Reconciling Formal Methods with Metrology

Improving Verification Verdicts of Traditional Hybrid Automata

Paul Kröger   Martin Fränzle

Research Group Hybrid Systems
Carl von Ossietzky Universität Oldenburg

2022-03-25

CARL VON OSSIETZKY universität OLDENBURG

Outline

1. The parking car:
   A toy example.
2. Why traditional hybrid automata models fail.
3. Bayesian hybrid automata:
   Hybrid automata incorporating Bayesian-style state-estimates.
   a) Incorporating probability density functions.
   b) The impact of hybrid dynamics.

The parking car

Setup

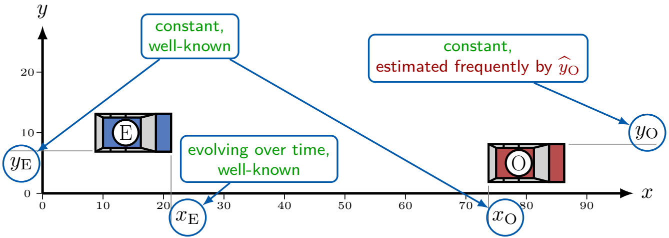

Car E can:

- travel straight-line with constant speed
- stop
- switch between dynamics instantaneously

Desired system properties:

- **safe** $:= (y_E \leq y_O) \Rightarrow$ **AG** $(x_E < x_O)$
- **live** $:= (y_E > y_O) \Rightarrow$ **AF** $(x_E \geq x_O)$

Nondeterministic hybrid automata

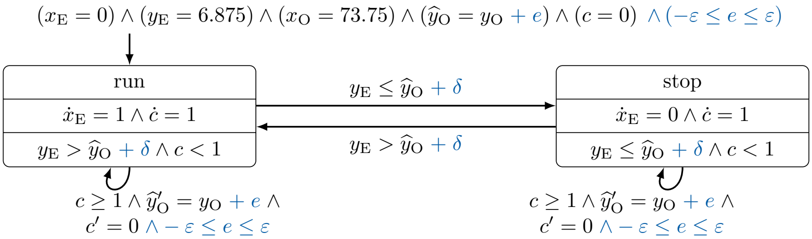

Stochastic hybrid automata

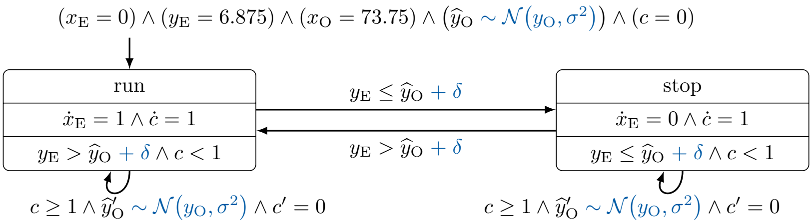

Why traditional models fail

Nondeterministic modelling

Estimated datum:

- uncontrollable measurement error: $\widehat{y}_{\mathrm{O}} = y_{\mathrm{O}} + e$
- error nondeterministic but bounded: $-\varepsilon \leq e \leq +\varepsilon$
- resolve nondeterminism demonically

Decision making:

- $y_{\mathrm{E}} > \widehat{y}_{\mathrm{O}} + \delta \Leftrightarrow$ go ahead

Nondeterministic modelling

Estimated datum:

- uncontrollable measurement error: $\widehat{y}_{\mathrm{O}} = y_{\mathrm{O}} + e$
- error nondeterministic but bounded: $-\varepsilon \leq e \leq +\varepsilon$
- resolve nondeterminism demonically

Decision making:

- $y_{\mathrm{E}} > \widehat{y}_{\mathrm{O}} + \delta \Leftrightarrow$ go ahead

A "pathological" case:

	safe	**live**
$\delta + \max(\widehat{y}_{\mathrm{O}}) \geq y_{\mathrm{E}} > y_{\mathrm{O}}$	trivial	unsat

