Myopically Verifiable Probabilistic Certificate for Long-term Safety

Yorie Nakahira

Control & Learning group @ CMU

理学

Stochastic safe control
Robust control
Optimization
Information theory …

科学

Neuroscience
Biomolecular control…

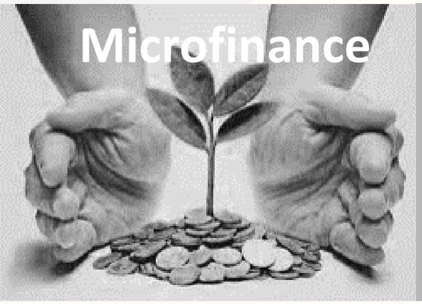

工学

Safety is critical for intelligent systems

Autonomous vehicles

Cobots
Intelligent manufacturing

Drones

Safety is critical for intelligent systems

Dealing with uncertainty

- Resilience
- Adaptability
- Collision Avoidance
- Regular Operation

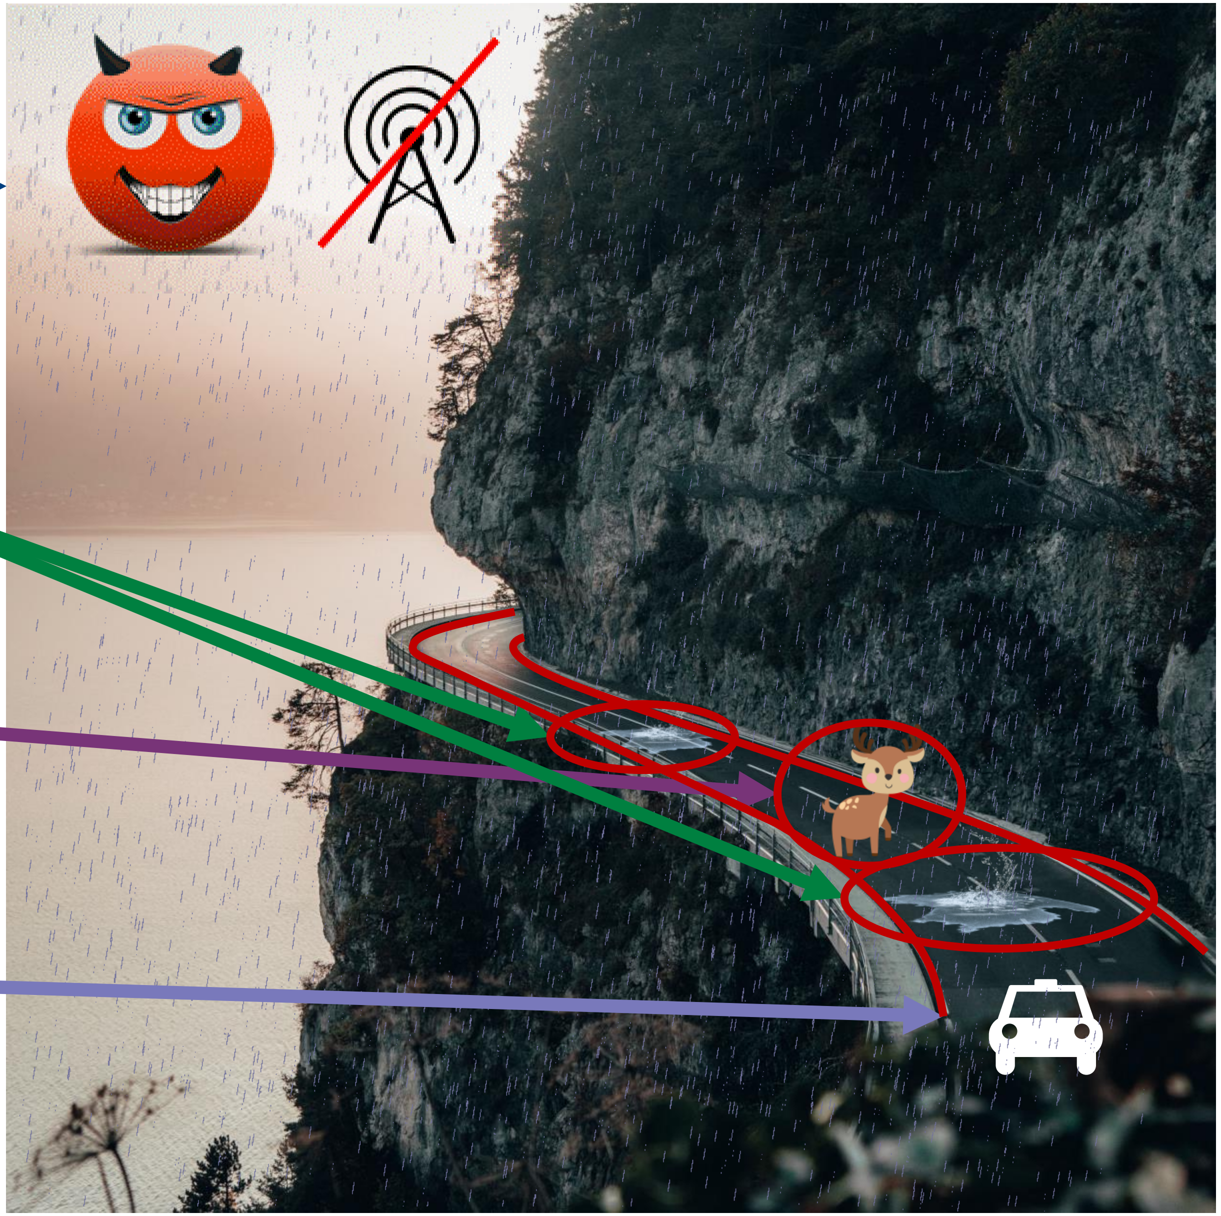

Challenges

Deterministic system

safe at each step → safe at all time

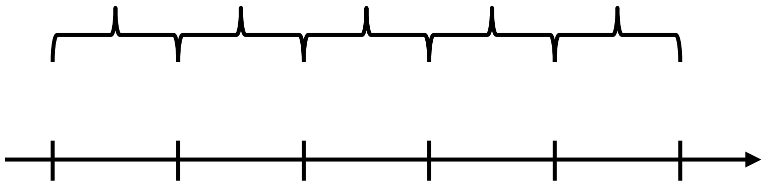

time

Stochastic system

safe at each step with probability $1-\delta$ → Long-term safe probability can scale according to $\lim_{t\to\infty}(1-\delta)^t=0$

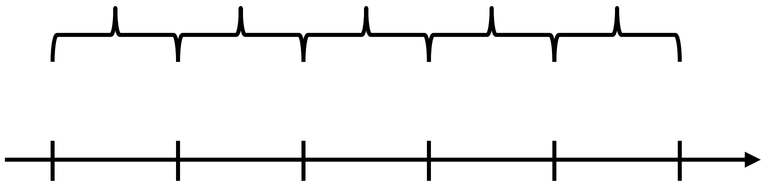

time

Challenges

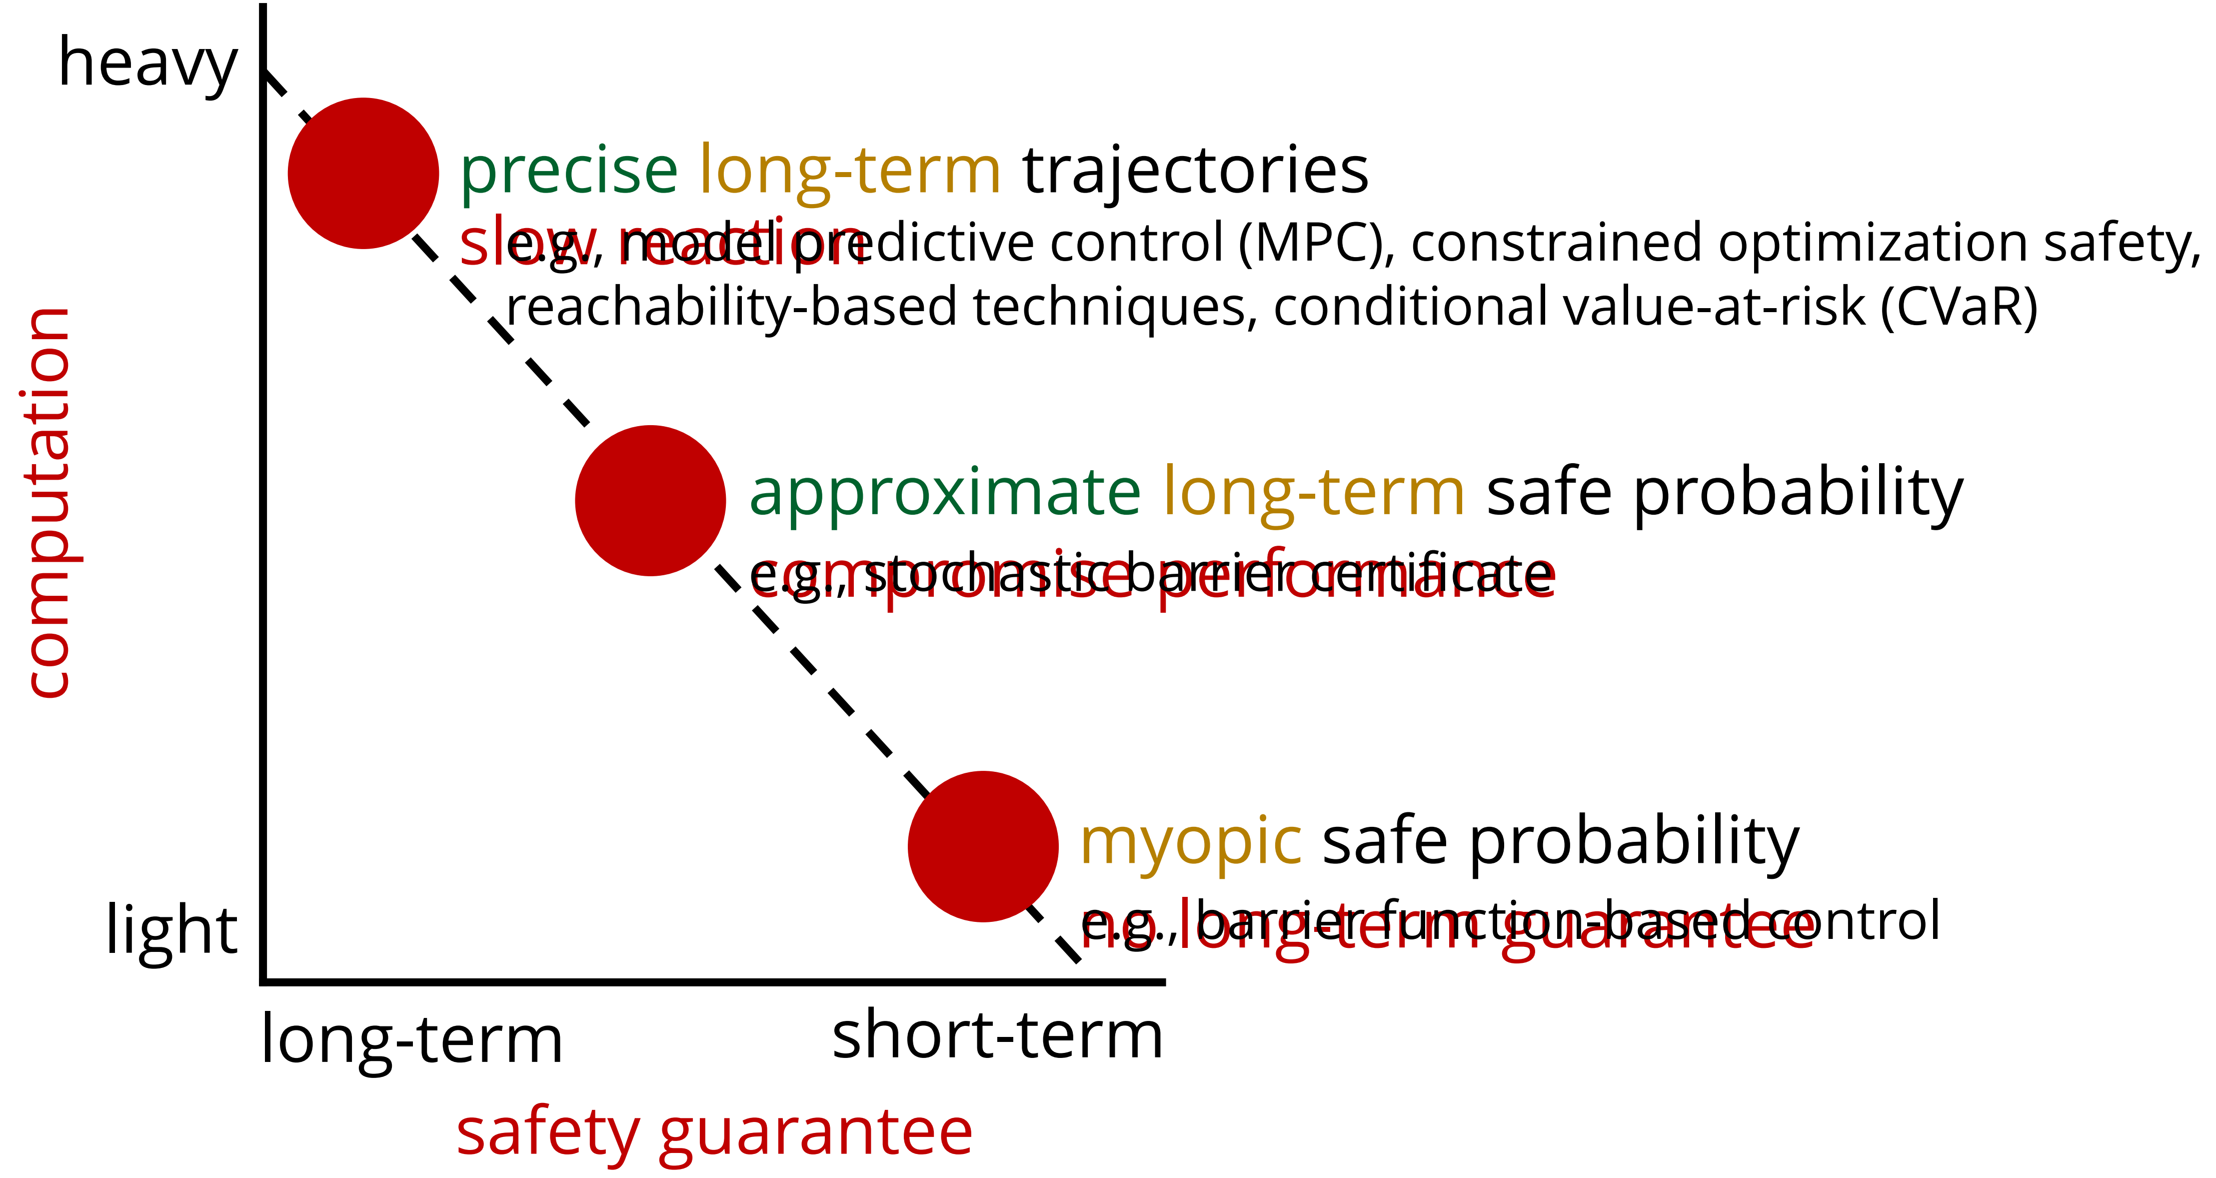
heavy
light
computation
precise long-term trajectories
slow reaction
e.g., model predictive control (MPC), constrained optimization safety,
reachability-based techniques, conditional value-at-risk (CVaR)
approximate long-term safe probability
compromise performance
e.g., stochastic barrier certificate
myopic safe probability
no long-term guarantee
e.g., barrier function-based control
long-term
short-term
safety guarantee

