomposition of an original temporal formula φ ; we decorate the resulting tree with conditional probabilities that we learn through sampling. These methods will be described in our forthcoming papers.

Related Work Besides search of better cost functions, an important direction in the study of falsification is improving optimization algorithms (that are used in Line 3 of Algorithm 1). In the falsification literature many different algorithms have been used and studied: they include *simulated annealing*, *ant-colony optimization*, the *cross-entropy method*, the *Nelder-Mead algorithm*, and so on [8, 13, 25]. In [11] a discrete algorithm of *Tabu search* is employed for enhanced coverage.

Yet another important direction is *multiple shooting falsification* [27, 28] where, unlike *single shooting* approaches like in this paper, a bunch of trajectories are investigated in a single iteration relying on suitable abstraction of a system model and/or a specification. We believe our idea of causality aid in falsification is orthogonal to the choice between single and multiple shooting; we will study as future work the effect of causality in multiple shooting falsification.

2 Backgrounds

2.1 STL and Robust Semantics

Here we present *signal temporal logic (STL)* [20] as our formalism for expressing (original, without causal information) specifications. We also present its *robust semantics* [14] that give the prototype of the cost function f_φ in Algorithm 1. Our cost function will be derived from the robust semantics of the formulas in a Bayesian network. At the same time we emphasize that our methodology of causality-aided falsification does not depend on the specific underlying specification formalism of **STL**.

Definition 2.1 (syntax of STL) The set of *STL formulas* are recursively defined as follows.

$$\varphi ::= g(\mathbf{y}) > 0 \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \mathcal{U}_I \varphi_2$$

Here $g(\mathbf{y})$ is some real-value function over the set of variables $\mathbf{y} = \{y_1, \dots, y_n\}$, and I is a closed non-singular interval in $\mathbb{R}_{\geq 0}$.

We also introduce the following standard temporal operators as abbreviations: the *eventually* operator $\diamond_I \varphi \triangleq (\infty > 0) \mathcal{U}_I \varphi$ and the *always* operator $\square_I \varphi \triangleq \neg \diamond_I \neg \varphi$.

Definition 2.2 (Boolean semantics of STL) Let $\sigma_{\mathbf{y}}: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^n$ be a *signal*, that is, a function that maps time τ to the values $\sigma_{\mathbf{y}}(\tau)$ of the variables \mathbf{y} at time τ . We define the (Boolean) validity of an **STL** formula

over a signal σ_y , as follows. Here σ_y^τ stands for the time-shifted signal such that $\sigma_y^\tau(\tau') \triangleq \sigma_y(\tau + \tau')$.

$$\begin{array}{lll}
\sigma_y \models g(\mathbf{y}) > 0 & \stackrel{\text{def.}}{\iff} & \text{the inequality } g(\sigma_y(0)) > 0 \text{ holds} \\
\sigma_y \models \neg\varphi & \stackrel{\text{def.}}{\iff} & \sigma_y \not\models \varphi \\
\sigma_y \models \varphi_1 \vee \varphi_2 & \stackrel{\text{def.}}{\iff} & \sigma_y \models \varphi_1 \text{ or } \sigma_y \models \varphi_2 \\
\sigma_y \models \varphi_1 \mathcal{U}_I \varphi_2 & \stackrel{\text{def.}}{\iff} & \exists \tau \in I. (\sigma_y^\tau \models \varphi_2 \text{ and } \forall \tau' \in [0, \tau]. \sigma_y^{\tau'} \models \varphi_1)
\end{array}$$

The following ‘‘quantitative refinement’’ of the semantics of **STL** initiated the research program of falsification by optimization [14, 15].

Definition 2.3 (robust semantics of STL) For a signal σ_y and an **STL** formula φ , we define the *robustness* $\llbracket \sigma_y, \varphi \rrbracket \in \mathbb{R} \cup \{\infty, -\infty\}$ inductively as follows. Here \sqcap and \sqcup denote infimums and supremums of real numbers, respectively.

$$\begin{array}{lll}
\llbracket \sigma_y, g(\mathbf{y}) > 0 \rrbracket & \triangleq & g(\sigma_y(0)) \\
\llbracket \sigma_y, \neg\varphi \rrbracket & \triangleq & -\llbracket \sigma_y, \varphi \rrbracket \\
\llbracket \sigma_y, \varphi_1 \vee \varphi_2 \rrbracket & \triangleq & \llbracket \sigma_y, \varphi_1 \rrbracket \sqcup \llbracket \sigma_y, \varphi_2 \rrbracket \\
\llbracket \sigma_y, \varphi_1 \mathcal{U}_I \varphi_2 \rrbracket & \triangleq & \sqcup_{\tau \in I} (\llbracket \sigma_y^\tau, \varphi_2 \rrbracket \sqcap \sqcap_{\tau' \in [0, \tau]} \llbracket \sigma_y^{\tau'}, \varphi_1 \rrbracket)
\end{array}$$

Note that the sign of robustness coincides with the Boolean semantics. That is, $\llbracket \sigma_y, \varphi \rrbracket > 0$ implies $\sigma_y \models \varphi$, and $\llbracket \sigma_y, \varphi \rrbracket < 0$ implies $\sigma_y \not\models \varphi$. Conversely, $\sigma_y \models \varphi$ implies $\llbracket \sigma_y, \varphi \rrbracket \geq 0$, and $\sigma_y \not\models \varphi$ implies $\llbracket \sigma_y, \varphi \rrbracket \leq 0$.

2.2 Gaussian Process Optimization

In this paper we follow the workflow in Algorithm 1, deriving the cost function f_φ in it from a Bayesian network. For the optimization step (Line 3 of Algorithm 1) we use *Gaussian process optimization*—we follow [4, 9, 10] about this choice. It has a feature that it suggests the global shape of an unknown function; this feature turns out to be convenient for our purpose of integrating causal information in falsification. We present a brief review of the topic; see e.g. [24] for details.