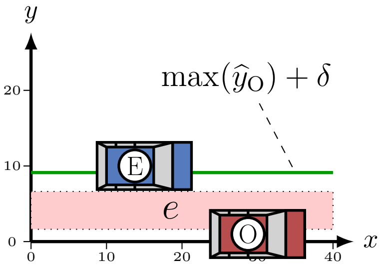

Stochastic modelling

Estimated datum:

- uncontrollable measurement error: $\widehat{y}_{\mathrm{O}} = y_{\mathrm{O}} + e$
- quantify errors by distribution, e.g. $e \sim \mathcal{N}(\mu, \sigma^2)$
- set safety margin δ s.t. $P(\widehat{y}_{\mathrm{O}} + \delta < y_{\mathrm{O}}) < \theta$

Decision making:

- $y_{\mathrm{E}} > \widehat{y}_{\mathrm{O}} + \delta \Leftrightarrow$ go ahead

Stochastic modelling

Estimated datum:

- uncontrollable measurement error: $\widehat{y}_{\mathrm{O}} = y_{\mathrm{O}} + e$
- quantify errors by distribution, e.g. $e \sim \mathcal{N}(\mu, \sigma^2)$
- set safety margin δ s.t. $P(\widehat{y}_{\mathrm{O}} + \delta < y_{\mathrm{O}}) < \theta$

Decision making:

- $y_{\mathrm{E}} > \widehat{y}_{\mathrm{O}} + \delta \Leftrightarrow$ go ahead

A “pathological” case:

	p(**safe**)	p(**live**)
$y_{\mathrm{E}} \leq y_{\mathrm{O}}$	$\to 0$	trivial

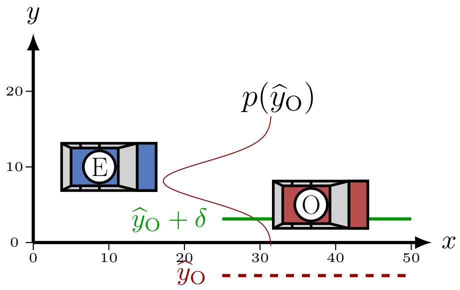

Basic idea

Existing flavours of HA are

- too optimistic, or
- too pessimistic.

Basic idea

Existing flavours of HA are

- too optimistic, or
- too pessimistic.

This does not reflect the behaviour of real-world systems.

Basic idea

Existing flavours of HA are

- too optimistic, or
- too pessimistic.

This does not reflect the behaviour of real-world systems.

What is done in real-world systems (sketch):

1. Convert (frequent) observations into proper estimates.

Basic idea

Existing flavours of HA are

- too optimistic, or
- too pessimistic.

This does not reflect the behaviour of real-world systems.

What is done in real-world systems (sketch):

1. Convert (frequent) observations into proper estimates.
   $\Rightarrow \quad \widehat{x} = 5 \quad \rightsquigarrow \quad p(x) \equiv \mathcal{N}(5, \sigma^2)$

Basic idea

Existing flavours of HA are

- too optimistic, or
- too pessimistic.

This does not reflect the behaviour of real-world systems.

What is done in real-world systems (sketch):

1. Convert (frequent) observations into proper estimates.
   $\Rightarrow \quad \widehat{x} = 5 \quad \rightsquigarrow \quad p(x) \equiv \mathcal{N}(5, \sigma^2)$
2. Combine all observations.

Basic idea

Existing flavours of HA are

- too optimistic, or
- too pessimistic.

This does not reflect the behaviour of real-world systems.

What is done in real-world systems (sketch):

1. Convert (frequent) observations into proper estimates.
   $\Rightarrow \quad \widehat{x} = 5 \quad \rightsquigarrow \quad p(x) \equiv \mathcal{N}(5, \sigma^2)$
2. Combine all observations.
   $\Rightarrow \quad p(x) = w_1 \cdot p_1(x) + w_2 \cdot p_2(x) + \cdots$

Basic idea

Existing flavours of HA are

- too optimistic, or
- too pessimistic.

