Peak Time-Windowed Mean Estimation using Convex Optimization

Jared Miller
Niklas Schmid
Matteo Tacchi
Didier Henrion
Roy S. Smith

December 16, 2024
IEEE CDC

AUTOMATIC CONTROL LABORATORY
ETH zürich

Safety/Risk Quantification

e.g. closest distance, max mean/CVAR current

Extremely context-dependent

Quantifying Risk

Many ways to quantify risk of state function $p(x)$

- Probability of entering unsafe set
- Mean of p
- 90% quantile of p
- Mean value above 90% quantile of p
- Other risk measures of p

Find peak (time-windowed) risk of p

Time-Windowed Risk Motivation: Signal

Oscillations near instanataneous peak ($t = 2$)

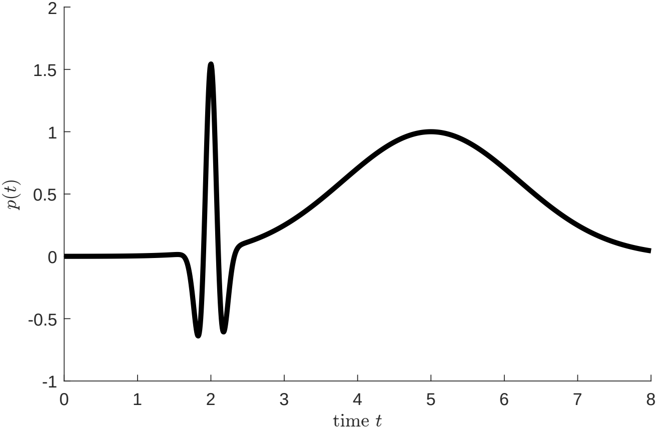

Time-Windowed Risk Example

Instantaneous maximal risks may not give full picture

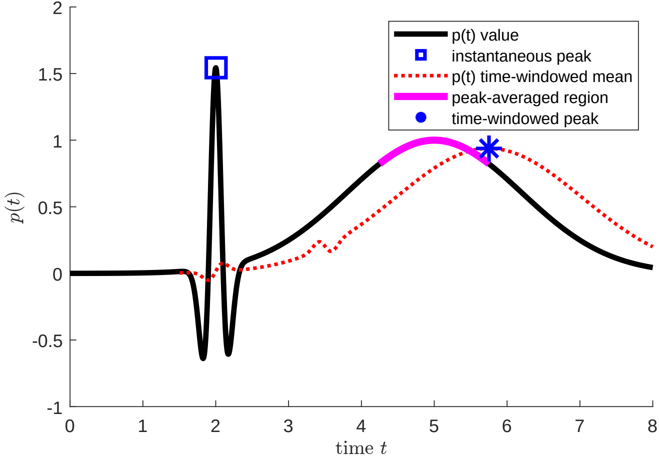

Large time-windowed avg. current on wire ≈ overheating

Time-Windowed Risk Example

Choose a time window h

Form a prob. dist. $\zeta(t)$ from $\{p(x(t'))\}_{t'=t-h}^{t}$

Analyze risk of $R(\zeta(t))$

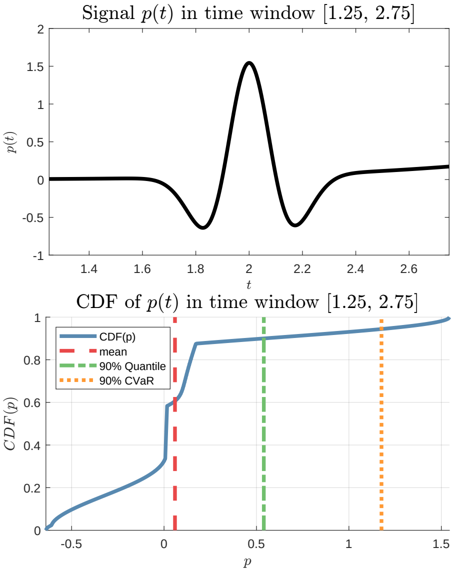

Time-Windowed Risk Problem

Given R and h, choose optimal t^*, x_0^*:

$$
\begin{aligned}
P^* = \sup_{t^*,\ x_0^*} \quad & R\left(\frac{1}{h}\int_{t^*-h}^{t^*} p(x(t'))dt'\right) \\
\text{s.t.} \quad & x(t) \text{ follows } \mathcal{L} \quad \forall t \in [0, \min(t^*, \tau_X)] \\
& x(0) = x_0^* \\
& t^* \in [h, T],\ x_0^* \in X_0
\end{aligned}
$$

Integral in objective collapses (marginalizes) time

CDC focus: R is **mean**

Solution Approach

Original risk estimation problem is nonconvex

Lift to *convex* but *infinite-dimensional* LP

Truncate infinite LP to computationally solve

How do we solve infinite LPs?