Figure 2: Gaussian process regression with the squared-exponential covariance kernel $k(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|^2 / 2l^2)$ with different length scale ($l = 1, 0.1, 0.01$). The black dots designate previous observations $D = \{(\mathbf{x}_1, f(\mathbf{x}_1)), \dots, (\mathbf{x}_t, f(\mathbf{x}_t))\}$. By the definition of Gaussian processes, when one looks at the vertical section of each figure at specific input $\mathbf{x} \in \mathbf{X}$, the shade of blue forms a normal distribution $\mathbf{GP}(\mu, k)(\mathbf{x}) = \mathcal{N}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}))$. The center of the pipe (the light-blue line in each figure) stands for the expected value $\mu(\mathbf{x})$ of the unknown function f ; the width of the pipe stands for its variance $k(\mathbf{x}, \mathbf{x})$.

2.2.1 Gaussian Process Regression

Let f be an unknown function, from a certain input domain to the set of real numbers, about which we wish to infer certain properties. (For Algorithm 1 we would take $f = f_\phi(\mathcal{M}(_))$). In *Gaussian process regression* the shape of f is estimated assuming that f is a probabilistic process called a *Gaussian process*.

We start with some formal definitions. For more detail, see e.g. [24].

Notation 2.4 We let $\mathcal{N}(\boldsymbol{\mu}, \mathbf{k})$ stand for the probability density function of the multivariate Gaussian distribution whose mean vector is $\boldsymbol{\mu}$ and covariance matrix is \mathbf{k} .

Definition 2.5 (Gaussian process) A *Gaussian process* is a family of probabilistic variables $(\mathbf{z}_{\mathbf{x}})_{\mathbf{x} \in \mathbf{X}}$ such that each of its finite subset $(\mathbf{z}_{\mathbf{x}_1}, \dots, \mathbf{z}_{\mathbf{x}_t})$ has a joint Gaussian distribution. A Gaussian process is determined by a pair $(\boldsymbol{\mu}, k)$ of its *mean function* $\boldsymbol{\mu} : \mathbf{X} \rightarrow \mathbb{R}$ and its *covariance function* $k : \mathbf{X} \times \mathbf{X} \rightarrow \mathbb{R}$; this Gaussian process is denoted by $\mathbf{GP}(\boldsymbol{\mu}, k)$. For this we have

$$(\mathbf{z}_{\mathbf{x}_1}, \dots, \mathbf{z}_{\mathbf{x}_t})^\top \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{k}) \quad \text{where} \quad \boldsymbol{\mu}_i = \boldsymbol{\mu}(\mathbf{x}_i) \quad \text{and} \quad \mathbf{k}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$$

for each finite subset $\{\mathbf{x}_1, \dots, \mathbf{x}_t\}$ of \mathbf{X} . We write $\mathbf{GP}(\boldsymbol{\mu}, k)(\mathbf{x}_1, \dots, \mathbf{x}_t)$ for the above multivariate Gaussian distribution $\mathcal{N}(\boldsymbol{\mu}, \mathbf{k})$.

In Fig. 2 is how an unknown function f can be guessed by Gaussian processes. The blue pipe designates the estimated values of the unknown function f : the farther input \mathbf{x} is from the observed points, the thicker the pipe is (that means bigger uncertainty).

In the regression of f using Gaussian processes, a choice of a covariance function $k : \mathbf{X} \times \mathbf{X} \rightarrow \mathbb{R}$ determines *smoothness* of f . A common template for covariance functions is the *squared-exponential kernel function* $k_l(\mathbf{x}, \mathbf{x}') \triangleq \exp(-l \cdot \|\mathbf{x} - \mathbf{x}'\|^2/2)$, where l is so-called the *length scale* parameter. In practice, we pick a good length scale parameter by cross validation. As we see in Fig. 2, the choice of a covariance function yields the following tendencies in Gaussian process regression:

- The bigger the distance $\|\mathbf{x} - \mathbf{x}'\|$ is, the smaller the covariance is, thus the harder it gets to estimate the value $f_\phi(\mathbf{x})$ from the observation of the value $f_\phi(\mathbf{x}')$.
- Covariance is smaller too when the length scale parameter l is bigger.

One advantage of Gaussian process regression is that, given a set of observations, the *posterior* process is described analytically. Let random variables $(\mathbf{z}_{\mathbf{x}})_{\mathbf{x} \in \mathbf{X}}$ obey a prior Gaussian process $\mathbf{GP}(\boldsymbol{\mu}, k)$; and $D = \{(\mathbf{x}_1, \mathbf{z}_1), \dots, (\mathbf{x}_t, \mathbf{z}_t)\}$ be a set of observations. Then the posterior distribution, denoted by $\mathbf{GP}(\boldsymbol{\mu}, k; D)$, is given by the Gaussian process $\mathbf{GP}(\boldsymbol{\mu}', k')$, where

$$\begin{aligned} \boldsymbol{\mu}'(\mathbf{x}) &= \boldsymbol{\mu}(\mathbf{x}) + \mathbf{k}_D(\mathbf{x})\mathbf{k}_{DD}^{-1}([\mathbf{z}_1 \dots \mathbf{z}_t]^\top - [\boldsymbol{\mu}(\mathbf{x}_1) \dots \boldsymbol{\mu}(\mathbf{x}_t)]^\top), \\ k'(\mathbf{x}, \mathbf{x}') &= k(\mathbf{x}, \mathbf{x}') - \mathbf{k}_D(\mathbf{x})\mathbf{k}_{DD}^{-1}\mathbf{k}_D(\mathbf{x}')^\top. \end{aligned}$$

Here $\mathbf{k}_D(\mathbf{x}) = [k(\mathbf{x}_1, \mathbf{x}) \dots k(\mathbf{x}_t, \mathbf{x})]$, and \mathbf{k}_{DD} is a $t \times t$ matrix whose i, j -component is $k(\mathbf{x}_i, \mathbf{x}_j)$. In practice, given observed data $D = \{(\mathbf{x}_1, f(\mathbf{x}_1)), \dots, (\mathbf{x}_t, f(\mathbf{x}_t))\}$ and a covariance kernel function k , we estimate the function f as $\mathbf{GP}(\mathbf{0}, k; D)$ where $\mathbf{0}$ denotes the function constantly zero.




