

Runtime Monitoring under Uncertainty in Autonomous Cyber-Physical Systems

Hazem Torfah


Department of Computer Science and Engineering
Chalmers University of Technology

**@Workshop on Formal Verification and Learning
for Quantitative Systems | FSTTCS 25 | Dec 16, 2025**

WASP | WALLENBERG AI,
AUTONOMOUS SYSTEMS
AND SOFTWARE PROGRAM


CHALMERS
UNIVERSITY OF TECHNOLOGY

Runtime monitoring




Runtime monitoring is a technique based on **extracting** information from a system at runtime and **evaluating** this information to determine whether an execution of the system satisfies or violates a given property, and consequently **deploying** the necessary recovery mechanisms.


Runtime monitoring in autonomous systems



Runtime monitoring in autonomous systems



Runtime monitor: Estimate the risk of colliding with one of the vehicles




In this talk


- ◎ **Runtime monitoring under uncertainty**
 - *formal problem statement*
- ◎ **Monitoring over Markov decision processes**
 - *from exponential to polynomial complexity*
- ◎ **Learning robust uncertainty models for monitoring**
 - *point-based models vs interval-based models*
 - *model-based vs model-free*

Problem statement

Runtime assurance in AI-based autonomy



Problem Statement



Let: $V_{sys} = V_E \cup V_C$ and $\Sigma_{sys}, \Sigma_{obs}$ denote the set of valuations over the variables V_{sys}, V_{obs}



Problem Statement: Given a distribution $d \in \text{Dist}(\Sigma_{sys}^*)$, a mapping $\mu: \Sigma_{sys}^* \rightarrow \text{Dist}(\Sigma_{obs}^*)$, a risk function $r: \Sigma_{sys}^* \rightarrow \mathbb{R}$, and a constant λ , learn a monitor M such that:

$$\forall \sigma_{obs}. M(\sigma_{obs}) \Leftrightarrow \sum_{\sigma_{sys} \in \Sigma_{sys}^*} d(\sigma_{sys}) \cdot \mu(\sigma_{sys})(\sigma_{obs}) \cdot r(\sigma_{sys}) > \lambda$$

Monitoring over Markov decision processes

Junges, Torfah, Seshia. CAV 2021

Monitoring MDPs



What is the risk after observing a trace τ ?

$$\text{Risk}(\tau) = \sup_{\sigma \in \text{Strategies}(M)} \sum_{\pi \in \text{Paths}(M, |\tau|)} \Pr_{\sigma}(\pi \mid \tau) \cdot r(\text{last}(\pi))$$

Algorithmic Approach

Filtering

Too expensive
Needs to track
exponential number of beliefs.

- *Algorithmic analysis of nonlinear hybrid systems. IEEE Trans. Autom Control 1998. T. Henzinger et al.*
- *Runtime Verification with State Estimation. RV 11. S. Stoller et. al -*
- *Runtime Verification of Stochastic, Faulty Systems. RV 10. M. Wilcox, B Williams -*
- *Monitoring Temporal Properties of Stochastic Systems. VMCAI 2008. A. Sistla, A.R. Srinivas -*


Algorithmic Approach

Filtering

Too expensive
Needs to track
exponential number of beliefs.

Trace length

Number of beliefs



Learning robust uncertainty models

Skurka, van der Maas, Junges, Torfah. AAMAS 2026