This does not reflect the behaviour of real-world systems.

What is done in real-world systems (sketch):

1. Convert (frequent) observations into proper estimates.
   $\Rightarrow \quad \widehat{x} = 5 \quad \rightsquigarrow \quad p(x) \equiv \mathcal{N}(5, \sigma^2)$
2. Combine all observations.
   $\Rightarrow \quad p(x) = w_1 \cdot p_1(x) + w_2 \cdot p_2(x) + \cdots$
3. Make “rational” decisions based on combined estimates.

Basic idea

Existing flavours of HA are

- too optimistic, or
- too pessimistic.

This does not reflect the behaviour of real-world systems.

What is done in real-world systems (sketch):

1. Convert (frequent) observations into proper estimates.
   $\Rightarrow \quad \widehat{x} = 5 \quad \rightsquigarrow \quad p(x) \equiv \mathcal{N}(5, \sigma^2)$
2. Combine all observations.
   $\Rightarrow \quad p(x) = w_1 \cdot p_1(x) + w_2 \cdot p_2(x) + \cdots$
3. Make "rational" decisions based on combined estimates.
   $\Rightarrow \quad$ stop if $p(\text{safe}) < \varepsilon$

Basic idea

Existing flavours of HA are

- too optimistic, or
- too pessimistic.

This does not reflect the behaviour of real-world systems.

What is done in real-world systems (sketch):

1. Convert (frequent) observations into proper estimates.
   $\Rightarrow \quad \widehat{x} = 5 \quad \rightsquigarrow \quad p(x) \equiv \mathcal{N}(5, \sigma^2)$
2. Combine all observations.
   $\Rightarrow \quad p(x) = w_1 \cdot p_1(x) + w_2 \cdot p_2(x) + \cdots$
3. Make "rational" decisions based on combined estimates.
   $\Rightarrow \quad \text{stop if } p(\text{safe}) < \varepsilon$

Can we adopt this for Hybrid-System Theory?

Combining observations: Bayesian inference

1. Build up evidence over measurement history via Bayesian inference.
   - For normally distributed measurement errors and linear dynamics: Kálmán filter.

Combining observations: Bayesian inference

1. Build up evidence over measurement history via Bayesian inference.
   - For normally distributed measurement errors and linear dynamics: Kálmán filter.

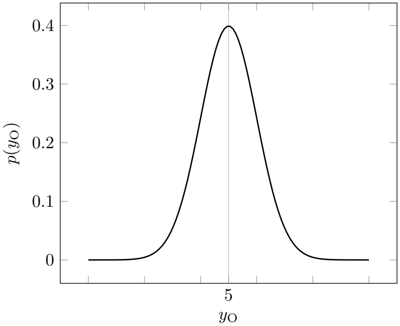

Combining observations: Bayesian inference

1. Build up evidence over measurement history via Bayesian inference.
   - For normally distributed measurement errors and linear dynamics: Kálmán filter.

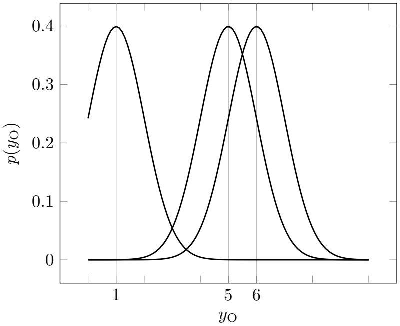

Combining observations: Bayesian inference

1. Build up evidence over measurement history via Bayesian inference.
   - For normally distributed measurement errors and linear dynamics: Kálmán filter.