Figure 3: Illustration of the **GP-PSat** algorithm. In each figure, the red line is for the unknown function f to minimize; and the blue cloud is the Gaussian process estimation $\mathbf{GP}(\mu, k)$ of f . At time $t = 1$ the input is chosen randomly (say $\mathbf{x}_1 = 5.5$). Subsequently we pick the point \mathbf{x}_t that minimizes the probability $\Pr_{\mathbf{GP}(\mu, k)}(f(\mathbf{x}) > 0)$. We observe that, as the algorithm proceeds, the estimate of f becomes finer too.

2.2.2 Gaussian Process Optimization and Acquisition Function

Gaussian process regression allows us to predict, based on observations in D , the value $f(\mathbf{x})$ for each input \mathbf{x} as a normal distribution $\mathbf{GP}(\mu, k)(\mathbf{x})$. To complete an optimization scenario, we wish to pick a candidate $\mathbf{x} \in \mathbf{X}$ for which $f(\mathbf{x})$ is small.

It is well-known that, for such choice, a balance is important between *exploration* (i.e. bias toward a bigger variance $k(\mathbf{x}, \mathbf{x})$) and *exploitation* (bias toward a smaller expected value $\mu(\mathbf{x})$). A criterion for this balance is called an *acquisition function*—we pick \mathbf{x} at which the acquisition function is minimum. Assuming that an acquisition function $\psi(\mathbf{x}; \mathbf{GP}(\mu, k))$ has been fixed, the whole procedure for Gaussian process optimization can be described as in Algorithm 2. Note that, in Line 3 of Algorithm 2, we usually employ another optimization solving method (such as simulated annealing).

In falsification, our goal would be to find \mathbf{x} such that $f(\mathbf{x}) < 0$. As a natural choice of acquisition functions, we focus on the following probability in this paper.

Definition 2.6 (Probability of Satisfaction)

$$\psi(\mathbf{x}; \mathbf{GP}(\mu, k)) \triangleq \Pr_{\mathbf{GP}(\mu, k)}(f(\mathbf{x}) > 0) \tag{1}$$

Here $\Pr_{\mathbf{GP}(\mu, k)}(f(\mathbf{x}) > c)$ is an abbreviation of $\Pr(f(\mathbf{x}) > c \mid f(\mathbf{x}) \sim \mathbf{GP}(\mu, k)(\mathbf{x}))$.

We write **GP-PSat** for Algorithm 2 under ψ as an acquisition function; Fig. 3 illustrates how it works.

The acquisition functions we will use are extension of this ψ . We note, however, that this acquisition function is commonly known as “pure and impractical” in the field of Gaussian process optimization. More sophisticated acquisition functions that are known include *probability improvement*, *expected improvement* [21], *upper confidence bound* [26] and so on. At the time of writing it is not clear how these acquisition functions can be used as part of our framework in §4.

Algorithm 2 Gaussian process optimization $\text{GPOptimization}(D, k, \psi)$ **Input:** a covariance function $k: \mathbf{X} \times \mathbf{X} \rightarrow \mathbb{R}$; an acquisition function ψ ; and an initial data set $D = \{(\mathbf{x}'_1, f(\mathbf{x}'_1)), \dots, (\mathbf{x}'_s, f(\mathbf{x}'_s))\}$ **Output:** input $\mathbf{x} \in \mathbf{X}$ for which $f(\mathbf{x})$ is small

```

1: for  $t = 1, 2, \dots$  do
2:    $\text{GP}(\mu', k') = \text{GP}(\mathbf{0}, k; D)$ ; ▷ Estimate the unknown function  $f$ 
3:    $\mathbf{x}_t = \arg \min_{\mathbf{x} \in \mathbf{X}} \psi(\mathbf{x}; \text{GP}(\mu', k'))$ ; ▷ Choose new sample input
4:    $D = D \cup \{(\mathbf{x}_t, f(\mathbf{x}_t))\}$ ; ▷ Observe the corresponding output
5: end for
6: return  $\mathbf{x}_t$ 

```

3 Causality in Falsification: Further Examples

In addition to Example 1, we shall exhibit two more examples of falsification problems; for each, we introduce a Bayesian network that encodes suitable causal information, too. The latter will be exploited in our causality-aided algorithm in §4.

3.1 Example Model 2: Coincidental Sine Waves

Let us consider the model in Fig. 4. In this simple model there are four sine waves $x_1(t), \dots, x_4(t)$ of different frequency, and we pick their initial phases i_1, \dots, i_4 as input of the system.

As a specification, we pick the following formula—it is falsified when the peaks of four sine waves correspond.

$$\varphi \equiv \square_{[0,10]} \left(\bigvee_{i=1,\dots,4} x_i < 0.99 \right) \quad (2)$$

We see that falsifying φ with pure random sampling is difficult because φ is false only in rare cases. What is worth, the (conventional) robustness of φ does not always guide us to the counterexamples.