Challenges

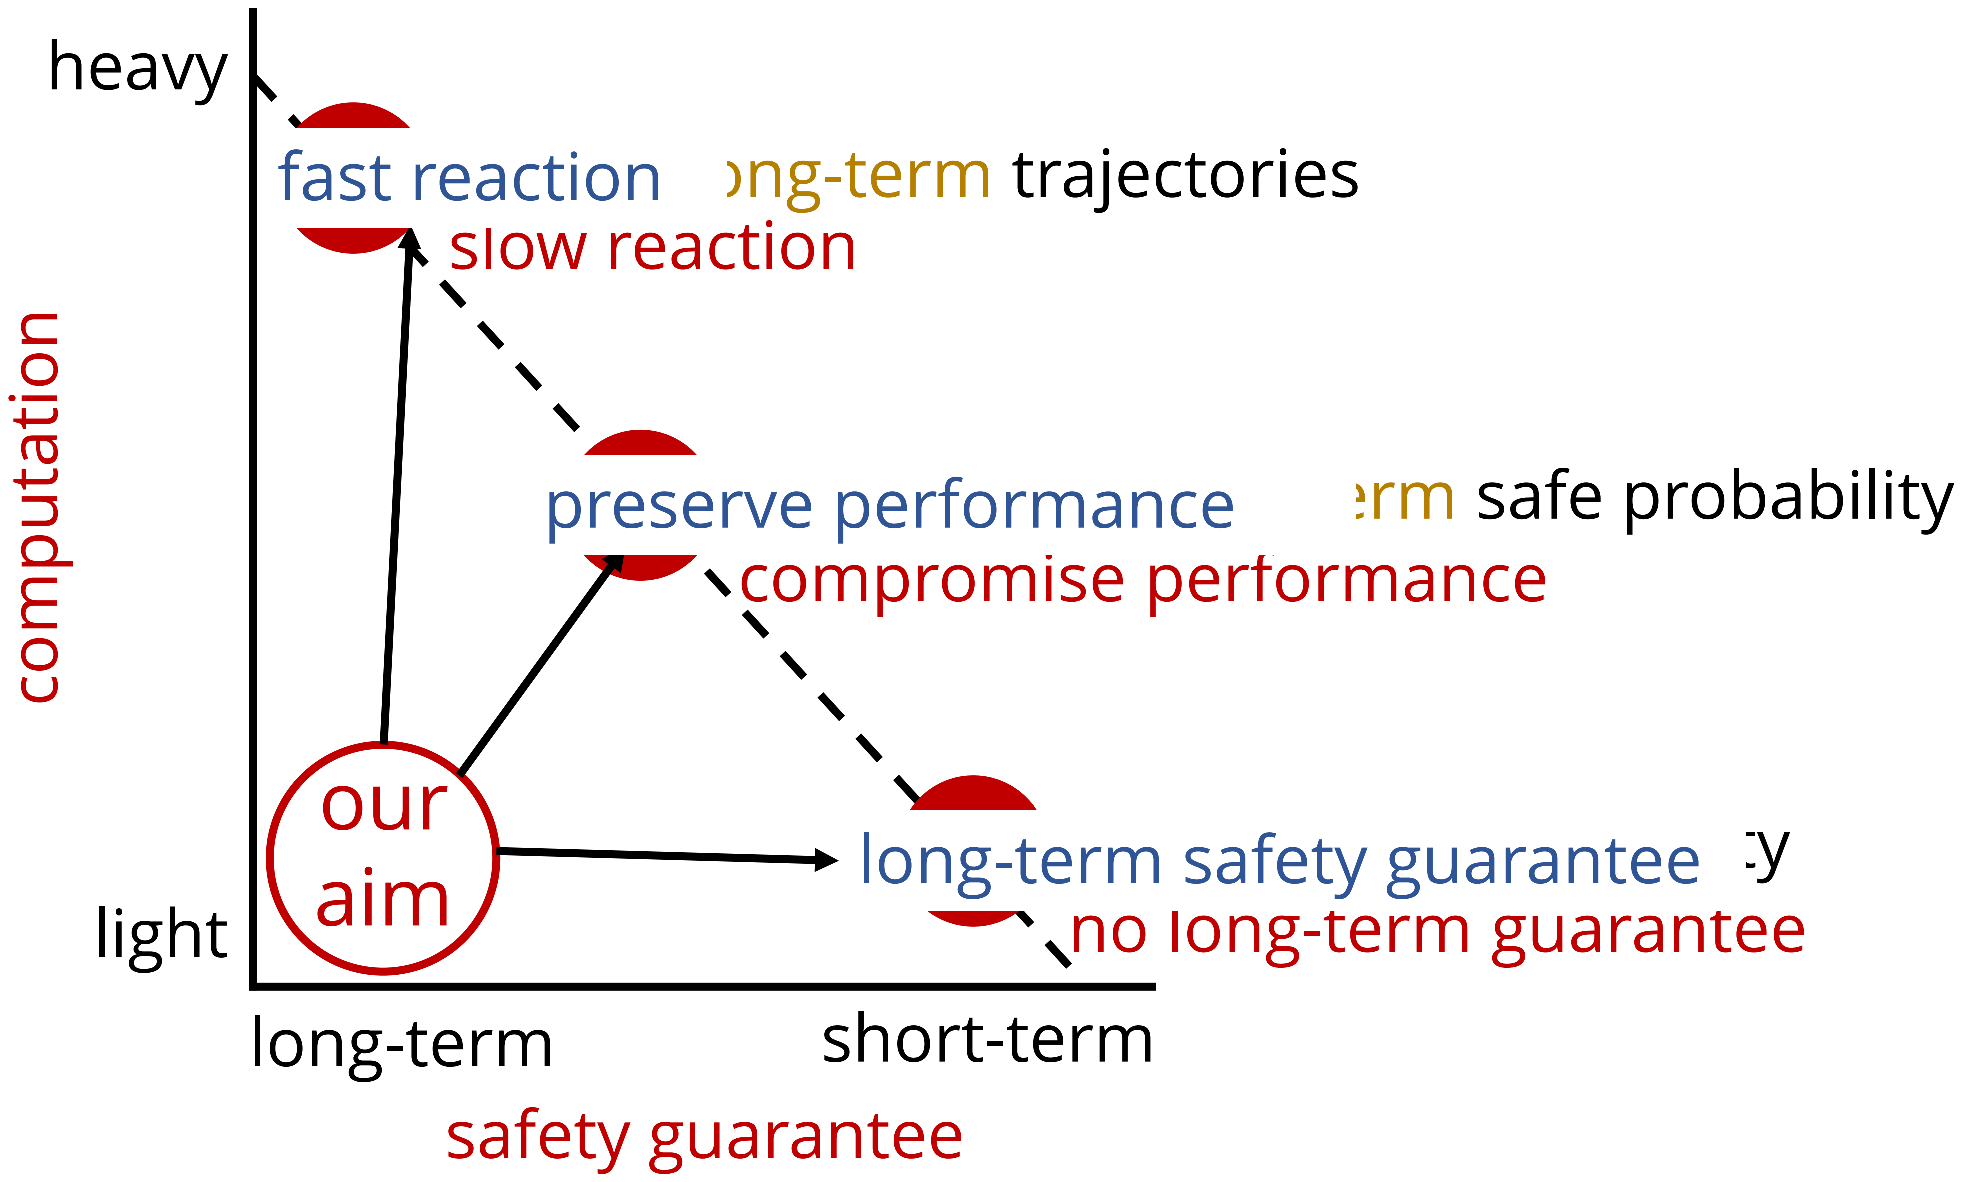
heavy
light
computation
fast reaction
ong-term trajectories
slow reaction
preserve performance
erm safe probability
compromise performance
our aim
long-term safety guarantee
y
no long-term guarantee
long-term
short-term
safety guarantee