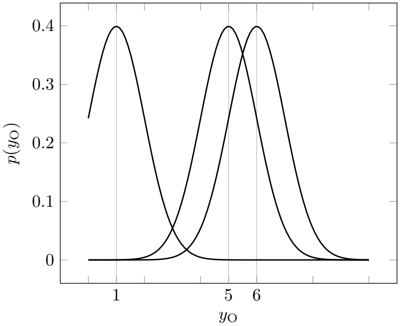

$\rightsquigarrow$

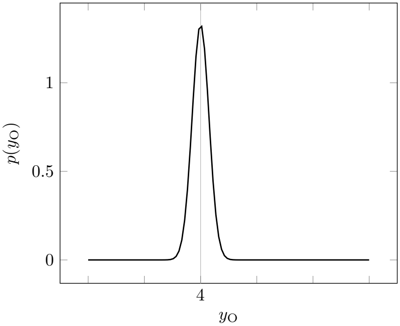

Combining observations: Bayesian inference

1. Build up evidence over measurement history via Bayesian inference.
   - For normally distributed measurement errors and linear dynamics: Kálmán filter.

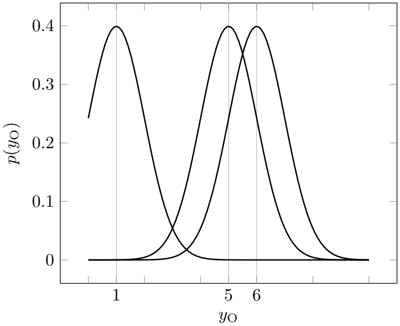

$\rightsquigarrow$

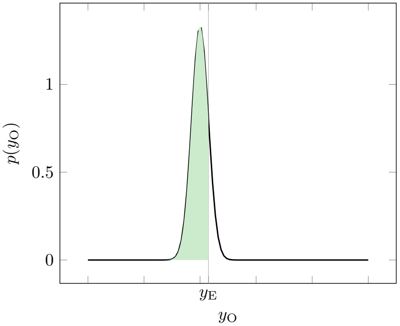

2. Make rational decisions: $P(y_{\mathrm{E}} > y_{\mathrm{O}}) > \theta \Leftrightarrow$ go ahead.

Bayesian hybrid automata

Incorporating probability density functions

Bayesian inference in hybrid automata

- State distributions become first class members of the state space.
- Transitions/locations are equipped with
  - mechanisms for applying Bayesian updates on measurements,
  - and guards/invariants accessing estimates.
- Prediction between measurements requires an application of the correct (!) dynamics to distributions.

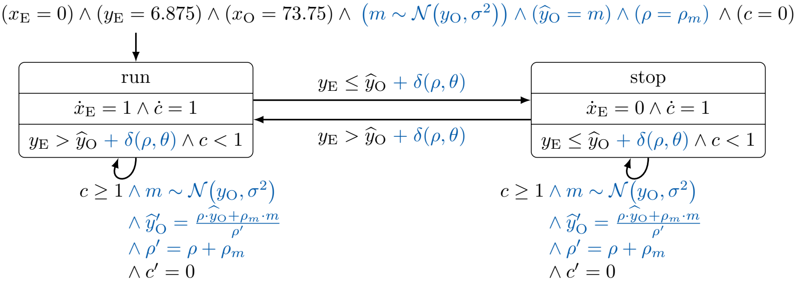

Bayesian cars

Estimated datum:

- is a probability density function $p(y_\mathrm{O})$
- updated by means of a Bayes filter

Decision making:

- $P(y_\mathrm{E} \leq y_\mathrm{O}) > \delta \Leftrightarrow$ stop

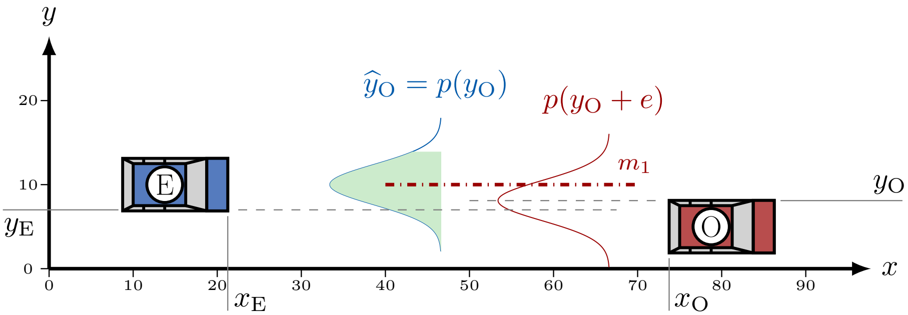

Bayesian cars

Estimated datum:

- is a probability density function $p(y_{\mathrm{O}})$
- updated by means of a Bayes filter

Decision making:

- $P(y_{\mathrm{E}} \leq y_{\mathrm{O}}) > \delta \Leftrightarrow$ stop

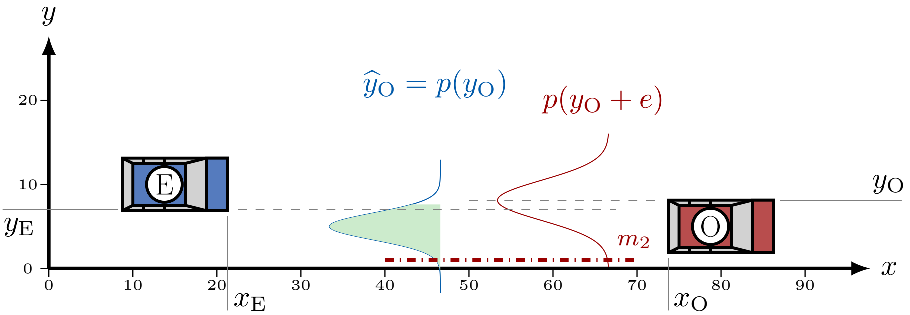

Bayesian cars

Estimated datum:

- is a probability density function $p(y_\mathrm{O})$
- updated by means of a Bayes filter

Decision making:

- $P(y_\mathrm{E} \leq y_\mathrm{O}) > \delta \Leftrightarrow$ stop

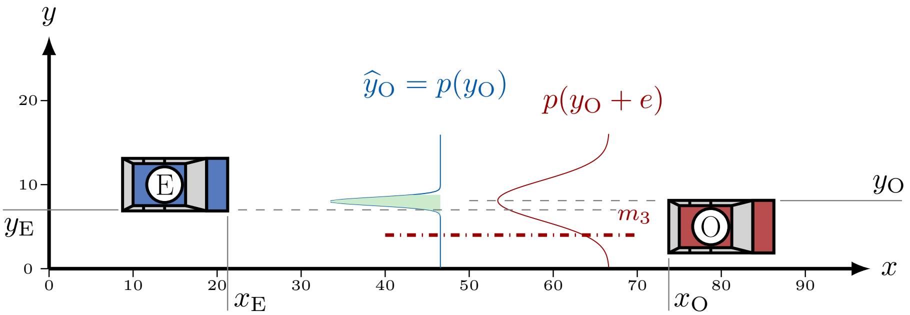

Bayesian cars

Estimated datum:

- is a probability density function $p(y_\mathrm{O})$
- updated by means of a Bayes filter

Decision making:

- $P(y_\mathrm{E} \leq y_\mathrm{O}) > \delta \Leftrightarrow$ stop

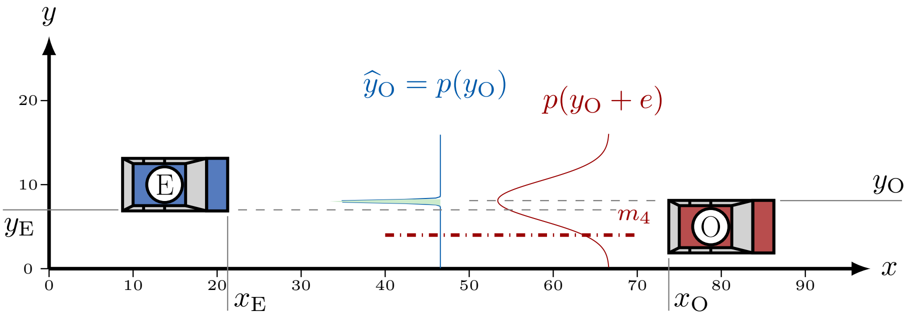

Experiment: safety

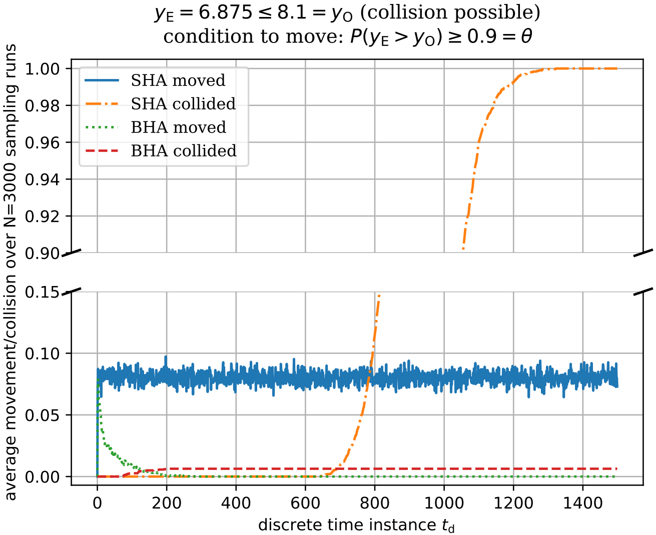

Experiment: liveness

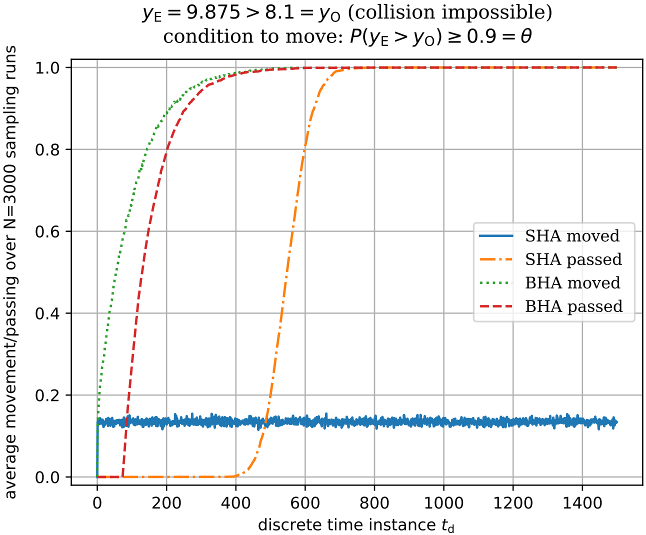

The impact of hybrid dynamics

Yet another toy example

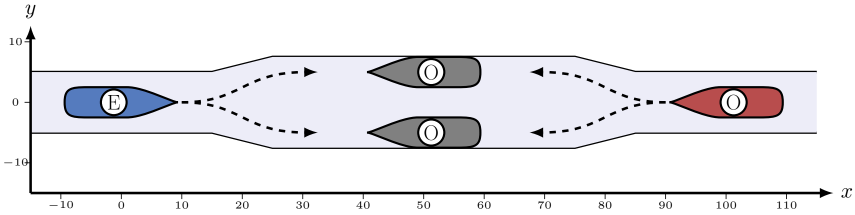

Ship O . . .

- chooses a direction for the evasive manoeuvre (left or right):
  - to the left, if $y_o \leq 0$
  - to the right, if $y_o > 0$

Ship E . . .