Example 2 Let us consider the subformula $\bigvee_{i=1,\dots,4} x_i < 0.99$. When we compare the values $(x_1, x_2, x_3, x_4) = (1, 1, 1, 0)$ and $(0, 0, 0, 1)$, we could say the former is “closer” to falsifying the subformula—the three out of four sine waves simultaneously at a peak. However, these robustness values are the same 0.99 in both cases.

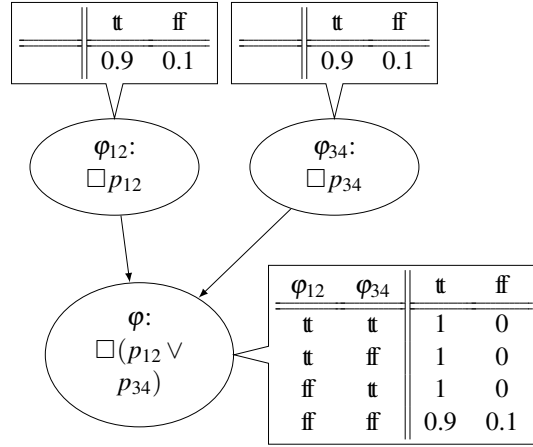
In this case, we sometimes divide the difficulty into small pieces—first get x_1 and x_2 simultaneously at a peak; then get x_3 and x_4 at a peak; finally, try to make them synchronize. Let us introduce formulas φ_{12} and φ_{34} such that falsifying them means matching the peak of x_1, x_2 , and x_3, x_4 respectively. Decomposing φ into φ_{12} and φ_{34} might help us in falsification for the following reasons.

- The small formulas φ_{12} and φ_{34} are much easier to falsify compared to φ .
- Moreover, the robustness mapping $f_{\varphi_{12}}(\mathcal{M}(_))$ and $f_{\varphi_{34}}(\mathcal{M}(_))$ have much simpler dynamics than the one of the original specification φ , so the Gaussian process regression for the small formulas tend to work better than the one for φ .

The Bayesian network \mathcal{B} in Fig. 5 is devised to express this intuition. For example, the formula φ is true with probability 1 when either φ_{12} or φ_{34} , otherwise φ becomes false with small probability 0.1. As shown in Fig. 6, the conditional joint distribution $\Pr_{\mathcal{B}}(- \mid \llbracket \varphi \rrbracket = \text{ff})$ tells us the fact that φ is false only if both φ_{12} and φ_{34} are false.

Input: $i_1, \dots, i_4 \in [0, 1]$
Output: $x_1(t), \dots, x_4(t)$
 for each $t \in \mathbb{R}_{\geq 0}$
 $x_1(t) = \sin(1.1t + i_1);$
 $x_2(t) = \sin(1.2t + i_2);$
 $x_3(t) = \sin(1.3t + i_3);$
 $x_4(t) = \sin(1.4t + i_4);$

Figure 4: System model for §3.1



where
 $\left(\begin{array}{l} p_{12} \equiv x_1 < 0.99 \vee x_2 < 0.99 \\ p_{34} \equiv x_3 < 0.99 \vee x_4 < 0.99 \end{array} \right)$

Figure 5: Bayesian network for §3.1

$\Pr_{\mathcal{B}}(-)$				$\Pr_{\mathcal{B}}(- \mid \llbracket \varphi \rrbracket = \text{ff})$			
φ_1	φ_2	φ		φ_1	φ_2	φ	
t	t	t	0.81	t	t	t	0
t	t	ff	0	t	t	ff	0
t	ff	t	0.09	t	ff	t	0
t	ff	ff	0	t	ff	ff	0
ff	t	t	0.09	ff	t	t	0
ff	t	ff	0	ff	t	ff	0
ff	ff	t	0.009	ff	ff	t	0
ff	ff	ff	0.001	ff	ff	ff	1

$\Pr_{\mathcal{B}}(-)$		$\Pr_{\mathcal{B}}(- \mid \llbracket \varphi \rrbracket = \text{ff})$	
$\varphi_{12} = \text{t}$	0.9	$\varphi_{12} = \text{t}$	0
$\varphi_{12} = \text{ff}$	0.1	$\varphi_{12} = \text{ff}$	1
$\varphi_{34} = \text{t}$	0.9	$\varphi_{34} = \text{t}$	0
$\varphi_{34} = \text{ff}$	0.1	$\varphi_{34} = \text{ff}$	1
$\varphi = \text{t}$	0.999	$\varphi = \text{t}$	0
$\varphi = \text{ff}$	0.001	$\varphi = \text{ff}$	1

Figure 6: Unconditional/conditional joint distributions in the Bayesian network of Fig. 5

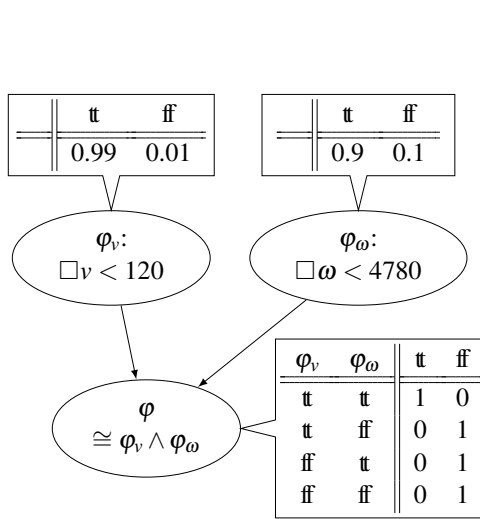


Figure 7: Bayesian network for §3.2

$\Pr_{\mathcal{B}}(-)$				$\Pr_{\mathcal{B}}(- \mid \llbracket \varphi \rrbracket = \text{ff})$			
φ_v	φ_ω	φ		φ_v	φ_ω	φ	
t	t	t	0.891	t	t	t	0
t	t	ff	0	t	t	ff	0
t	ff	t	0	t	ff	t	0
t	ff	ff	0.099	t	ff	ff	0.908
ff	t	t	0	ff	t	t	0
ff	t	ff	0.009	ff	t	ff	0.083
ff	ff	t	0	ff	ff	t	0
ff	ff	ff	0.001	ff	ff	ff	0.009