Discretization necessary to solve on computer

More complexity: more accurate solutions

Method	Increasing Complexity
Gridding	# Grid Points
Basis Functions	# Functions
Random Sampling	# Samples
$\star$ Sum-of-Squares (SOS)	Polynomial Degree
Your Favorite Method	Some Accuracy Parameter

Runtime usually exponential in dimension, complexity

Infeasibility: unsolvable problem or not enough compute?

Occupation Measures (stochastic)

μ: stochastic kernel from $\{\mu_t\}$

Set $\mapsto$ Avg. time spent in the set

Average: μ_0 and stoch. dynamics

Averaged value of $v \in \mathcal{C}$:
$\langle v, \mu \rangle = \int_0^T \mathbb{E}_{x \sim X_t}[v(t,x)]dt$

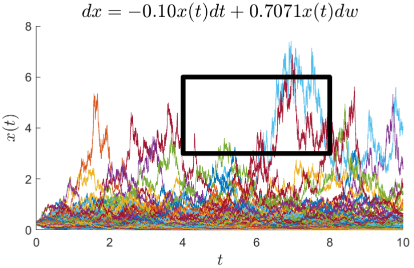

Box: set in (t, x)

Martingale Relation

End = Start + Accumulated Change (in $\mathbb{E}$)

$$\forall v \in \mathcal{C} : \ \mathbb{E}[v(t+s,x) \mid \mu_{t+s}] = \mathbb{E}[v(t,x) \mid \mu_t] + \int_{t'=t}^{t+s} \mathbb{E}[\mathcal{L}v(t',x) \mid \mu_{t'}]dt'$$

Relation between measures (μ_t, μ_{t+s}, μ) for all $v \in \mathcal{C}$

$$\langle v(t+s,x), \mu_{t+s}(x) \rangle = \langle v(t,x), \mu_t(x) \rangle + \langle \mathcal{L}v, \mu \rangle$$

Compress notation using adjoint $\mathcal{L}^\dagger$ (implicitly express $\forall v$)

$$\mu_{t+s} = \mu_t + \mathcal{L}^\dagger \mu$$

Augmented Time Coordinate

We can stick to ODE methods by adding a new time s

Two continuous times (t, s):

$$\begin{aligned} &\text{Active time} && t \in [0, T] && \dot{t} = 1 \\ &\text{Stopping time} && s \in [h, T] && \dot{s} = 0 \end{aligned}$$

Temporal support sets $\Omega_\pm$:

$$\Omega_- : \quad t \in [0, s-h] \qquad \qquad \Omega_+ : \quad t \in [s-h, s]$$

Risk evaluated in Ω_+, similar process in discrete-time

Two Time Coordinates?

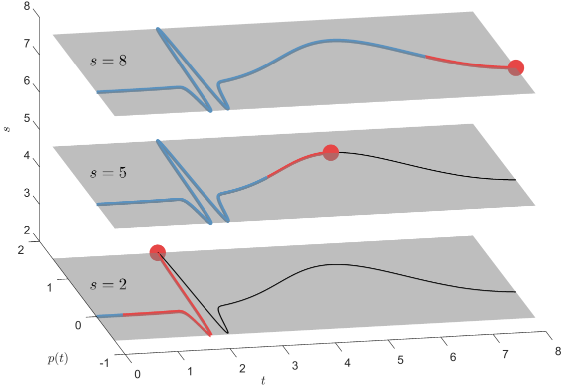

Curves $(t, p(t), s)$: time intervals $[0, s - h]$, $[s - h, s]$, $[s, T]$

Measures for Risk Estimation

Mean-type risk estimation measures (with constant state s)

$$
\begin{aligned}
\mu_0(s, x) &\in \mathcal{M}_+([h, T] \times X_0) && \text{Initial} \\
\mu_\tau(s, x) &\in \mathcal{M}_+([h, T] \times X) && \text{Terminal} \\
\mu_+(s, t, x) &\in \mathcal{M}_+(\Omega_+ \times X) && \text{Risk Occ.} \\
\mu_-(s, t, x) &\in \mathcal{M}_+(\Omega_- \times X) && \text{Past Occ.}
\end{aligned}
$$

Time-windowed risk evaluation: $\frac{1}{h}\int_{s-h}^{s} p(x(t'))dt' \to \frac{1}{h} p_\# \mu_+$