- is not aware of O's decision.
- chooses direction based on $\widehat{y}_{\mathrm{O}}$.

The consequence of hybrid dynamics

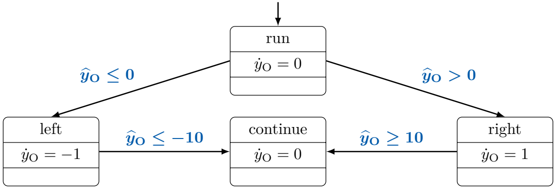

Apply correct mode dynamics to right part of the continous state space:

The consequence of hybrid dynamics

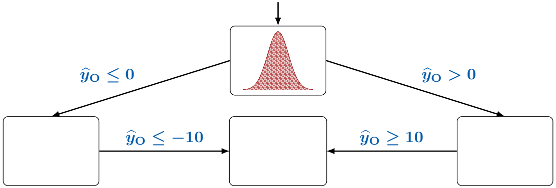

Apply correct mode dynamics to right part of the continous state space:

- guards enabled with some probability (yields probability of the mode)
- successor mode is ambiguous
- distribute the distribution over enabled transitions
  - $\longrightarrow$ mixture distributions for continuous state

The consequence of hybrid dynamics

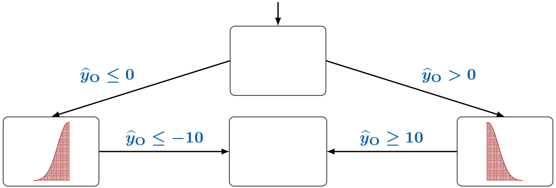

Apply correct mode dynamics to right part of the continous state space:

- guards enabled with some probability (yields probability of the mode)
- successor mode is ambiguous
- distribute the distribution over enabled transitions
  $\rightarrow$ mixture distributions for continuous state

The consequence of hybrid dynamics

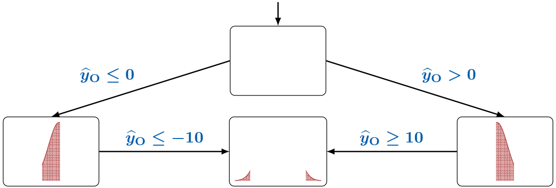

Apply correct mode dynamics to right part of the continous state space:

- guards enabled with some probability (yields probability of the mode)
- successor mode is ambiguous
- distribute the distribution over enabled transitions
  - $\longrightarrow$ mixture distributions for continuous state

The consequence of hybrid dynamics

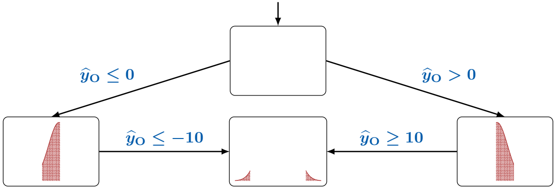

Apply correct mode dynamics to right part of the continous state space:

- guards enabled with some probability (yields probability of the mode)
- successor mode is ambiguous
- distribute the distribution over enabled transitions
  $\longrightarrow$ mixture distributions for continuous state

Estimate at time t:

- continuous state: weighted re-assembly from (partial) distributions
- discrete state: derived from probability mass shifted “into” the mode

Guess what I'm doing: hybrid estimation

So far, upon new measurements

- mixture components are updated (via filtering)
- but mode probabilites remain unchanged.

Guess what I'm doing: hybrid estimation

So far, upon new measurements

- mixture components are updated (via filtering)
- but mode probabilites remain unchanged.

However, measurements yield information about the true mode:

- Given a measurement, obtain the distribution of the true continuous state according to that measurement: $\widehat{x} = 5 \quad \rightsquigarrow \quad p(x) \equiv \mathcal{N}(5, \sigma^2)$.
- Reweighted probability mass of mode invariant under this distribution yields probability of the mode according to the measurement result, e.g. via $\int_{\mathsf{inv}(\mathrm{run})} p(x)$.
- This gives raise to a filter process for modes (e.g. using Bayes' rule).