$\Pr_{\mathcal{B}}(-)$		$\Pr_{\mathcal{B}}(- \mid \llbracket \varphi \rrbracket = \text{ff})$	
$\varphi_v = \text{t}$	0.99	$\varphi_v = \text{t}$	0.908
$\varphi_v = \text{ff}$	0.01	$\varphi_v = \text{ff}$	0.092
$\varphi_\omega = \text{t}$	0.9	$\varphi_\omega = \text{t}$	0.083
$\varphi_\omega = \text{ff}$	0.1	$\varphi_\omega = \text{ff}$	0.917
$\varphi = \text{t}$	0.891	$\varphi = \text{t}$	0
$\varphi = \text{ff}$	0.109	$\varphi = \text{ff}$	1

Figure 8: Unconditional/conditional joint distributions in the Bayesian network of Fig. 7

3.2 Example Model 3: Automatic Transmission

The last example is the automatic transmission model from the benchmark of temporal logic verification [16]. This model is still miniature, but an elaborate mimicry of the systems in the real world hence suitable for our purpose.

As a specification φ to falsify, we use the following formula. It is taken from [16] (it is φ_2^{AT} there).

$$\Box(v < 120 \wedge \omega < 4780)$$

Here the variable v and ω stand for the speed of the vehicle and the angular velocity of the engine rotation respectively.

We know that we can falsify φ either by violating the speed limit ($v < 120$) or the engine rotation limit ($\omega < 4780$). In this model, ω takes the values in the range around $[0, 4800]$ while v does around $[0, 120]$. Note that their scales are very different: hence in the most of the cases, the robustness of the ω -component is likely to be shadowed by the one of the v -component. As a consequence, we expect that conventional falsification solver only try to falsify by the violation of the speed limit $v < 120$.

The Bayesian network annotation is also effective in such a situation. That is, we can add the information about “which is more likely to happen, the violation of the speed and the rotation limit.” (In actual deployment such insights will be provided by engineers’ domain knowledge.) Let assume that the probabilities of the violation of the speed and the rotation limit are 0.01 and 0.1 respectively. This information is expressed in Fig. 7, where the conditional probabilities for φ simply encode logical relationship (note that φ is semantically equivalent to $\varphi_v \wedge \varphi_\omega$) however, the probabilities at leaves reflect the above insight.

Remark 3.1 In §3.1 and §3.2, as an indicator of robustness, we employed the (space) robust semantics of **STL** in [14] and shown that it is not sensitive enough for some falsification scenarios. In contrast to [14], the metric-based robustness of **MITL** in [15] has a degree of freedom to capture the lacked notions. For example in §3.2, we could solve the falsification problem more efficiently if we could re-scale v and ω appropriately, and this re-scaling is nothing but the defining the metric space in [15]. However, defining such a metric space itself is challenging and needs expert’s domain knowledge—similarly as our framework needs suitable causal information. We expect that our causality-aided framework is a viable option compare to finding a suitable metric.

4 Falsification with Causality Annotation

Given the backgrounds in §2 and the examples in §3, we are now ready to ask the question: *given a falsification problem and a Bayesian network annotation about causality, what cost function should we optimize?* In this section, we will give some answers to the question by lifting up the conventional notion of acquisition functions which we reviewed in §2.2 to the multi-formula setting.

Consider one of the Bayesian networks that we have seen in the paper. Let \mathcal{B} denote the Bayesian network; and let $\Phi = \{\varphi_1, \dots, \varphi_N\}$ be the set of formulas that appear there. Now assume that we are running the Gaussian regression not only for $f_\varphi = \llbracket \mathcal{M}(_), \varphi \rrbracket$ but also $f_{\varphi_i} = \llbracket \mathcal{M}(_), \varphi_i \rrbracket$ for all the formulas φ_i in the Bayesian network.

The regression result for f_{φ_i} gives us the probabilistic “forecast” of the truth values assignment of the formulas $\Theta \in 2^\Phi$ as follows.

Notation 4.1 Let $\mathbf{GP}(\mu_i, k_i) \sim f_{\varphi_i}$ be our estimate for f_{φ_i} ; we can use this data to estimate the probability of obtaining Θ as the truth assignment, under an input value \mathbf{x} . Precisely: let Θ be the assignment $(\varphi_1 = \theta_1, \dots, \varphi_N = \theta_N)$ where $\theta_i \in \{\mathbf{t}, \mathbf{ff}\}$; then

$$\Pr_{\mathbf{GP}(\mathbf{x})}(\Theta) \triangleq \Pr_{\mathbf{GP}(\mu_1, k_1)}(f_{\varphi_1}(\mathbf{x}) R_1 0) \cdots \Pr_{\mathbf{GP}(\mu_N, k_N)}(f_{\varphi_N}(\mathbf{x}) R_N 0), \quad (3)$$

where R_i is $>$ if $\theta_i = \mathbf{t}$, and $<$ otherwise.

4.1 KL Divergence based acquisition function

Recall the scenario in §3.1—from the conditional joint distribution $\Pr_{\mathcal{B}}(- \mid \llbracket \varphi \rrbracket = \mathbf{ff})$, we see that the both small formulas φ_{12} and φ_{34} also should be false to synchronize all the peaks of the sine waves.

Inspired from the above example, we propose the following criteria to choose the next candidate \mathbf{x} as falsifying input.