Time-Duplication Map

The last technical detail needed: a time-duplicating map φ

$$\varphi : (s, x) \mapsto (s, s, x)$$

For all test functions $\omega(s, t, x) \in C([h, T] \times [0, T] \times X)$

$$\langle \omega(s, t, x), \varphi_{\#}\mu_\tau(s, t, x) \rangle = \langle \omega(s, s, x), \mu_\tau(s, x) \rangle$$

Relaxed occupation measure of $\hat{\mathcal{L}}$: $(\mu_0, \varphi_{\#}\mu_\tau, \mu_+ + \mu_-)$

Time-Windowed Mean Estimation

Non-conservative infinite LP with generator $\hat{\mathcal{L}} : (\mathcal{L}, \dot{s} = 0)$

$$
\begin{aligned}
p^* = \ & \sup \quad \langle p, \mu_+ \rangle / h \\
& \text{s.t. } \varphi_{\#}\mu_\tau = \delta_0 \otimes \mu_0 + \hat{\mathcal{L}}^\dagger(\mu_- + \mu_+) \\
& \qquad \langle 1, \mu_0 \rangle = 1 \\
& \qquad \langle 1, \mu_+ \rangle = h \\
& \qquad \text{Mean-type time-windowed support constraints}
\end{aligned}
$$

Constraint $\langle 1, \mu_+ \rangle = h$ imposes that h time units elapse

CVaR modification: $\sup \mathrm{mean}(\psi) : \ \epsilon\psi + \hat{\psi} = (p_{\#}\mu_+)/h$

Time-Windowed Deterministic Mean $(h = 1.5)$

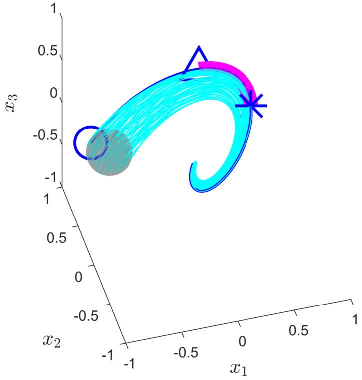

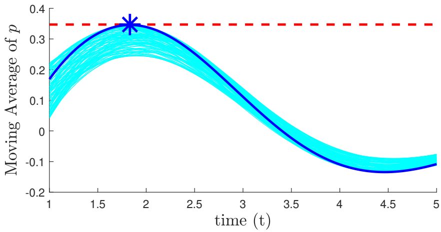

Time-Windowed Stoch. Mean Example ($h = 1.5$)

Instantaneous and time-windowed mean are separated ($p(x) = x_2$)

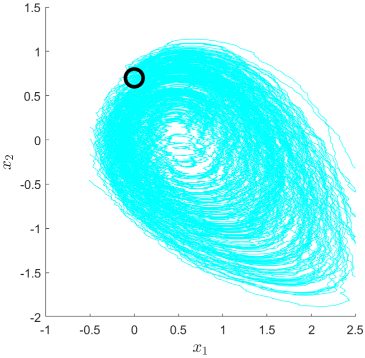

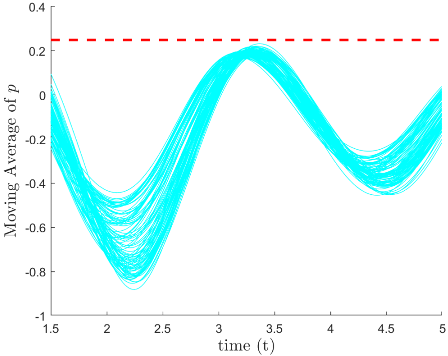

Time-Windowed Stoch. CVaR Example ($h = 1.5$)

Peak CVaR is close to peak instantaneous p (with $\epsilon = 0.15$)

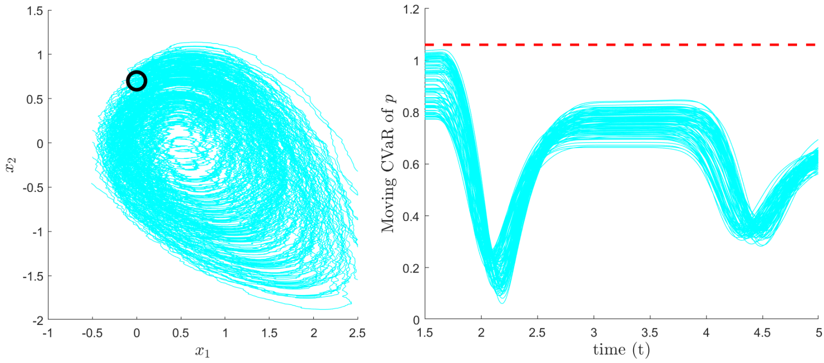

Take-aways

Conclusion

Time-windowed risk estimation

Solved using infinite-dimensional LPs/SOCPs in measures

Certified outer-approximations of risk

Thanks!