Proposed Method: Intuitions

Imposing forward invariance on

State space

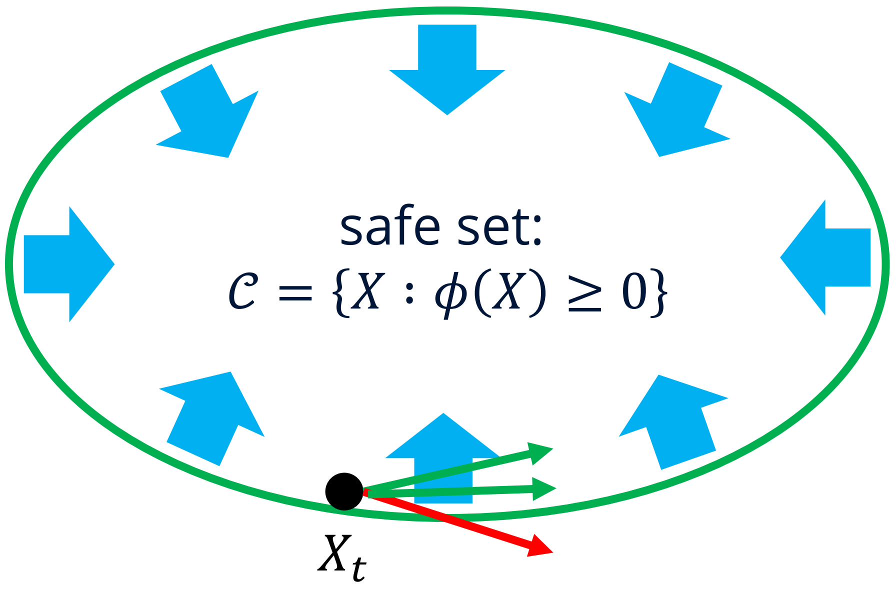

Tail probability can accumulate over time

Probability space

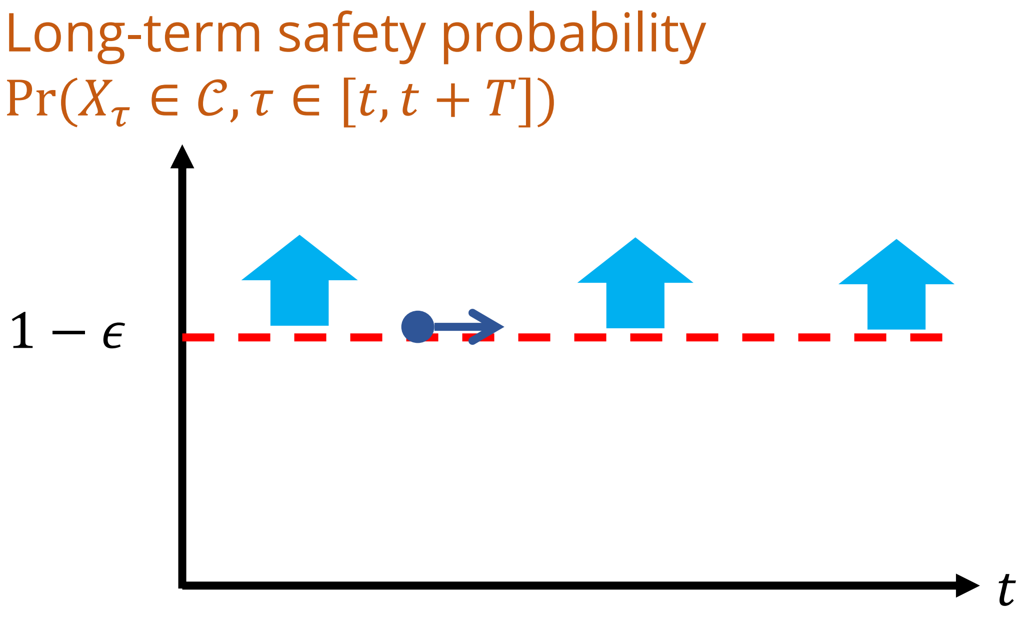

Direct control over accumulation of tail events

Proposed Method: Intuition

Barrier function based
-> Myopic evaluation

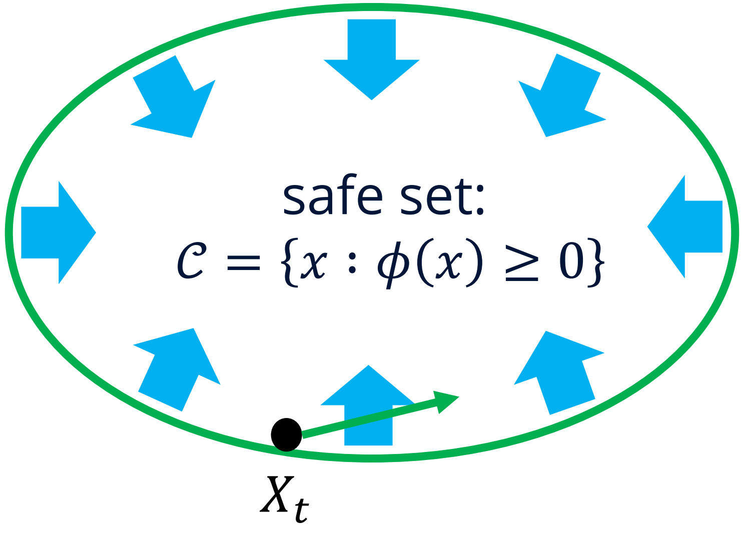

Reachability based
-> Ensures long-term safety

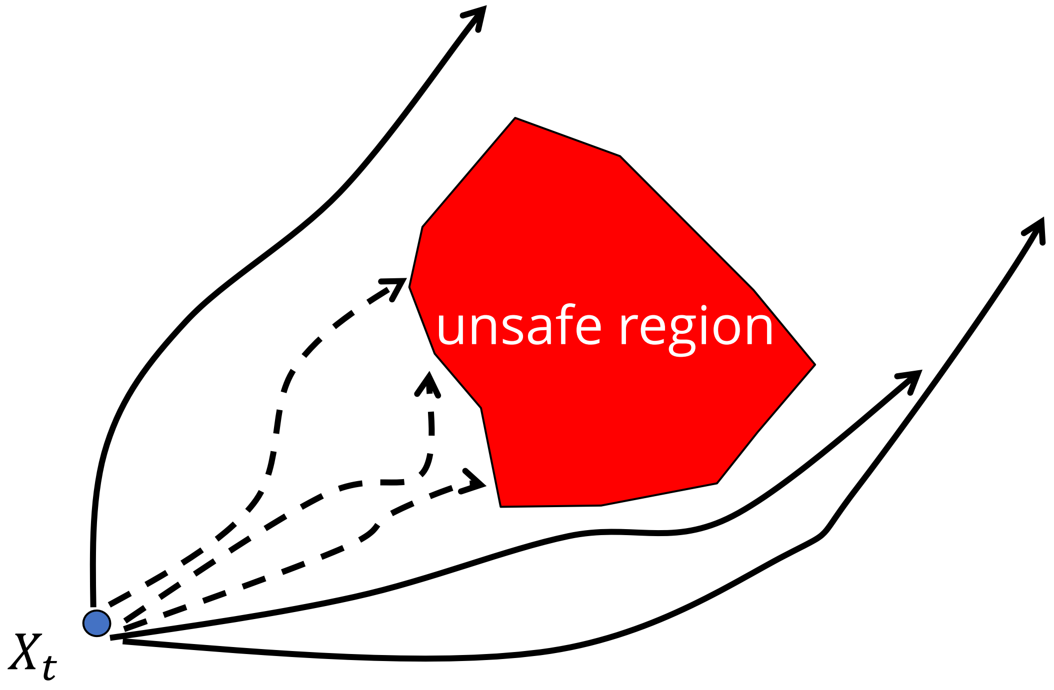

Proposed Method: Intuition

Barrier function based
-> Myopic evaluation

Reachability based
-> Ensures long-term safety

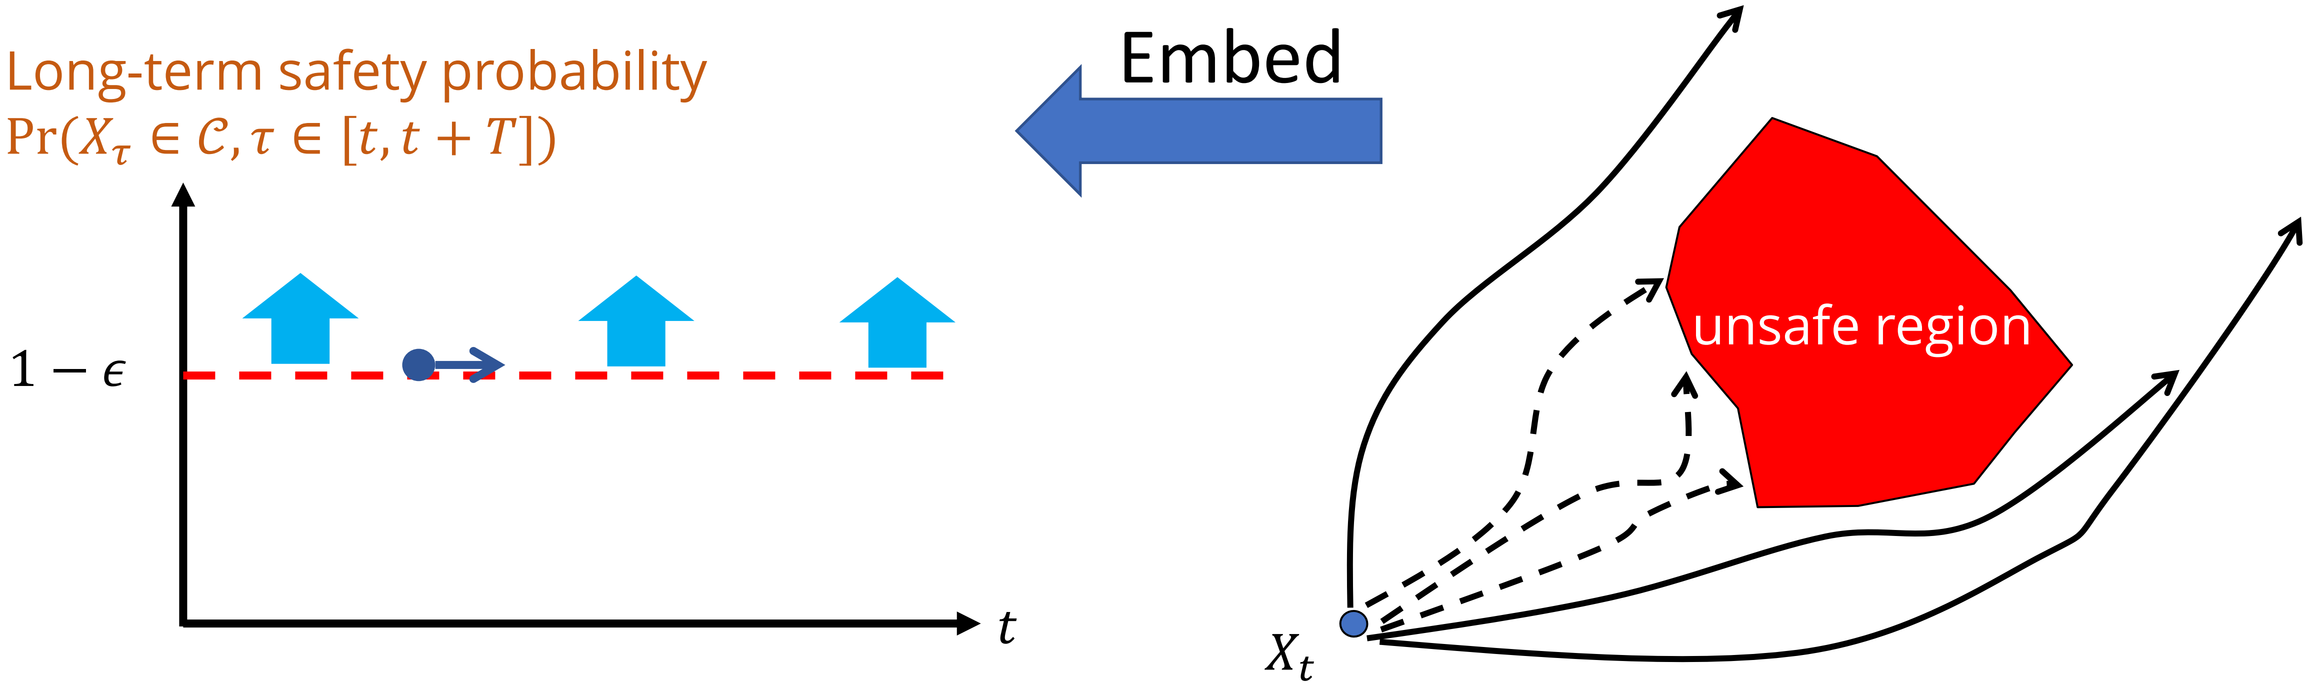

Direct control over
accumulation of tail events

Proposed Method

Long-term safety probability
$\Pr(X_\tau \in \mathcal{C}, \tau \in [t, t+T])$

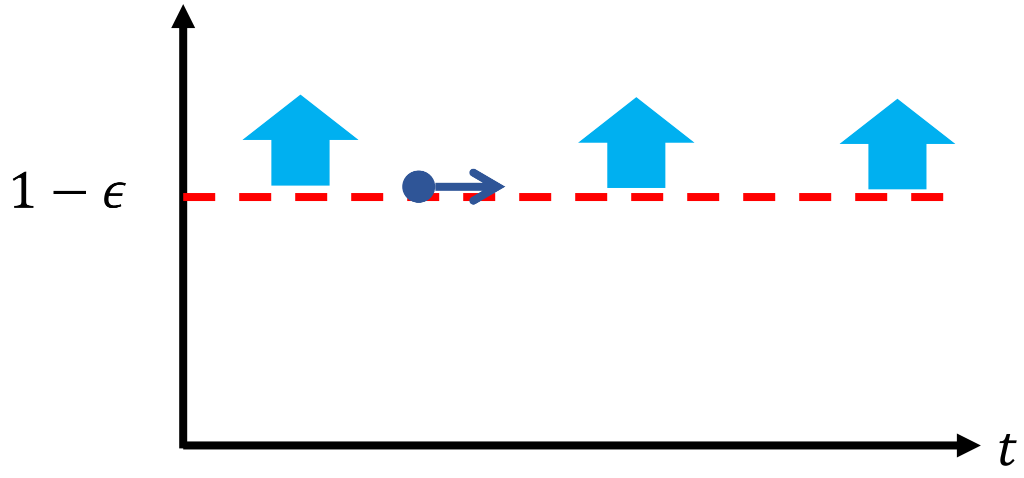

Direct control over
accumulation of tail events

$$\boldsymbol{F}(X_t) = \Pr(X_\tau \in \mathcal{C}, \tau \in [t, t+T] | X_t)$$

$$A\boldsymbol{F}(X_t) \geq -\alpha(\boldsymbol{F}(X_t) - (1-\epsilon))$$

time derivative of safety probability

desired safety probability

A: infinitesimal generator
α: monotonically increasing, concave, $\alpha(0) \leq 0$

Theoretical Guarantees

Theorem: Given

$$\boldsymbol{F}(X_0) > 1 - \epsilon,$$

if we choose the control action to satisfy

$$A\boldsymbol{F}(X_t) \geq -\alpha(\boldsymbol{F}(X_t) - (1 - \epsilon)) \text{ for } t > 0,$$

then we have

$$\Pr(X_\tau \in \mathcal{C}, \tau \in [t, t + T]) \geq 1 - \epsilon \text{ for } \forall t > 0$$

$\alpha: \mathbb{R} \to \mathbb{R}$ is a monotonically increasing concave function that satisfies $\alpha(0) \leq 0$.

Proposed Safety Condition

Step 1: compute F
Step 2: compute A, B based on X, F(X_t) and other

$$A\boldsymbol{F}(X_t) \geq -\alpha(\boldsymbol{F}(X_t) - (1-\epsilon))$$

$$dX_t = (f(X_t) + g(X_t)U_t)dt + \sigma(X_t)dW$$

Affine control

$$\mathcal{L}_f\boldsymbol{F}(X_t) + \left(\mathcal{L}_g\boldsymbol{F}(X_t)\right)\boldsymbol{U_t} + \frac{1}{2}\mathrm{tr}([\sigma(X_t)]^\intercal \mathrm{Hess}\boldsymbol{F}(X_t)[\sigma(X_t)]) \geq -\alpha(\boldsymbol{F}(X_t) - (1-\epsilon))$$

$$\left(\mathcal{L}_g\boldsymbol{F}(X_t)\right)\boldsymbol{U_t} \geq -\alpha(\boldsymbol{F}(X_t) - (1-\epsilon))\text{-}\ \mathcal{L}_f\boldsymbol{F}(X_t)\text{-}\frac{1}{2}\mathrm{tr}([\sigma(X_t)]^\intercal \mathrm{Hess}\boldsymbol{F}(X_t)[\sigma(X_t)])$$

Control constraints: A U_t >= B

A= $\left(\mathcal{L}_g\boldsymbol{F}(X_t)\right)$

linear constraints of U_t

B=$-\alpha(\boldsymbol{F}(X_t) - (1-\epsilon))$- $\mathcal{L}_f\boldsymbol{F}(X_t)$-$\frac{1}{2}\mathrm{tr}([\sigma(X_t)]^\intercal \mathrm{Hess}\boldsymbol{F}(X_t)[\sigma(X_t)])$

Proposed Safety Condition

Step 0: define f(x), g(x),
Sigma = 1, replace G = 1, w(t) -N(0,sampling time)
dW is normal (mean zero, variance sigma)
Step 1: compute F
replace their system dynamics with your in the monte carlo simulation of F
Step 2: compute A in (2), B in (3)
Step 3: replace safety condition using (1)

$$dX_t = (f(X_t) + g(X_t)U_t)dt + \sigma(X_t)dW$$

F(x)= Ax

G(X) = Bu

$$\mathcal{L}_f \boldsymbol{F}(X_t) + \left(\mathcal{L}_g \boldsymbol{F}(X_t)\right)\boldsymbol{U_t} + \frac{1}{2}\mathrm{tr}([\sigma(X_t)]^\intercal \mathrm{Hess}\boldsymbol{F}(X_t)[\sigma(X_t)]) \geq -\alpha(\boldsymbol{F}(X_t) - (1-\epsilon))$$

$$\left(\mathcal{L}_g \boldsymbol{F}(X_t)\right)\boldsymbol{U_t} \geq -\alpha(\boldsymbol{F}(X_t) - (1-\epsilon)) - \mathcal{L}_f \boldsymbol{F}(X_t) - \frac{1}{2}\mathrm{tr}([\sigma(X_t)]^\intercal \mathrm{Hess}\boldsymbol{F}(X_t)[\sigma(X_t)])$$

Control constraints: A U_t >= B – (1)

A= $\left(\mathcal{L}_g \boldsymbol{F}(X_t)\right)$ -(2)

B=$-\alpha(\boldsymbol{F}(X_t) - (1-\epsilon))$- $\mathcal{L}_f \boldsymbol{F}(X_t)$- $\frac{1}{2}\mathrm{tr}([\sigma(X_t)]^\intercal \mathrm{Hess}\boldsymbol{F}(X_t)[\sigma(X_t)])$ –(3)

linear constraints of U_t

Simulation

system dynamic: $dx_t = (2x_t + 2.5u_t)\, dt + 2dw_t$

initial state: $x_0 = 3$

safe set: $\mathcal{C} = \{x \in \mathbb{R} : x - 1 > 0\}$

nominal controller: $N(x_t) = 2.5x_t$

desired safety probability: $1 - \epsilon = 0.9$

Simulation

Proposed: $A\boldsymbol{F}(X_t) \geq -\alpha(\boldsymbol{F}(X_t) - (1-\epsilon))$

Clark: $A\phi(X_t) \geq -\alpha\phi(X_t)$

Luo et al.: $\mathbb{P}(d\phi(X_t, U_t) + \alpha\phi(X_t) \geq 0) \geq 1-\epsilon$