Definition 4.2 (An acquisition function $\psi_{\mathcal{B}}(\mathbf{x})$)

$$\mathbf{x} = \arg \min_{\mathbf{x}} \psi_{\mathcal{B}}(\mathbf{x}) \text{ where } \psi_{\mathcal{B}}(\mathbf{x}) = D_{\text{KL}} \left(\Pr_{\mathcal{B}}(\Theta \mid \llbracket \varphi \rrbracket = \mathbf{ff}) \parallel \Pr_{\mathbf{GP}(\mathbf{x})}(\Theta) \right)$$

Here D_{KL} is the *Kullback Leibler divergence*—a measure of the difference between two probabilistic distributions. Intuitively, with this criteria, we pick the next input \mathbf{x} with which the probabilistic forecast $\Pr_{\mathbf{GP}(\mathbf{x})}(\Theta)$ by regression becomes “closer to the conditional joint distribution $\Pr_{\mathcal{B}}(\Theta \mid \llbracket \varphi \rrbracket = \mathbf{ff})$ ”.

Example 3 Let us consider the sine waves model in §3.1. From simple calculation, we see that the acquisition function $\psi_{\mathcal{B}}$ is as follows.

$$\psi_{\mathcal{B}}(\mathbf{x}) = -\log \Pr_{\mathbf{GP}}(f_{\varphi}(\mathbf{x}) < 0) - \log \Pr_{\mathbf{GP}}(f_{\varphi_{12}}(\mathbf{x}) < 0) - \log \Pr_{\mathbf{GP}}(f_{\varphi_{34}}(\mathbf{x}) < 0)$$

Hence minimizing $\psi_{\mathcal{B}}(\mathbf{x})$ means trying to falsify all the formulas φ , φ_{12} , and φ_{34} .

Remark 4.3 In this paper, we assume that $\Pr_{\mathcal{B}}(\llbracket \varphi \rrbracket = \mathbf{ff})$ is not 0 nor 1 on the given Bayesian network \mathcal{B} . In the former case, $\psi_{\mathcal{B}}(\mathbf{x})$ is undefined because $\Pr_{\mathcal{B}}(\llbracket \varphi \rrbracket = \mathbf{ff})$ is 0, and the latter case, $\psi_{\mathcal{B}}(\mathbf{x})$ is constantly 0 because $\Pr_{\mathcal{B}}(\Theta \mid \llbracket \varphi \rrbracket = \mathbf{ff}) = \Pr_{\mathcal{B}}(\Theta)$. We believe this is reasonable if we believe the given annotation \mathcal{B} is correct—in case $\Pr_{\mathcal{B}}(\llbracket \varphi \rrbracket = \mathbf{ff})$ is 0 (or 1) falsification never succeeds (or always succeeds, respectively).

The resulting extension of the **GP-PSat** algorithm (§2.2) with Bayesian networks is presented in Algorithm 3.

4.2 Another acquisition function based on the difference of KL divergence

Aside from the acquisition function $\psi_{\mathcal{B}}$ in Def. 4.2, we propose another criteria.

Definition 4.4 (Another acquisition function $\psi'_{\mathcal{B}}(\mathbf{x})$)

$$\mathbf{x} = \arg \min_{\mathbf{x}} \psi'_{\mathcal{B}}(\mathbf{x}) \text{ where}$$

$$\psi'_{\mathcal{B}}(\mathbf{x}) = D_{\text{KL}} \left(\Pr_{\mathcal{B}}(\Theta \mid \llbracket \varphi \rrbracket = \mathbf{ff}) \parallel \Pr_{\mathbf{GP}(\mathbf{x})}(\Theta) \right) - D_{\text{KL}} \left(\Pr_{\mathcal{B}}(\Theta) \parallel \Pr_{\mathbf{GP}(\mathbf{x})}(\Theta) \right).$$

Gaussian process regression Line 2 in Algorithm 3 is done by GPML MATLAB Code version 4.0 [3], a widely used library for computation about Gaussian processes.

Inference on Bayesian networks We employ Bayes Net Toolbox for Matlab [1] for inference on Bayesian networks.

The algorithms GP-PSat and GP-PI aided by Bayesian networks This part is new. Optimization of an acquisition function ψ is done by the following two steps: 1) we randomly pick initial samples $\mathbf{x}_1, \dots, \mathbf{x}_{100}$ and compute the corresponding values of ψ ; and 2) from the minimum \mathbf{x}_i of the one hundred, we further do greedy hill-climbing search.

5.2 Experiments

Using our implementation we conducted the following experiments. We do experiments for the three falsification problems; Problem 1 is from Examples 1, Problem 2 from §3.1 and Problem 3 from §3.2. For the automatic transmission example (in §3.2) we used two different parameters; Problem 3-1 is with the specification $\varphi = \square(v > -1 \wedge \omega < 4780)$; and Problem 3-2 is with $\varphi = \square(v < 120 \wedge \omega < 4780)$.

The experiments were done on a ThinkPad T530 with Intel Core i7-3520M 2.90GHz CPU with 3.7GB memory. The OS was Ubuntu14.04 LTS (64-bit). A single falsification trial consists of a number of iterations—iterations of for-loop in line 2 in Algorithm 1—before it succeeds or times out (after 100 seconds). For each problem we made ten falsification trials. We made multiple trials because of the stochastic nature of the optimization algorithm. We measured the performance by the following criteria:

- **Success rate:** The number of successful trials (out of ten).
- **The number of iteration loops:** The average number of iteration loops to find the counterexample.
- **The computational time:** The average time to find the counterexample.

Besides our two extended algorithms with the acquisition functions (in Def. 4.2 and 4.4), we measured the performance of the conventional Gaussian process optimization algorithms **GP-PSat** and compare them.