Bonus Slides

Assumptions for Stochastic LPs

Assumptions used in all presented programs[1]:

1. Trajectories stop upon the first exit from X ($\tau_X \wedge T$).
2. The test function set $\mathcal{C} = \mathrm{dom}(\mathcal{L})$ satisfies $\mathcal{C} \subseteq C([t_0, T] \times X)$ with $1 \in \mathcal{C}$ and $\mathcal{L}1 = 0$.
3. The set $\mathcal{C}$ separates points and is multiplicatively closed.
4. There exists a countable set $\{v_k\} \in \mathcal{C}$ such that $\forall v \in \mathcal{C}$: $(v, \mathcal{L}v)$ is contained in the bounded pointwise closure of the linear span of $\{(v_k, \mathcal{L}v_k)\}$.

[1]Cho, Moon Jung, and Richard H. Stockbridge. "Linear programming formulation for optimal stopping problems." SICON 40.6 (2002): 1965-1982.

Occupation Measure (Deterministic)

Time trajectories spend in set

Test function
$v(t, x) \in C([0, T] \times X)$

Single trajectory:
$\langle v, \mu \rangle = \int_0^T v(t, x(t \mid x_0))dt$

Averaged trajectory: $\langle v, \mu \rangle = \int_X \left(\int_0^T v(t, x)dt \right) d\mu_0(x)$

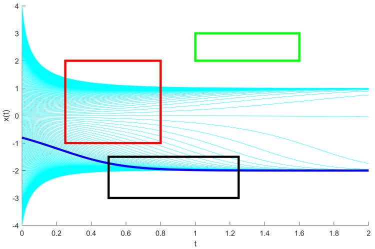

$x' = -x(x+2)(x-1)$

Unsafe Probability using Measures

Maximize prob. $\langle 1, \mu_p \rangle$ of ending in X_u (with $\mu_p + \mu_c = \mu_\tau$)

$$
\begin{aligned}
p^*(t_0, X_0) = \sup \quad & \langle 1, \mu_p \rangle \\
\text{s.t.} \quad & \mu_p + \mu_c = \delta_{t_0} \otimes \mu_0 + \mathcal{L}^\dagger \mu \\
& \langle 1, \mu_0 \rangle = 1 \\
& \mu_0 \in \mathcal{M}_+(X_0) \\
& \mu, \ \mu_c \in \mathcal{M}_+([t_0, T] \times X) \\
& \mu_p \in \mathcal{M}_+([t_0, T] \times X_u)
\end{aligned}
$$

Relaxed occupation measure $(\mu_0, \mu_u + \mu_c, \mu)$,

Strongly dual to previous continuous-function program

SOS Expectation-Peak

$$
\begin{aligned}
d^*_{\mathbb{E}} = \min \quad & \int_X v(0,x)\; d\mu_0(x) && \text{(2a)} \\
\text{s.t.} \quad & -\mathcal{L}v(t,x) \in \Sigma[[0,T] \times X] && \text{(2b)} \\
& v(t,x) - p(x) \in \Sigma[[0,T] \times X] && \text{(2c)} \\
& v \in \mathbb{R}[t,x] && \text{(2d)}
\end{aligned}
$$

SOS Concentration-Peak

Second-order cone $\mathbb{L}^n : \{(u, q) \in \mathbb{R}^n \times \mathbb{R}_{\geq 0} \mid q \geq \|u\|_2\}$

$$
\begin{aligned}
d_r^* = \min \quad & u_1 + 2u_3 + \int_{X_0} v(0, x_0) d\mu_0(x_0) && \text{(3a)} \\
\text{s.t.} \quad & -\mathcal{L}v(t, x) \in \Sigma[[0, T] \times X] && \text{(3b)} \\
& v(t, x) + u_1\, p^2(x) - 2\, u_2\, p(x) - p(x) && \text{(3c)} \\
& \qquad \in \Sigma[[0, T] \times X] \\
& ([u_1 + u_3, -(r/2), u_2], u_3) \in \mathbb{L}^3 && \text{(3d)} \\
& u \in \mathbb{R}^3, \;\; v \in \mathbb{R}[t, x]
\end{aligned}
$$