Ahmadi et al.: $\mathrm{CVaR}_\beta(\phi(X_{t+1})) \geq \gamma\phi(X_t)$

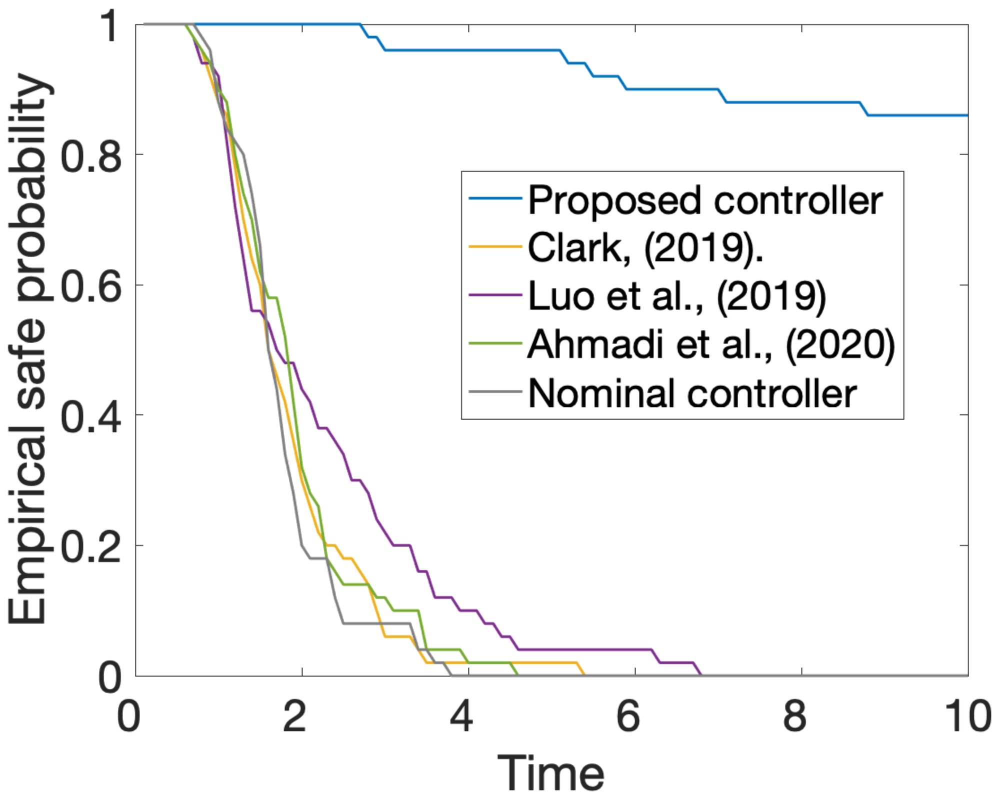

Simulation

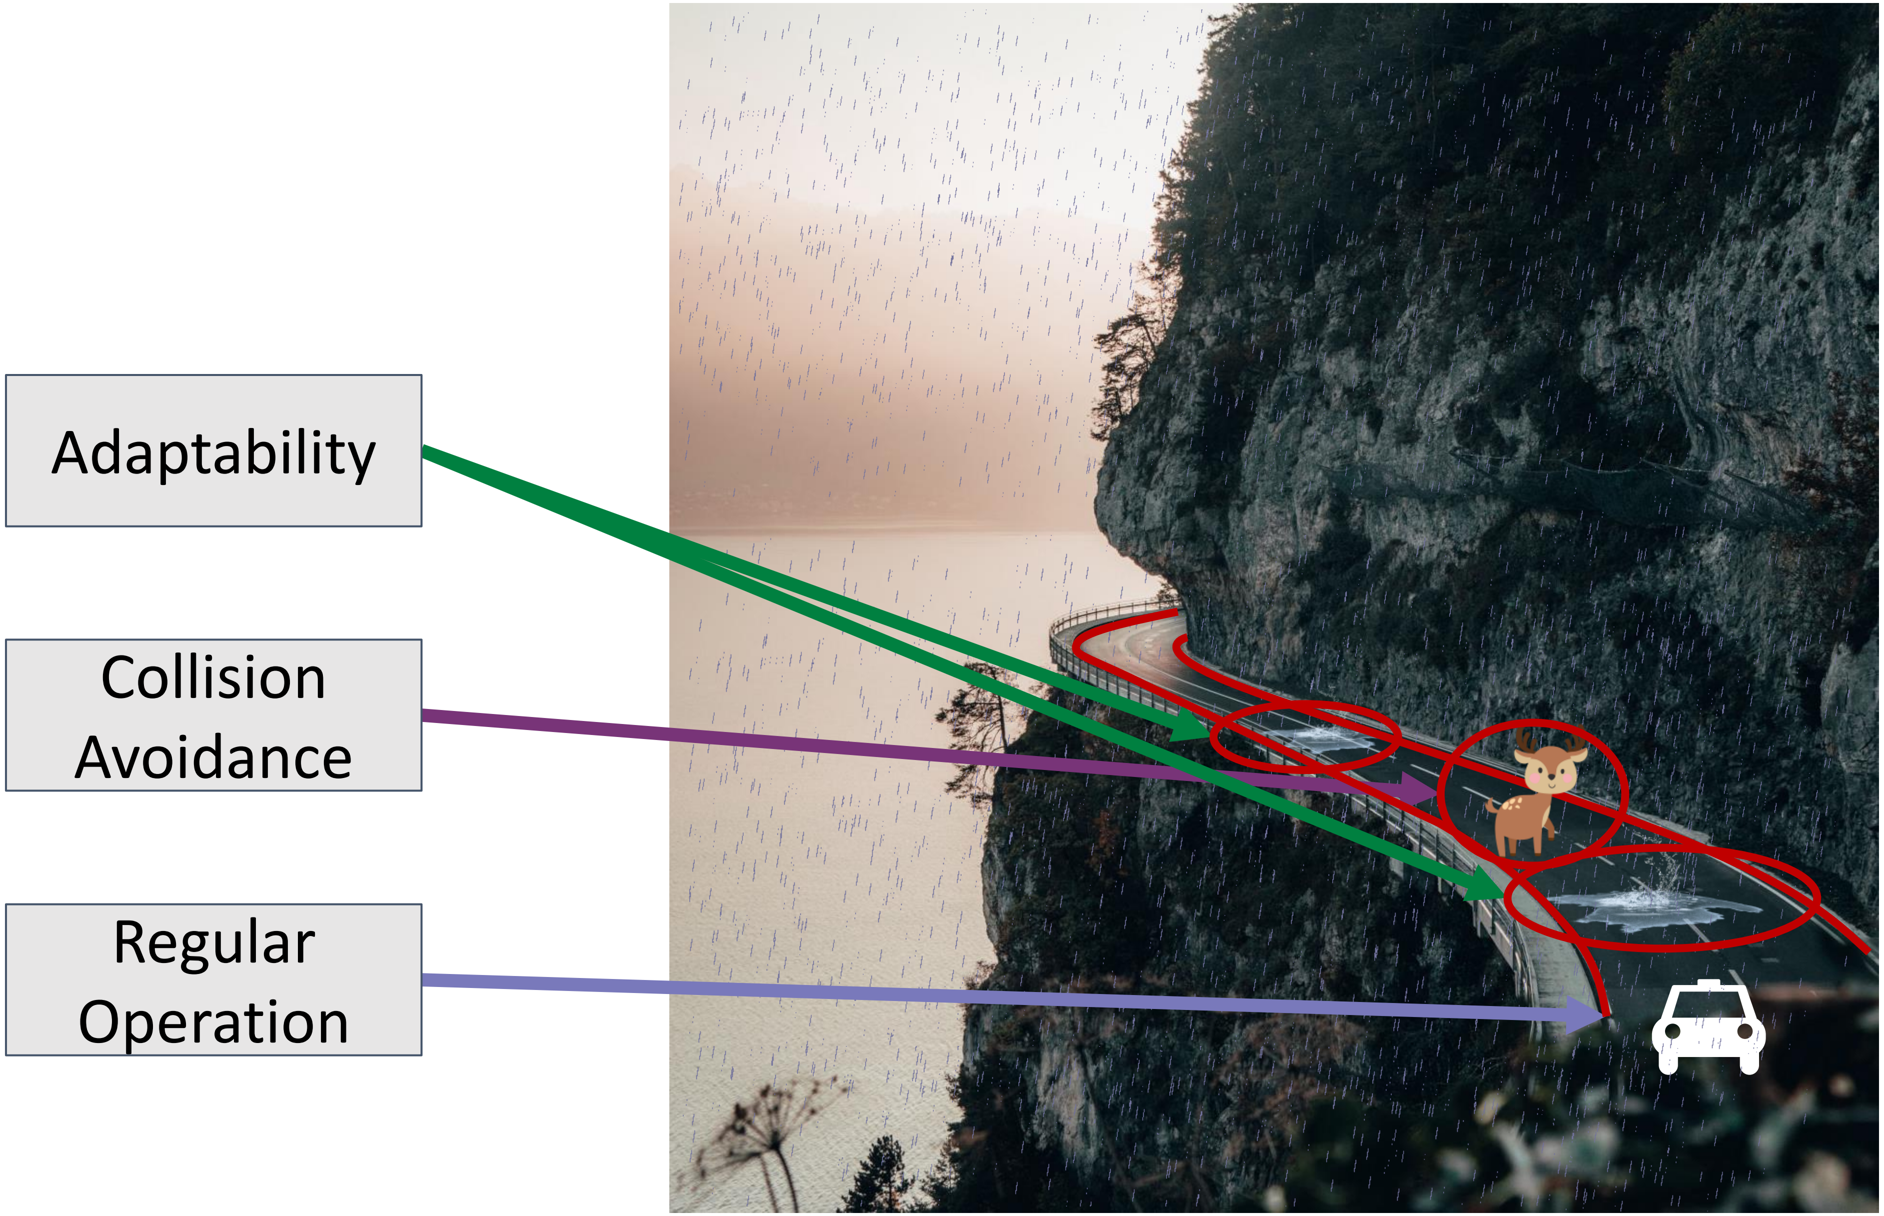

Advantage 1: Long-term Safety Guarantee

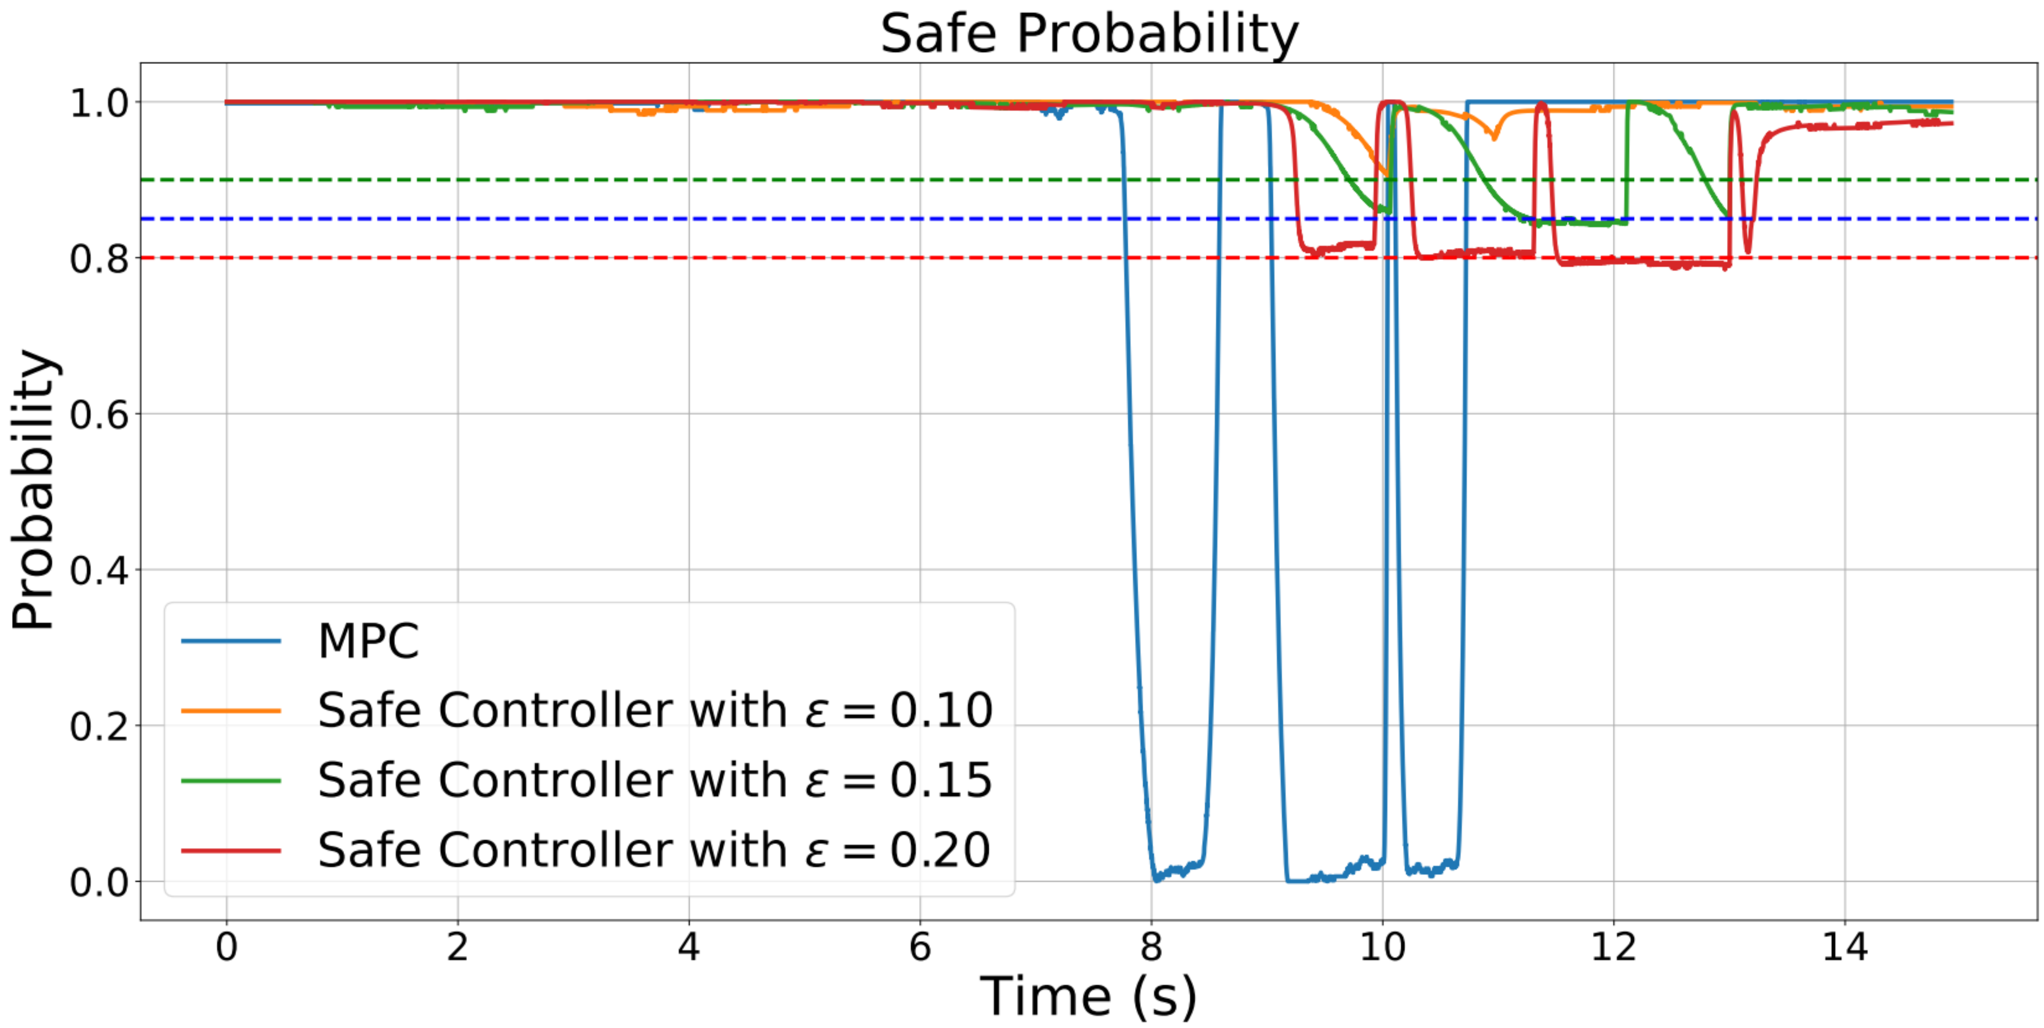

Advantage 1: Long-term Safety Guarantee (Cont'd)