The experimental results are in Table 1. We see that our causality-aided approach (**GP-PSat** with $\psi_{\mathcal{B}}$ and $\psi'_{\mathcal{B}}$) significantly outperformed others for Example 1. This suggests promising potential of the proposed approach in the context of probabilistic programs—all the more because Bayesian networks like in Fig. 1 could be systematically derived using probabilistic predicate transformers.

Our algorithms performed at least as well as the conventional **GP-PSat**, for the other examples (Problem 2, 3-1 and 3-2). In Problem 3-1 and 3-2 we observe that our algorithms took fewer iterations before successful falsification. This is potentially an advantage when we wish to deal with bigger Simulink models as system models \mathcal{M} (their numerical simulation, i.e. computation of $\mathcal{M}(\sigma)$, is computationally expensive). That said, we believe the idea of causality aid in falsification can be a breaking one, with a potential of accelerating falsification by magnitudes. Its current performance for Problem 3 (that is from cyber-physical systems, a main application domain of falsification) is therefore not satisfactory. We will therefore pursue further improvement of our algorithm (Algorithm 3).

6 Future Work

In this paper, we show that the causality information given in the form of a Bayesian network helps us to solve falsification problems efficiently. However, we still have many challenges in constructing

Table 1: Experiment results

Algorithm	Problem 1			Problem 2			Problem 3-1			Problem 3-2		
	Succ.	Iter. (Succ.)	Time (Succ.)	Succ.	Iter. (Succ.)	Time (Succ.)	Succ.	Iter. (Succ.)	Time (Succ.)	Succ.	Iter. (Succ.)	Time (Succ.)
GP-PSat	0	208.2	100.0	5	122.4	77.2	10	49.5	41.5	10	23.0	6.7
		-	-		93.6	54.4		49.5	41.5		23.0	6.7
GP-PSat with ψ_{\emptyset}	7	109.3	81.3	6	62.5	77.4	8	32.0	56.7	10	15.7	7.2
		105.0	73.2		52.7	62.4		28.0	45.8		15.7	7.2
GP-PSat with ψ'_{\emptyset}	5	104.5	76.3	5	63.2	81.0	7	36.6	64.3	10	13.7	25.0
		92.0	52.7		51.2	62.1		29.8	49.0		13.7	25.0

such helpful Bayesian networks. As we discussed in §1, we expect that the theory of probabilistic programming languages will shed light on the problem, but at any rate we need more practical example scenarios to evaluate the viability of our approach.

Moreover, we conceive that our proposed algorithm in §4 contains the potential for many improvements. As we note in §2.2.2, the acquisition function in **GP-PSat** is simple, but not the state-of-the-art in the field of Gaussian process optimization. Extending our approach to other type of the acquisition function is not straightforward, but we think it is within possibility.

References

- [1] Bayes net toolbox for matlab. <http://www.gaussianprocess.org/gpml/code/matlab/doc/>. Accessed: 2017-01-25.
- [2] Breach toolbox. <https://bitbucket.org/donze/breach>. Accessed: 2017-01-25.
- [3] Gpml matlab code. <https://github.com/bayesnet/bnt>. Accessed: 2017-01-25.
- [4] Takumi Akazaki. Falsification of conditional safety properties for cyber-physical systems with gaussian process regression. In Yliès Falcone and César Sánchez, editors, *Runtime Verification - 16th International Conference, RV 2016, Madrid, Spain, September 23-30, 2016, Proceedings*, volume 10012 of *Lecture Notes in Computer Science*, pages 439–446. Springer, 2016, doi:10.1007/978-3-319-46982-9_27.
- [5] Takumi Akazaki and Ichiro Hasuo. Time robustness in MTL and expressivity in hybrid system falsification. In Daniel Kroening and Corina S. Pasareanu, editors, *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part II*, volume 9207 of *Lecture Notes in Computer Science*, pages 356–374. Springer, 2015, doi:10.1007/978-3-319-21668-3_21.
- [6] Takumi Akazaki, Ichiro Hasuo, and Kohei Suenaga. Input synthesis for sampled data systems by program logic. In Manuela L. Bujorianu and Rafael Wisniewski, editors, *Proceedings 4th Workshop on Hybrid Autonomous Systems, HAS 2014, Grenoble, France, 12-13 April 2014.*, volume 174 of *EPTCS*, pages 22–39, 2014, doi:10.4204/EPTCS.174.3.
- [7] Rajeev Alur, Tomás Feder, and Thomas A. Henzinger. The benefits of relaxing punctuality. *J. ACM*, 43(1):116–146, January 1996, doi:10.1145/227595.227602.
- [8] Yashwanth Annpureddy, Che Liu, Georgios E. Fainekos, and Sriram Sankaranarayanan. S-TaLiRo: A tool for temporal logic falsification for hybrid systems. In Parosh Aziz Abdulla and K. Rustan M. Leino, editors, *TACAS*, volume 6605 of *Lecture Notes in Computer Science*, pages 254–257. Springer, 2011, doi:10.1007/978-3-642-19835-9_21.
- [9] Gang Chen, Zachary Sabato, and Zhaodan Kong. Active learning based requirement mining for cyber-physical systems. In *55th IEEE Conference on Decision and Control, CDC 2016, Las Vegas, NV, USA, December 12-14, 2016*, pages 4586–4593. IEEE, 2016, doi:10.1109/CDC.2016.7798967.