Guess what I'm doing: hybrid estimation

So far, upon new measurements

- mixture components are updated (via filtering)
- but mode probabilites remain unchanged.

However, measurements yield information about the true mode:

- Given a measurement, obtain the distribution of the true continuous state according to that measurement: $\widehat{x} = 5 \quad \rightsquigarrow \quad p(x) \equiv \mathcal{N}(5, \sigma^2)$.
- Reweighted probability mass of mode invariant under this distribution yields probability of the mode according to the measurement result, e.g. via $\int_{\mathsf{inv}(\mathrm{run})} p(x)$.
- This gives raise to a filter process for modes (e.g. using Bayes' rule).

This sketches of the idea of currently ongoing work only.

Example

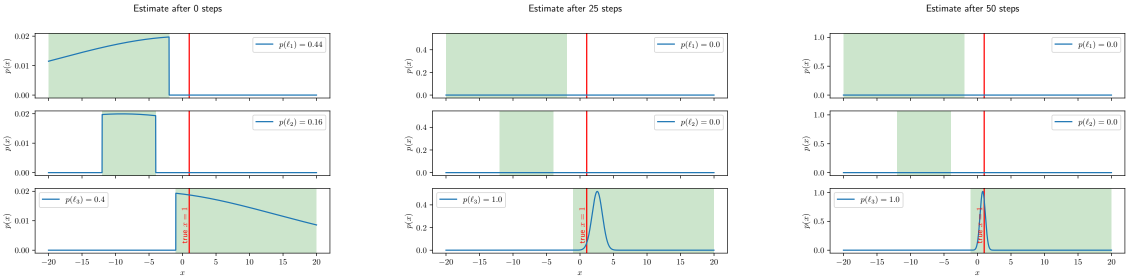

Estimate after 0 steps
Estimate after 25 steps
Estimate after 50 steps
$p(\ell_1) = 0.44$
$p(\ell_2) = 0.16$
$p(\ell_3) = 0.4$
$p(\ell_1) = 0.0$
$p(\ell_2) = 0.0$
$p(\ell_3) = 1.0$
$p(\ell_1) = 0.0$
$p(\ell_2) = 0.0$
$p(\ell_3) = 1.0$
true $x = 1$
$p(x)$
x

The decomposed model

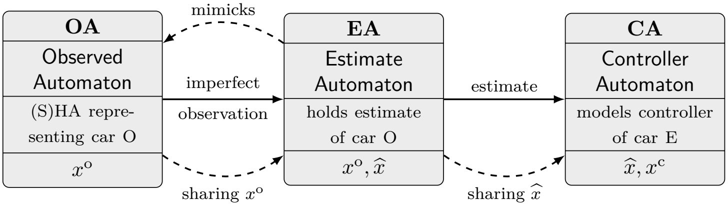

In case of more complex measurement processes another automaton modelling this process may be introduced between OA and EA.

Some papers

- M. Fränzle and P. Kröger.

  The demon, the gambler, and the engineer – reconciling hybrid-system theory with metrology.

  In *Symposium on Real-Time and Hybrid Systems*, volume 11180 of *Theoretical Computer Science and General Issues*, pages 165–185, Cham, 2018. Springer International Publishing.

- M. Fränzle and P. Kröger.

  Guess what I'm doing! Rendering Formal Verification Methods Ripe for the Era of interacting Intelligent Systems

  In *Leveraging Applications of Formal Methods, Verification and Validation: Applications*, pages 255–272, Cham, 2020. Springer International Publishing.

- P. Kröger and M. Fränzle.

  Bayesian hybrid automata: A formal model of justified belief in interacting hybrid systems subject to imprecise observation.

  accepted for LITES. 2021.