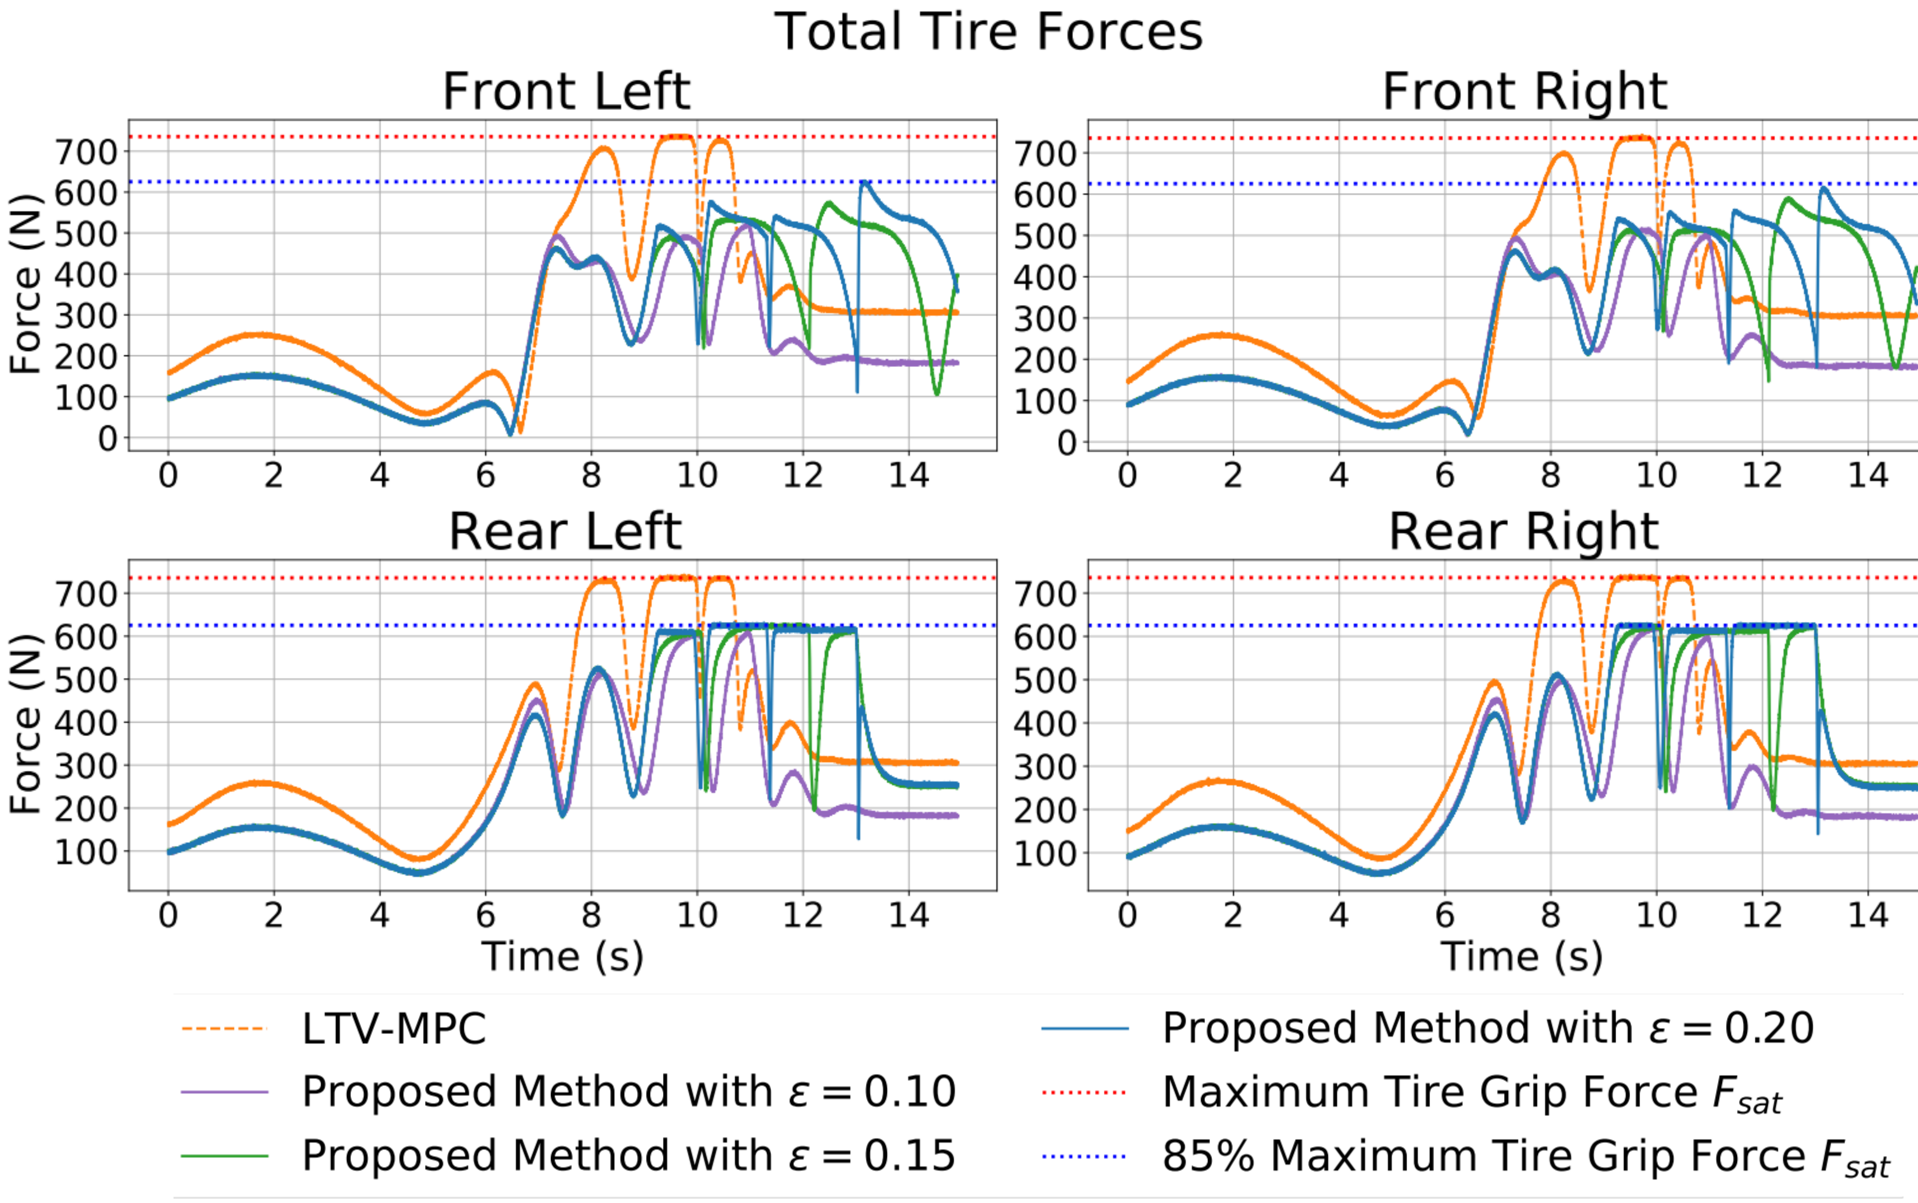

Advantage 2: Better Performance Tradeoffs

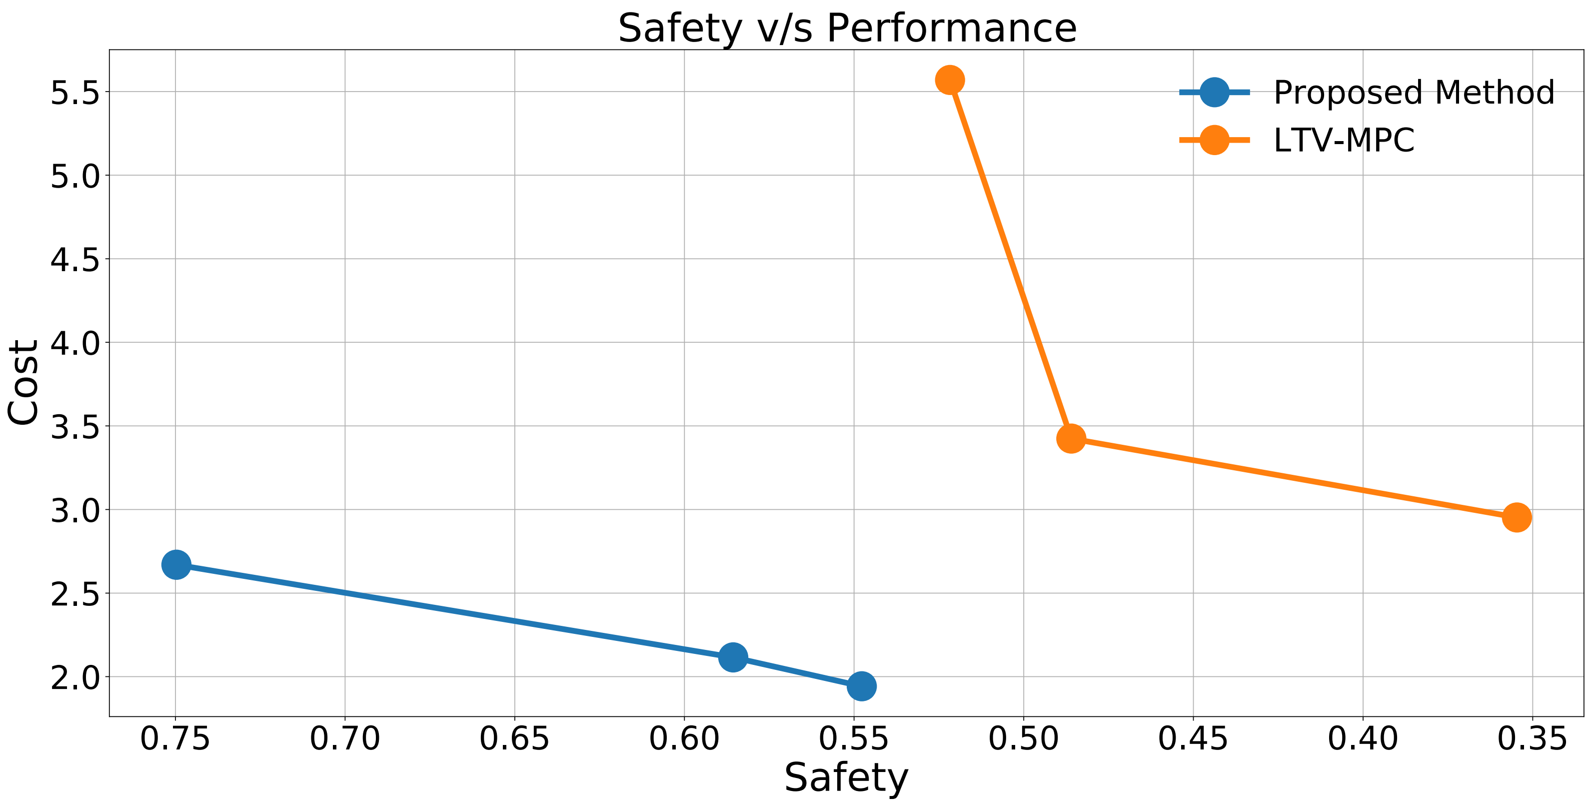

cost:
deviation from the reference trajectory

safety: satisfaction of the tire force limits

Advantage 3: Less Computation Costs

- Computation of MPC grows in $O(H^3)$
- Safety will not be compromised even with short outlook horizons

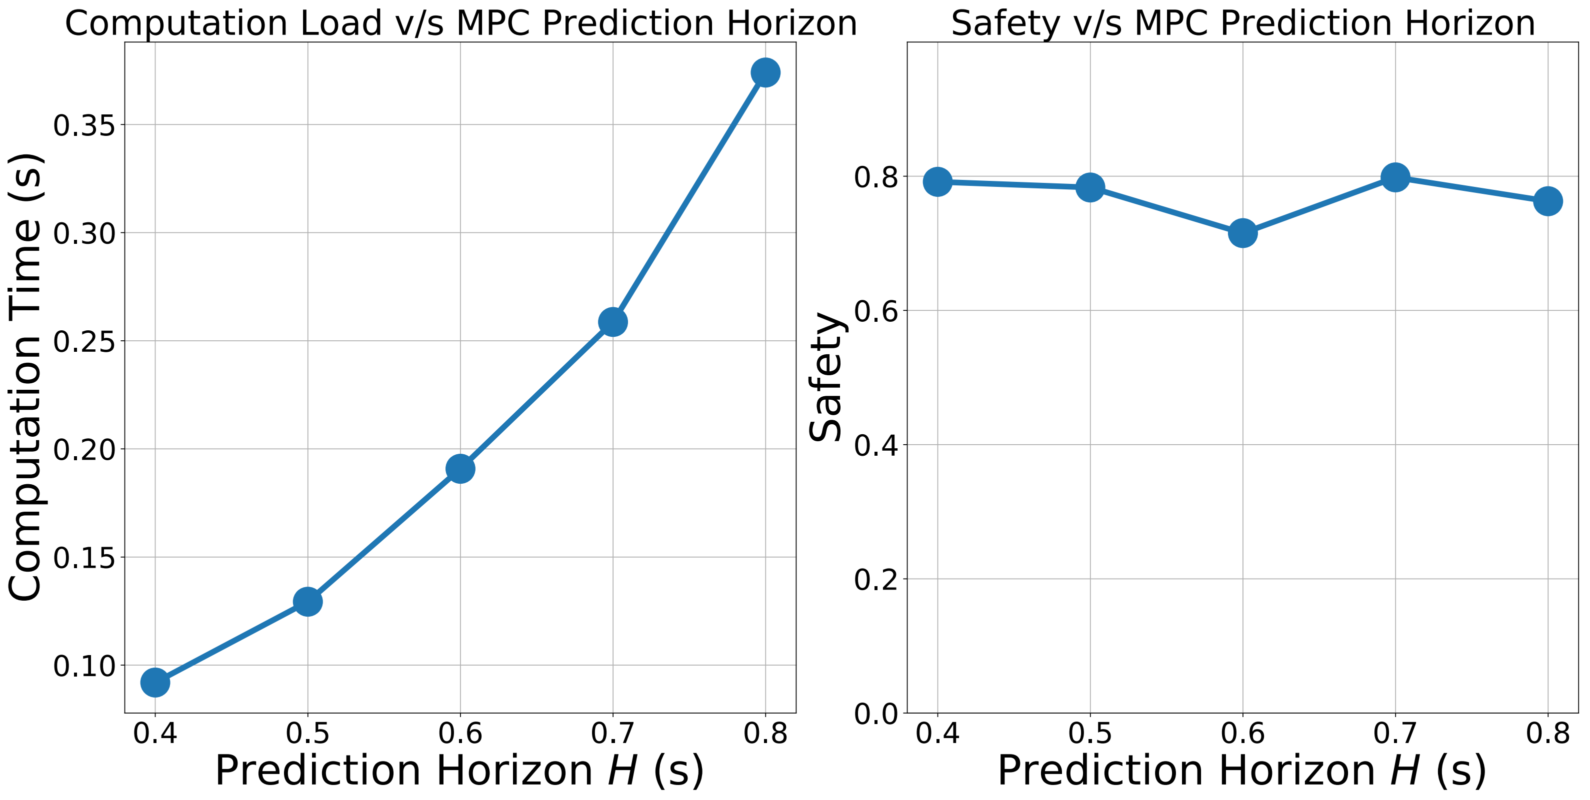

Random variables that inform safety

Characterized the distribution of:

- **Worst-case margin:** $\Phi_x(T) \coloneqq \inf\{\phi(X_t) \in \mathbb{R}: t \in [0,T], X_0 = x\}$
- **First exit time:** $\Gamma_x(\ell) \coloneqq \inf\{t \in \mathbb{R}_+: \phi(X_t) < \ell, X_0 = x\}$
- **Distance to the safe set:** $\Theta_x(T) \coloneqq \sup\{\phi(X_t) \in \mathbb{R},: t \in [0,T], X_0 = x\}$
- **Recovery time:** $\Psi_x(\ell) \coloneqq \inf\{t \in \mathbb{R}_+: \phi(X_t) \geq \ell, X_0 = x\}$

All distributions are given by the deterministic convection-diffusion equations

Theorem 1: Worst-case margin $\Phi_x(T)$

The complementary cumulative distribution function of the safety margin $\Phi_x(T)$

$$F(z,T;\ell) = P(\Phi_x(T) \geq \ell), \ell \in \mathbb{R}$$

is the solution to

$$\begin{cases} \frac{\partial F}{\partial t} = \frac{1}{2}\nabla \cdot (D\nabla F) + \mathcal{L}_{\rho - \frac{1}{2}\nabla \cdot D} F & z[1] \geq \ell, T > 0 \\ F(z,t) = 1 & z[1] \geq \ell, T > 0 \\ F(z,0) = \mathbf{1}_{\{z[1] < \ell\}}(z) & z \in \mathbb{R}^{n+1} \end{cases}$$

Safe set $\mathcal{C} = \{x: \phi(x) \geq \ell\}$

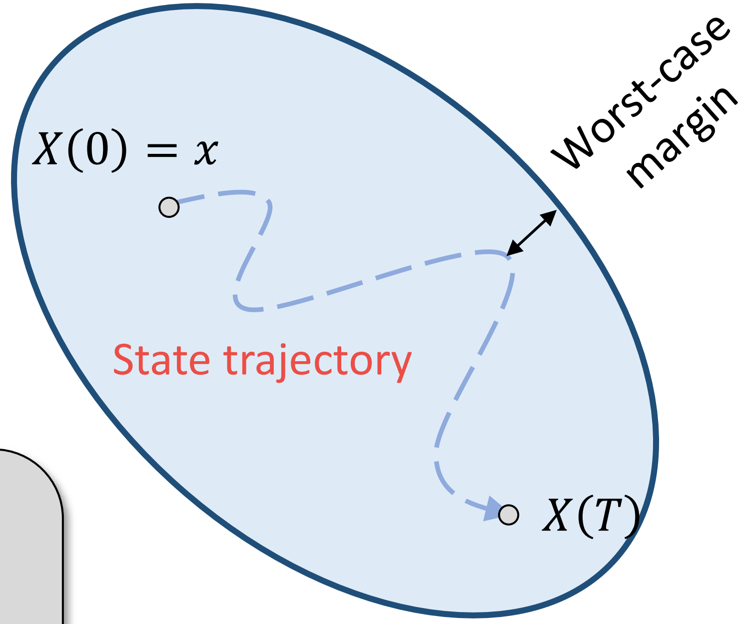

$F(z,T;0)$ = safe probability during [0,T]

$z = [\phi(x), x]$

Theorem 2: First exit time $\Gamma_x(\ell)$

Safe set $\mathcal{C} = \{x: \phi(x) \geq \ell\}$

The cumulative distribution function of the first exit time $\Gamma_x(\ell)$

$$G(z, t; \ell) = P(\Gamma_x(\ell) \leq t)$$

is the solution to

$$\begin{cases} \dfrac{\partial G}{\partial t} = \dfrac{1}{2}\nabla \cdot (D\nabla G) + \mathcal{L}_{\rho - \frac{1}{2}\nabla \cdot D} G & z[1] \geq \ell, t > 0 \\ G(z, t) = 1 & z[1] < \ell, t > 0 \\ G(z, 0) = \mathbb{1}_{\{z[1] < \ell\}}(z) & z \in \mathbb{R}^{n+1} \end{cases}$$

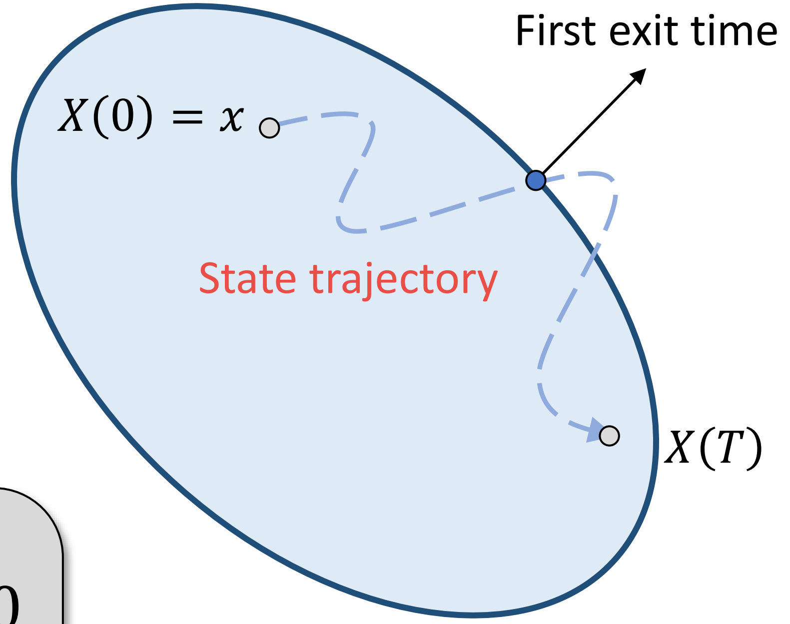

$1 - G(z, T; 0)$ = safe probability during [0,T]

$z = [\phi(x), x]$

Theorem 3: Distance to the safe set $\Theta_x(T)$

Safe set $\mathcal{C} = \{x: \phi(x) \geq \ell\}$

The **cumulative distribution function of the** distance to the safe set $\Theta_x(T)$

$$Q(z, T; -\ell) = P(\Theta_x(T) \geq \ell), \ell \in \mathbb{R}$$

is the solution to

$$\begin{cases} \dfrac{\partial Q}{\partial t} = \dfrac{1}{2}\nabla \cdot (D\nabla Q) + \mathcal{L}_{\rho - \frac{1}{2}\nabla \cdot D} Q & z[1] < -\ell, T > 0 \\ Q(z, T; -\ell) = 0 & z[1] \geq -\ell, T > 0 \\ Q(z, 0; -\ell) = \mathbb{1}_{\{z[1] < \ell\}}(z) & z \in \mathbb{R}^{n+1} \end{cases}$$

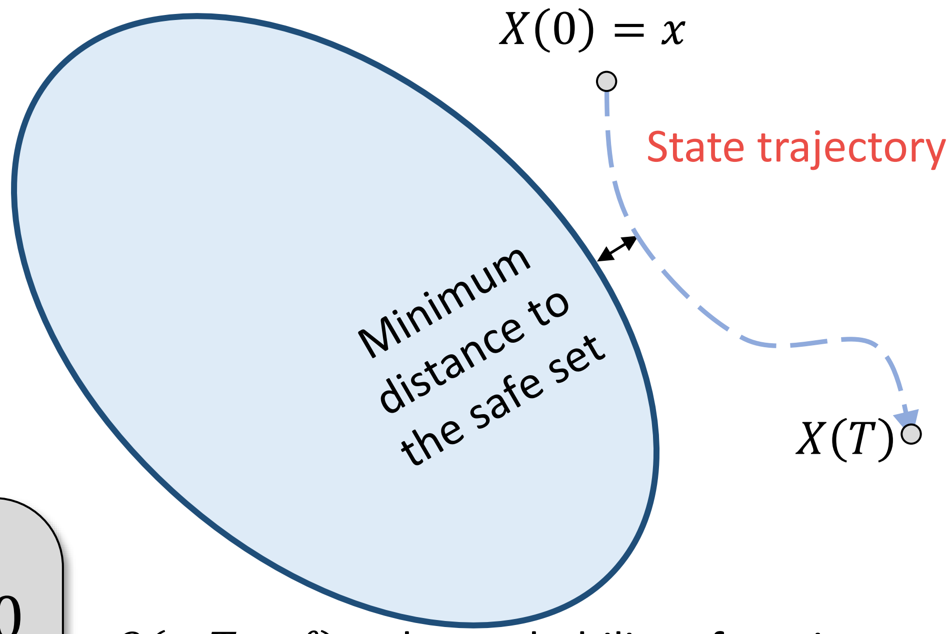

$Q(z, T; -\ell)$ = the probability of getting within ℓ distance to C during [0,T]

$z = [\phi(x), x]$

Theorem 4: First recovery time $\Psi_x(\ell)$

Let $D = \zeta\zeta^T$,the cumulative distribution function of recovery time $\Psi_x(\ell)$

$$N(z, T; \ell) = P(\Psi_x(\ell) \le t)$$

is the solution to

$$\begin{cases} \dfrac{\partial N}{\partial t} = \dfrac{1}{2}\nabla \cdot (D\nabla N) + \mathcal{L}_{\rho - \frac{1}{2}\nabla \cdot D} N & z[1] < \ell, t > 0 \\ N(z,t) = 1 & z[1] \ge \ell, t > 0 \\ N(z, 0) = \mathbb{1}_{\{z[1] \ge \ell\}}(z) & z \in \mathbb{R}^{n+1} \end{cases}$$