- [10] Gang Chen, Zachary Sabato, and Zhaodan Kong. Active requirement mining of bounded-time temporal properties of cyber-physical systems. *CoRR*, abs/1603.00814, 2016.
- [11] Jyotirmoy V. Deshmukh, Xiaoqing Jin, James Kapinski, and Oded Maler. Stochastic local search for falsification of hybrid systems. In Bernd Finkbeiner, Geguang Pu, and Lijun Zhang, editors, *Automated Technology for Verification and Analysis - 13th International Symposium, ATVA 2015, Shanghai, China, October 12-15, 2015, Proceedings*, volume 9364 of *Lecture Notes in Computer Science*, pages 500–517. Springer, 2015, doi:10.1007/978-3-319-24953-7_35.
- [12] Adel Dokhanchi, Aditya Zutshi, Rahul T. Sriniva, Sriram Sankaranarayanan, and Georgios E. Fainekos. Requirements driven falsification with coverage metrics. In Alain Girault and Nan Guan, editors, *2015 International Conference on Embedded Software, EMSOFT 2015, Amsterdam, Netherlands, October 4-9, 2015*, pages 31–40. IEEE, 2015, doi:10.1109/EMSOFT.2015.7318257.
- [13] Alexandre Donzé. Breach, A toolbox for verification and parameter synthesis of hybrid systems. In Tayssir Touili, Byron Cook, and Paul Jackson, editors, *Computer Aided Verification, 22nd International Conference, CAV 2010, Edinburgh, UK, July 15-19, 2010. Proceedings*, volume 6174 of *Lecture Notes in Computer Science*, pages 167–170. Springer, 2010, doi:10.1007/978-3-642-14295-6_17.
- [14] Alexandre Donzé and Oded Maler. Robust satisfaction of temporal logic over real-valued signals. In Krishnendu Chatterjee and Thomas A. Henzinger, editors, *Formal Modeling and Analysis of Timed Systems - 8th International Conference, FORMATS 2010, Klosterneuburg, Austria, September 8-10, 2010. Proceedings*, volume 6246 of *Lecture Notes in Computer Science*, pages 92–106. Springer, 2010, doi:10.1007/978-3-642-15297-9_9.
- [15] Georgios E. Fainekos and George J. Pappas. Robustness of temporal logic specifications for continuous-time signals. *Theor. Comput. Sci.*, 410(42):4262–4291, 2009, doi:10.1016/j.tcs.2009.06.021.
- [16] Bardh Hoxha, Houssam Abbas, and Georgios Fainekos. Benchmarks for temporal logic requirements for automotive systems. In Goran Frehse and Matthias Althoff, editors, *1st and 2nd International Workshop on Applied veRification for Continuous and Hybrid Systems*, volume 34 of *EPiC Series in Computing*, pages 25–30, EasyChair, 2014.
- [17] Bardh Hoxha, Houssam Abbas, and Georgios E. Fainekos. Using S-TaLiRo on industrial size automotive models. In Goran Frehse and Matthias Althoff, editors, *1st and 2nd International Workshop on Applied veRification for Continuous and Hybrid Systems*, volume 34 of *EPiC Series in Computing*, pages 113–119. EasyChair, 2014.
- [18] Claire Jones. *Probabilistic Non-Determinism*. PhD thesis, Univ. Edinburgh, 1990.
- [19] Dexter Kozen. Semantics of probabilistic programs. *J. Comput. Syst. Sci.*, 22(3):328–350, 1981, doi:10.1016/0022-0000(81)90036-2.
- [20] Oded Maler and Dejan Nickovic. Monitoring temporal properties of continuous signals. In Yassine Lakhnech and Sergio Yovine, editors, *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems, Joint International Conferences on Formal Modelling and Analysis of Timed Systems, FORMATS 2004 and Formal Techniques in Real-Time and Fault-Tolerant Systems, FTRTFT 2004, Grenoble, France, September 22-24, 2004, Proceedings*, volume 3253 of *Lecture Notes in Computer Science*, pages 152–166. Springer, 2004, doi:10.1007/978-3-540-30206-3_12.
- [21] Jonas Mockus. *Bayesian approach to global optimization: theory and applications*. Mathematics and its applications (Kluwer Academic Publishers).: Soviet series. Kluwer Academic, 1989.
- [22] Carroll Morgan, Annabelle McIver, and Karen Seidel. Probabilistic predicate transformers. *ACM Trans. Program. Lang. Syst.*, 18(3):325–353, 1996, doi:10.1145/229542.229547.
- [23] Federico Olmedo, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Christoph Matheja. Reasoning about recursive probabilistic programs. In Martin Grohe, Eric Koskinen, and Natarajan Shankar, editors, *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*, pages 672–681. ACM, 2016, doi:10.1145/2933575.2935317.

- [24] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.
- [25] Sriram Sankaranarayanan and Georgios Fainekos. Falsification of temporal properties of hybrid systems using the cross-entropy method. In *Proceedings of the 15th ACM International Conference on Hybrid Systems: Computation and Control*, HSCC '12, pages 125–134, New York, NY, USA, 2012. ACM, doi:10.1145/2185632.2185653.
- [26] Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias W. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In Johannes Fürnkranz and Thorsten Joachims, editors, *Proceedings of the 27th International Conference on Machine Learning (ICML-10), June 21-24, 2010, Haifa, Israel*, pages 1015–1022. Omnipress, 2010, doi:10.1109/TIT.2011.2182033.
- [27] Aditya Zutshi, Jyotirmoy V. Deshmukh, Sriram Sankaranarayanan, and James Kapinski. Multiple shooting, cegar-based falsification for hybrid systems. In *Proceedings of the 14th International Conference on Embedded Software*, EMSOFT '14, pages 5:1–5:10, New York, NY, USA, 2014. ACM, doi:10.1145/2656045.2656061.
- [28] Aditya Zutshi, Sriram Sankaranarayanan, Jyotirmoy V. Deshmukh, and James Kapinski. A trajectory splicing approach to concretizing counterexamples for hybrid systems. In *Proceedings of the 52nd IEEE Conference on Decision and Control, CDC 2013, December 10-13, 2013, Firenze, Italy*, pages 3918–3925. IEEE, 2013, doi:10.1109/CDC.2013.6760488.