SOS CVaR-Peak

$$
\begin{aligned}
d_c^* = \min \quad & u + \int_X v(0,x)\ d\mu_0(x) && \text{(4a)}\\
\text{s.t.} \quad & -\mathcal{L}v(t,x) \in \Sigma[[0,T] \times X] && \text{(4b)}\\
& v(t,x) - w(p(x)) \in \Sigma[[0,T] \times X] && \text{(4c)}\\
& u + \epsilon\, w(q) - q \in \Sigma[p_{min}, p_{max}] && \text{(4d)}\\
& w(q) \in \Sigma[p_{min}, p_{max}] && \text{(4e)}\\
& u \in \mathbb{R}, v \in \mathbb{R}[t,x] && \text{(4f)}
\end{aligned}
$$

Time-Delay Approach (Bad, Don't Do This)

Embed as non-Markovian stochastic process:

$$
\begin{aligned}
P^* = \sup_{t^*,\ x_0^*} \quad & R\left(\beta(t^*)\right) \\
\text{s.t.} \quad & x(t) \text{ follows } \mathcal{L} \quad \forall t \in [0, \min(t^*, \tau_X)] \\
& d\beta = [p(x(t)) - p(x(t-h))](1/h)dt \\
& \beta(h) = (1/h) \textstyle\int_0^h p(x(t'))dt' \\
& x(0) = x_0^* \\
& t^* \in [h, T],\ x_0^* \in X_0
\end{aligned}
$$

Could introduce relaxation gap, requires $2n + 2$ states

SOS Time-Window Mean

$$
\begin{aligned}
d_k^* = \; & \min_{v,\gamma,\xi} \gamma + h\xi && \text{(5a)}\\
\text{s.t.} \quad & \gamma - v(s,0,x) \in \Sigma[[h,T] \times X_0] && \text{(5b)}\\
& v(t,t,x) \in \Sigma[[h,T] \times X]_{\leq 2k} && \text{(5c)}\\
& \xi - p(x)/h - \hat{\mathcal{L}} v(s,t,x) \in \Sigma[\Omega_+ \times X] && \text{(5d)}\\
& -\mathcal{L}_f v(s,t,x) \in \Sigma[\Omega_- \times X] && \text{(5e)}\\
& v \in \mathbb{R}[s,t,x] && \text{(5f)}\\
& \gamma, \xi \in \mathbb{R} && \text{(5g)}
\end{aligned}
$$

SOS Time-Window CVAR

$$d_k^* = \min_{v,\gamma,\xi,\beta,w} \gamma + h\xi + \beta \tag{6a}$$

$$\text{s.t.} \quad \gamma - v(s,0,x) \in \Sigma[[h,T] \times X_0] \tag{6b}$$

$$v(t,t,x) \in \Sigma[[h,T] \times X] \tag{6c}$$

$$\xi - w(p(x))/h - \hat{\mathcal{L}} v(s,t,x) \in \Sigma[\Omega_+ \times X] \tag{6d}$$

$$-\mathcal{L}_f v(s,t,x) \in \Sigma[\Omega_- \times X] \tag{6e}$$

$$w(q),\ \epsilon w(q) + \beta \in \Sigma[[p_{\min}, p_{\max}]] \tag{6f}$$

$$v \in \mathbb{R}[s,t,x] \tag{6g}$$

$$w \in \mathbb{R}[q] \tag{6h}$$

$$\gamma, \xi, \beta \in \mathbb{R} \tag{6i}$$

Time-Windowed Mean Example ($h = 1.5$)

Instantaneous and time-windowed mean are separated ($p(x) = x_2$)

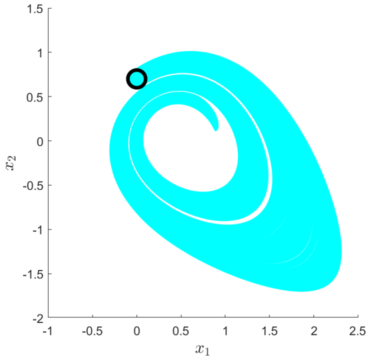

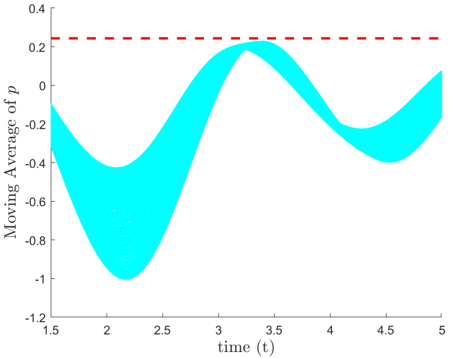

Time-Windowed CVaR Example ($h = 1.5$)

Peak CVaR is close to peak instantaneous p (with $\epsilon = 0.15$)