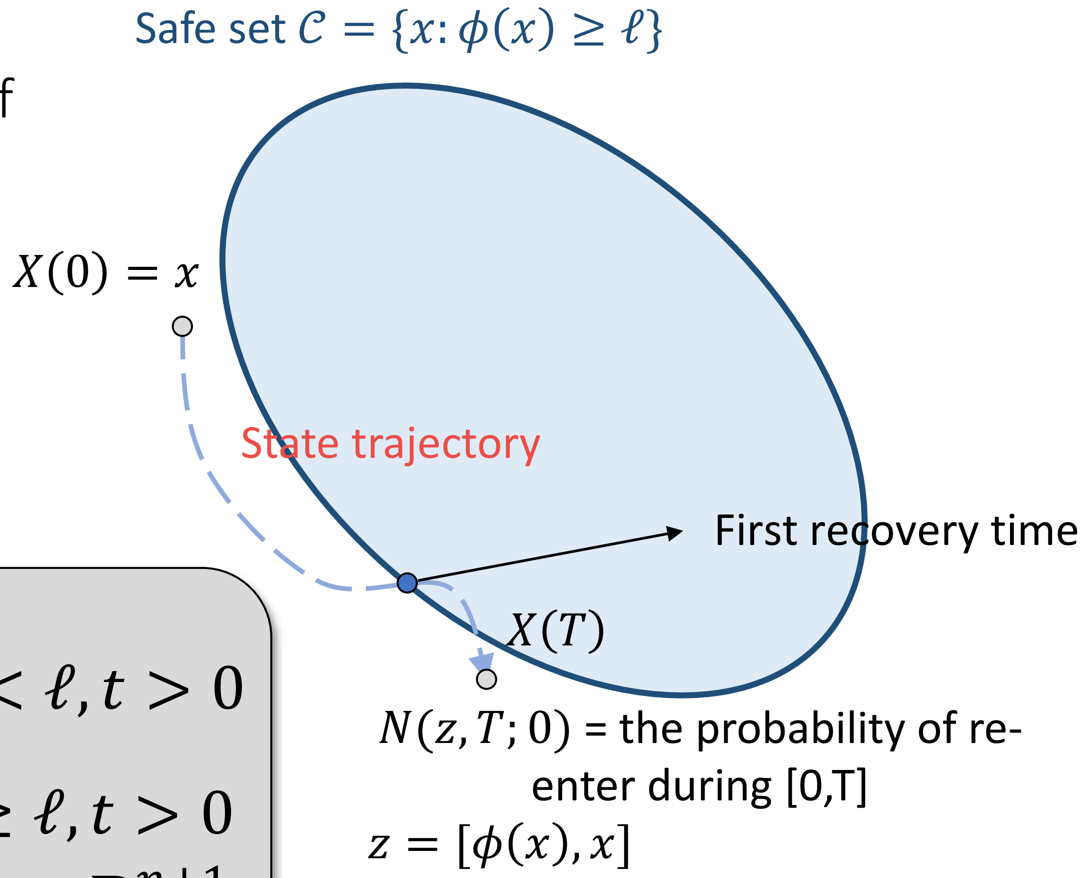

Example use case

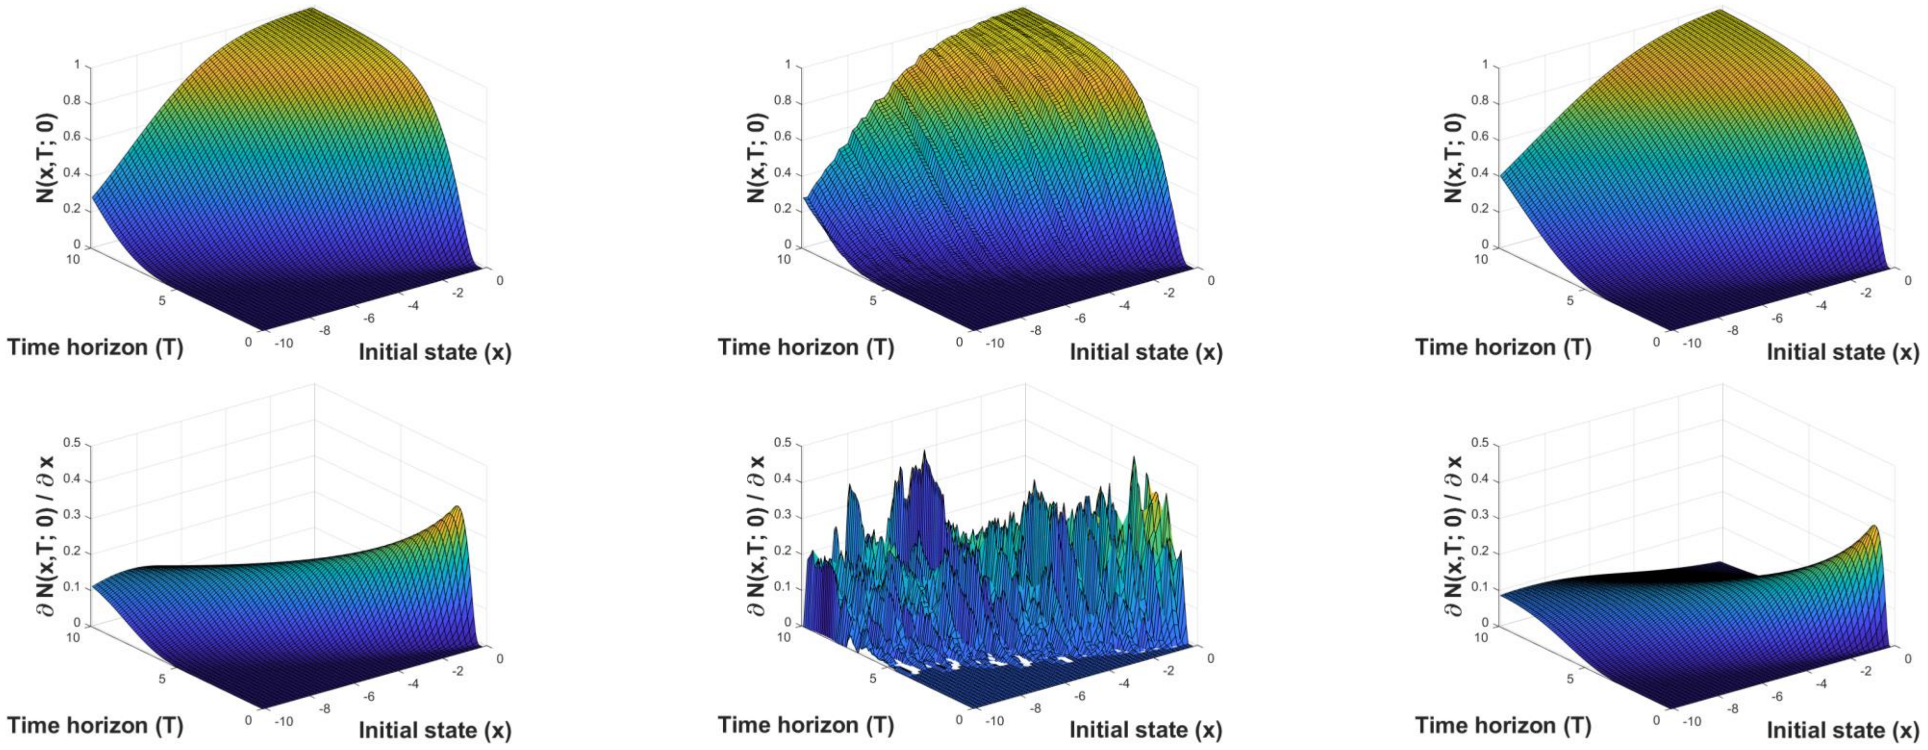

Ground truth　　Monte Carlo　　PDE solver

Today's talk

科学

Neuroscience
Biomolecular control…

工学

Constraints vs robust performance in human

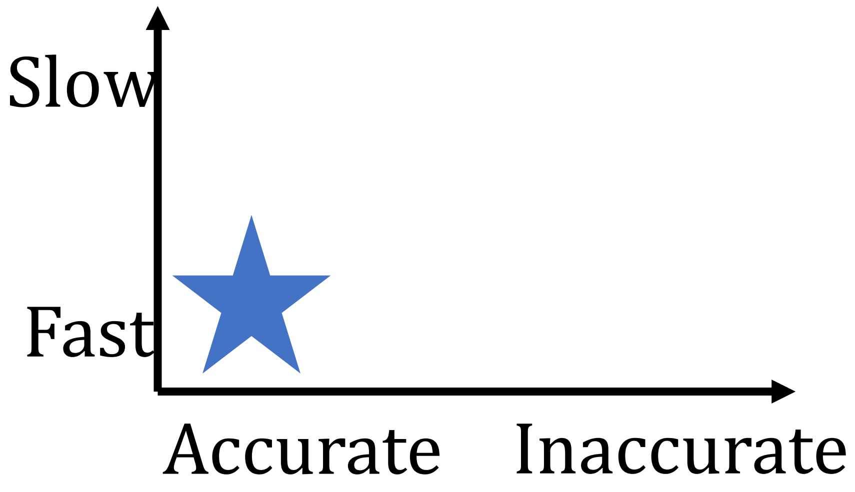

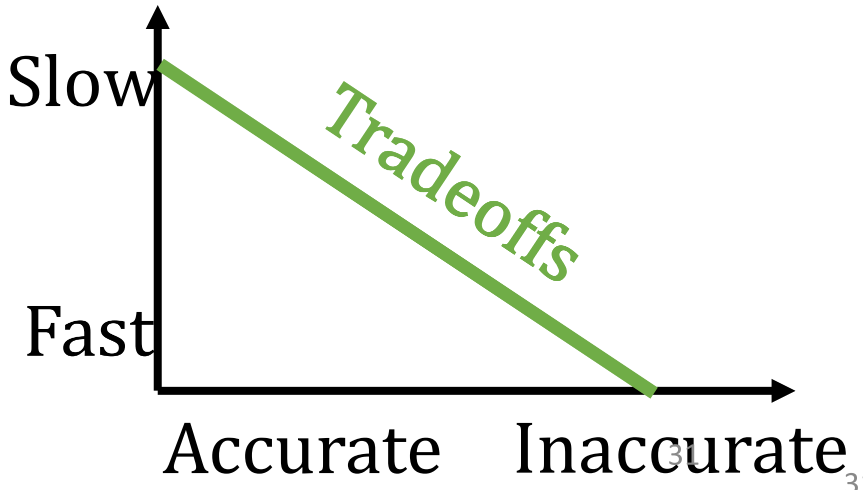

31

Constraints vs robust performance in human

Task 1: Compensate for the head motion

Kyoto

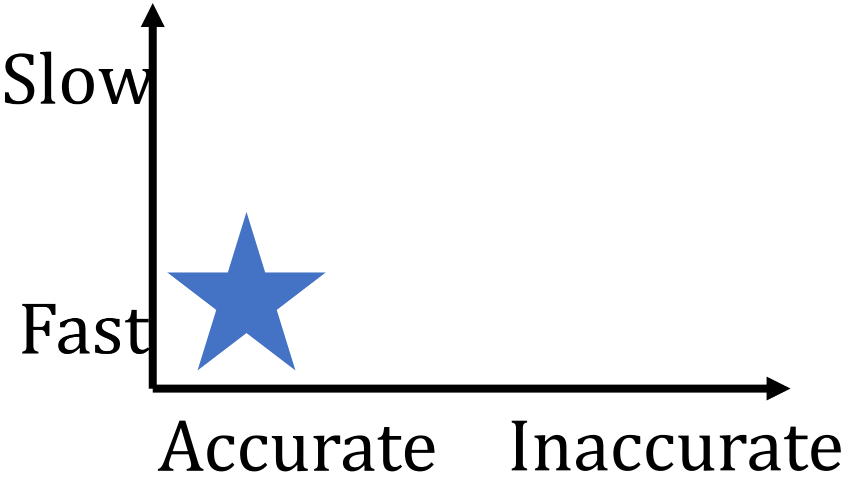

Task 2: Tracking a moving object

Kyoto

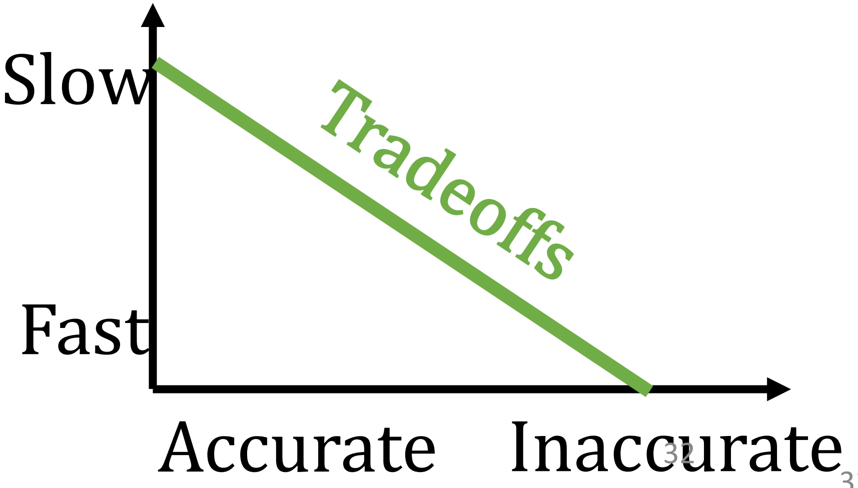

Constraints vs robust performance in human

Sensorimotor control

Neurophysiology

Biking,
eye movement, etc.

↑

A feedback loop
(e.g. VOR, reflex)

↑

Hardware
(neurons, muscles)

↑

Biological resources

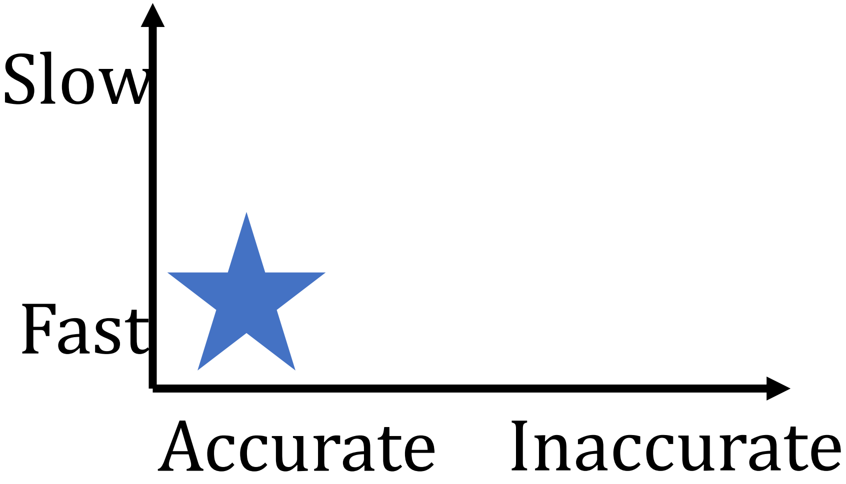

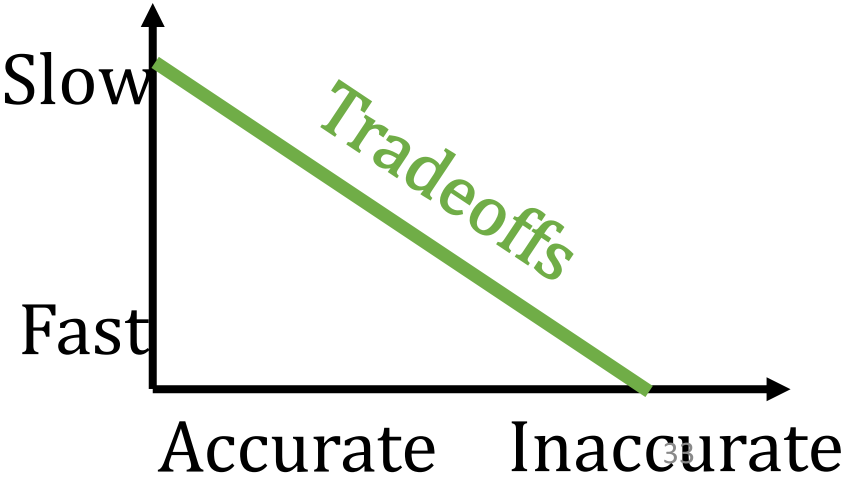

Component speed-accuracy tradeoffs

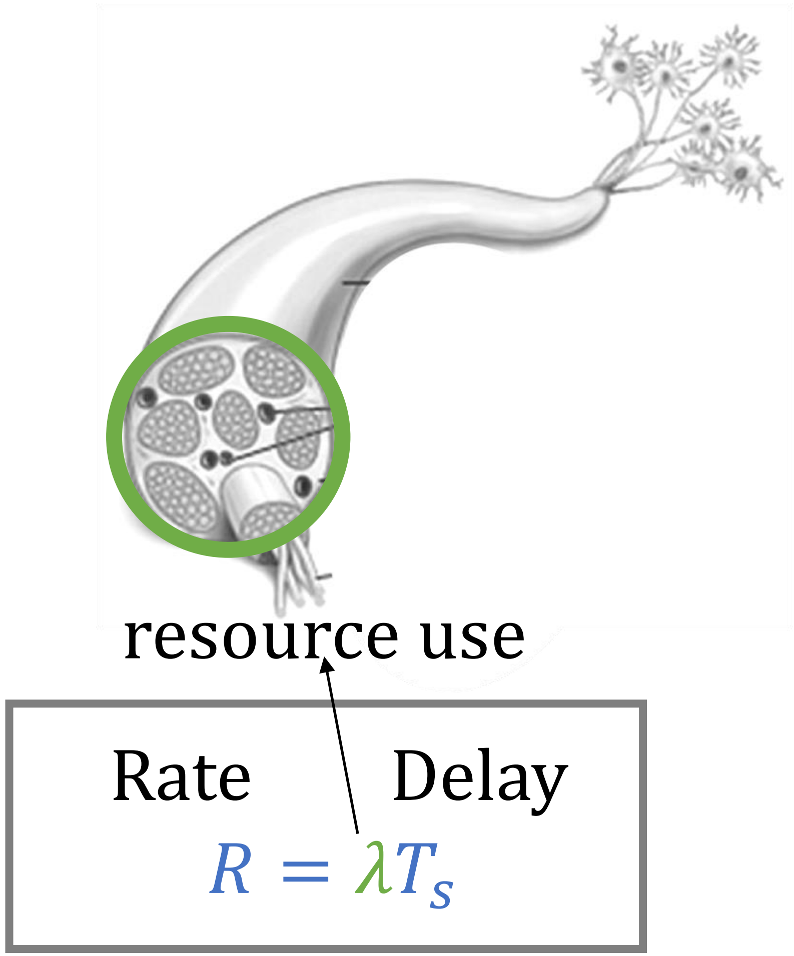

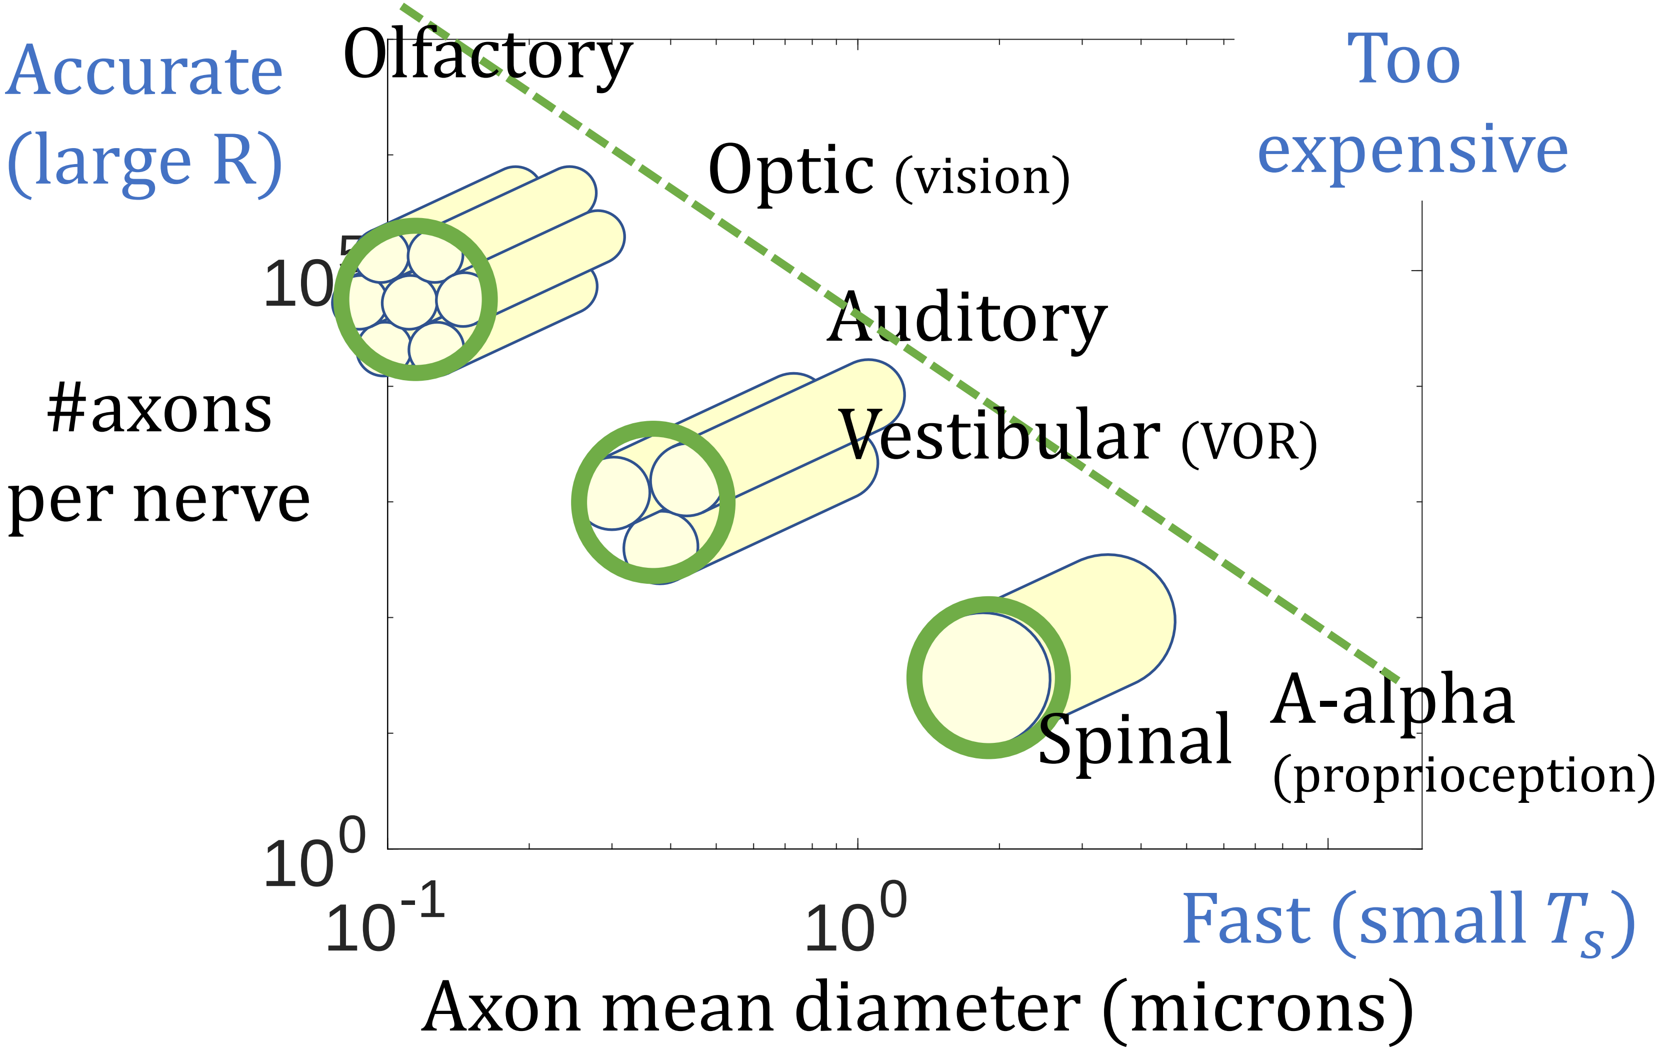

Diversity in axon radius

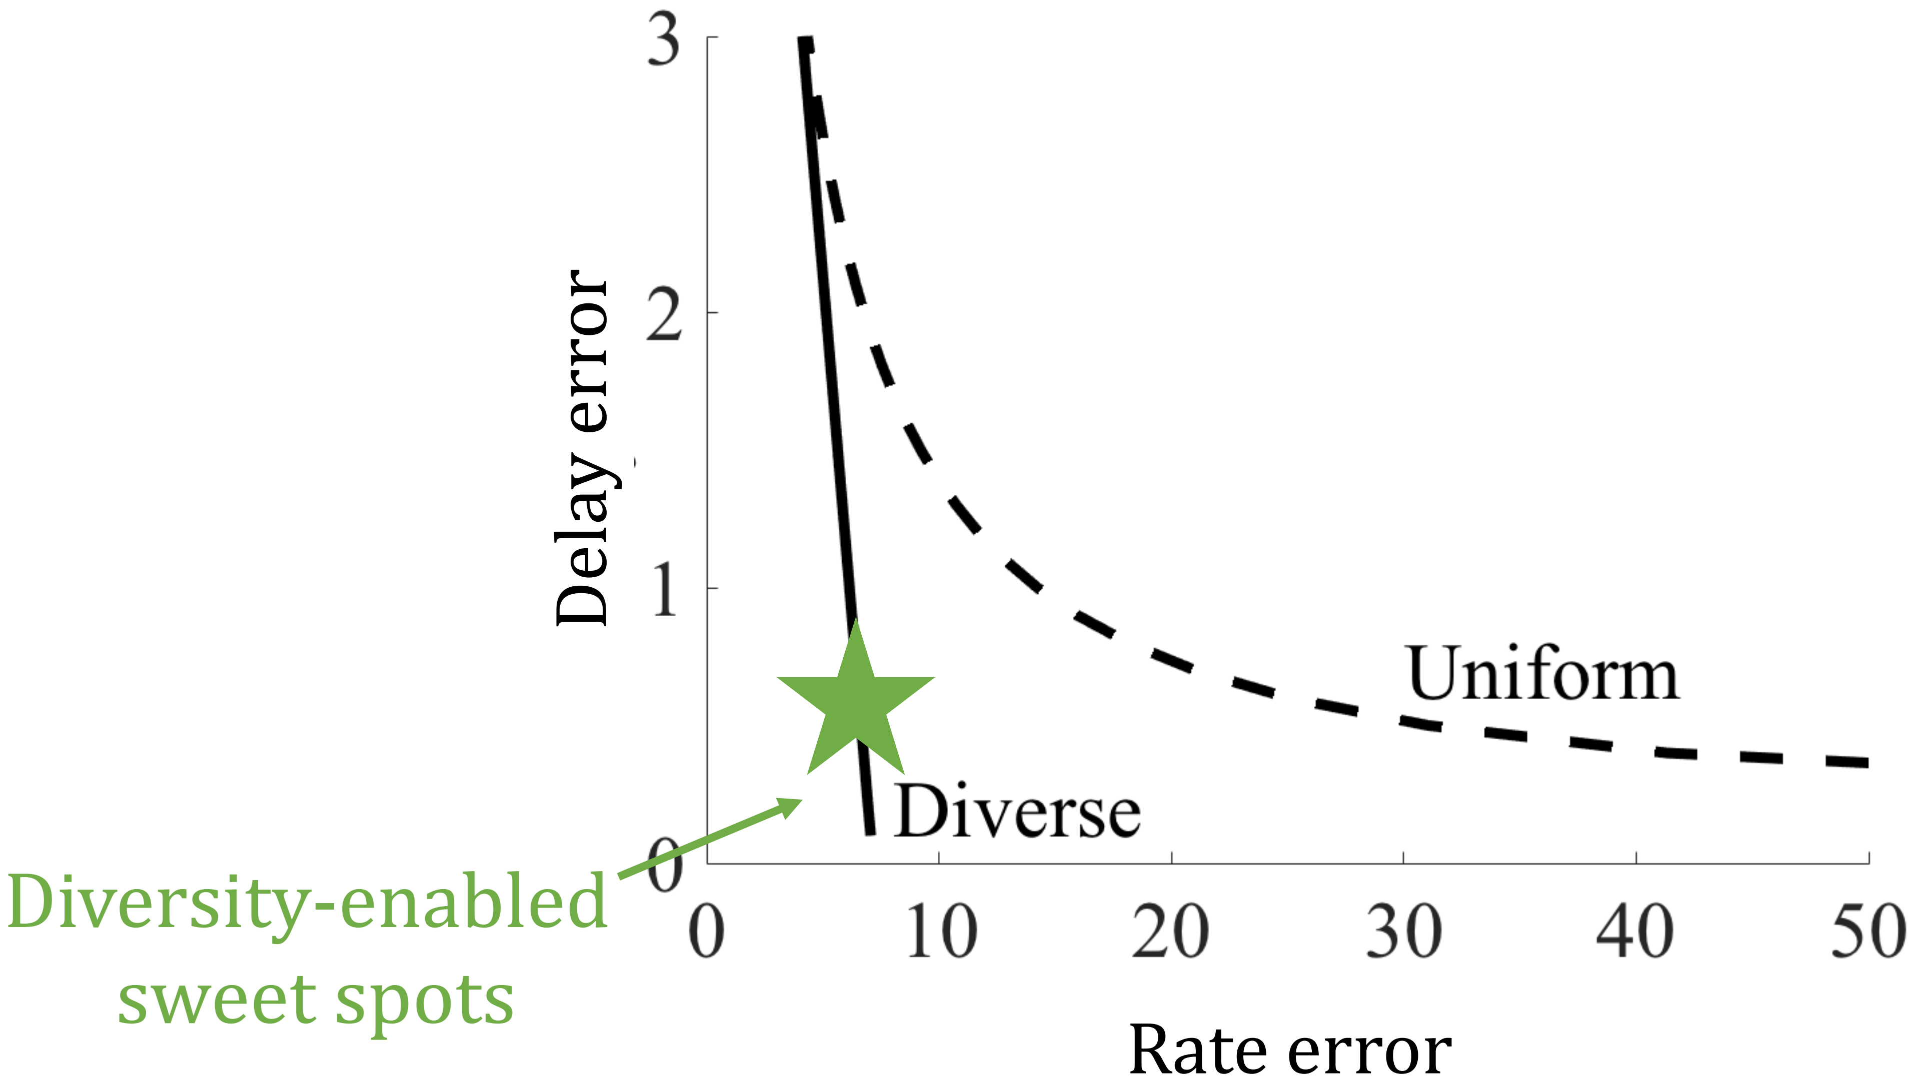

Diversity in muscle compositions

Different types of muscle

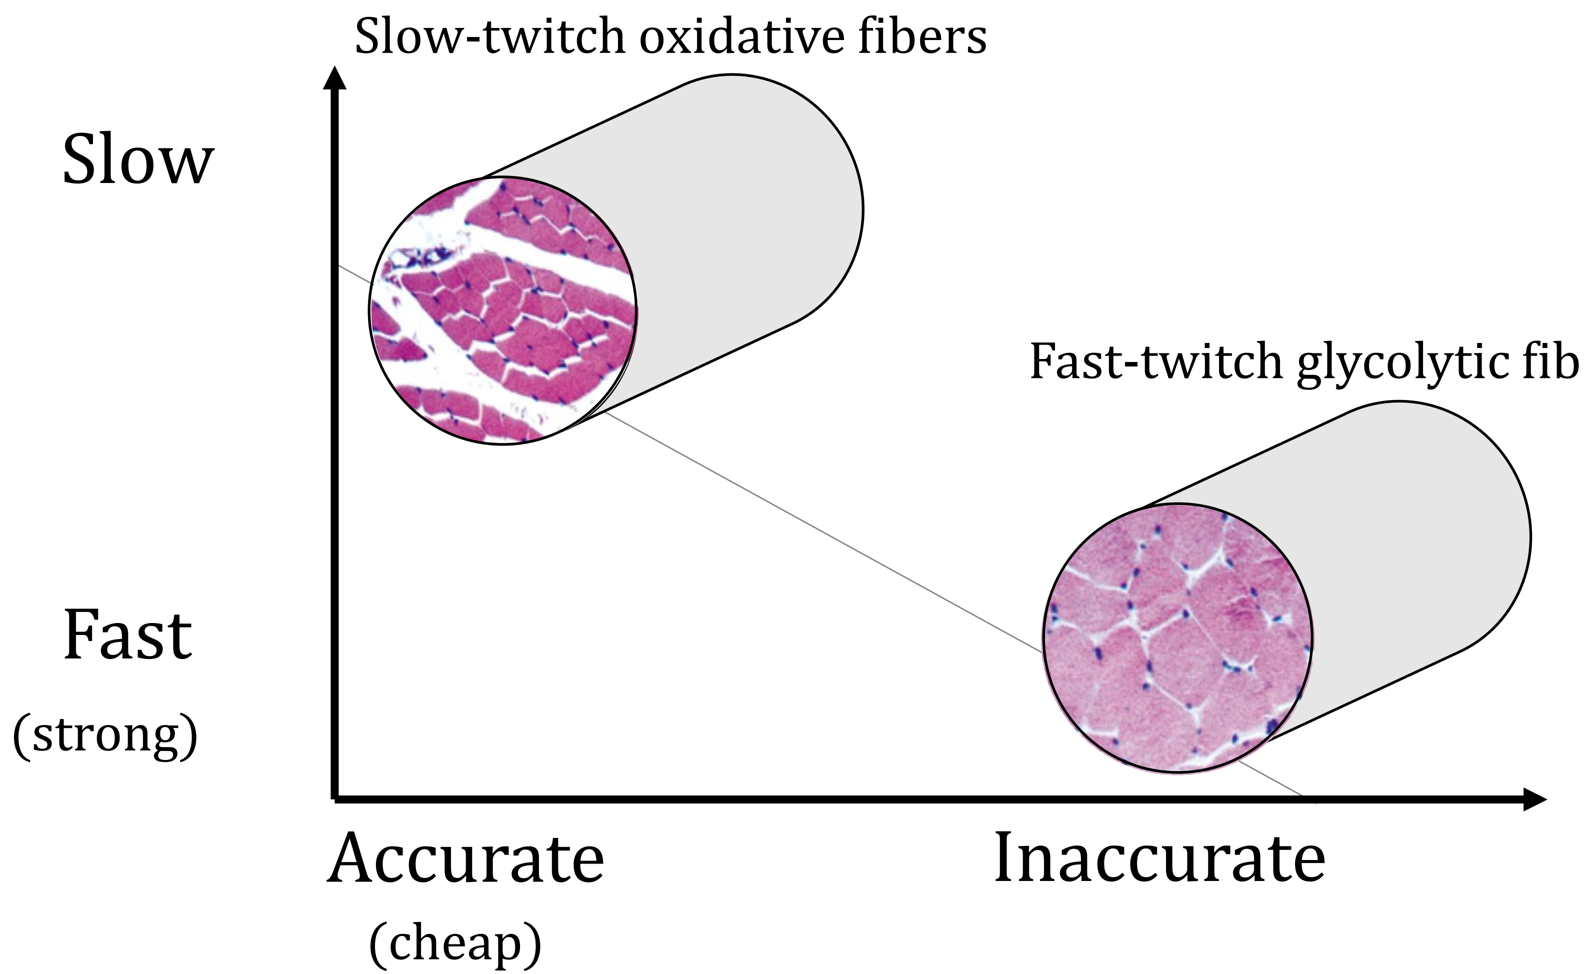

Diversity in muscle compositions

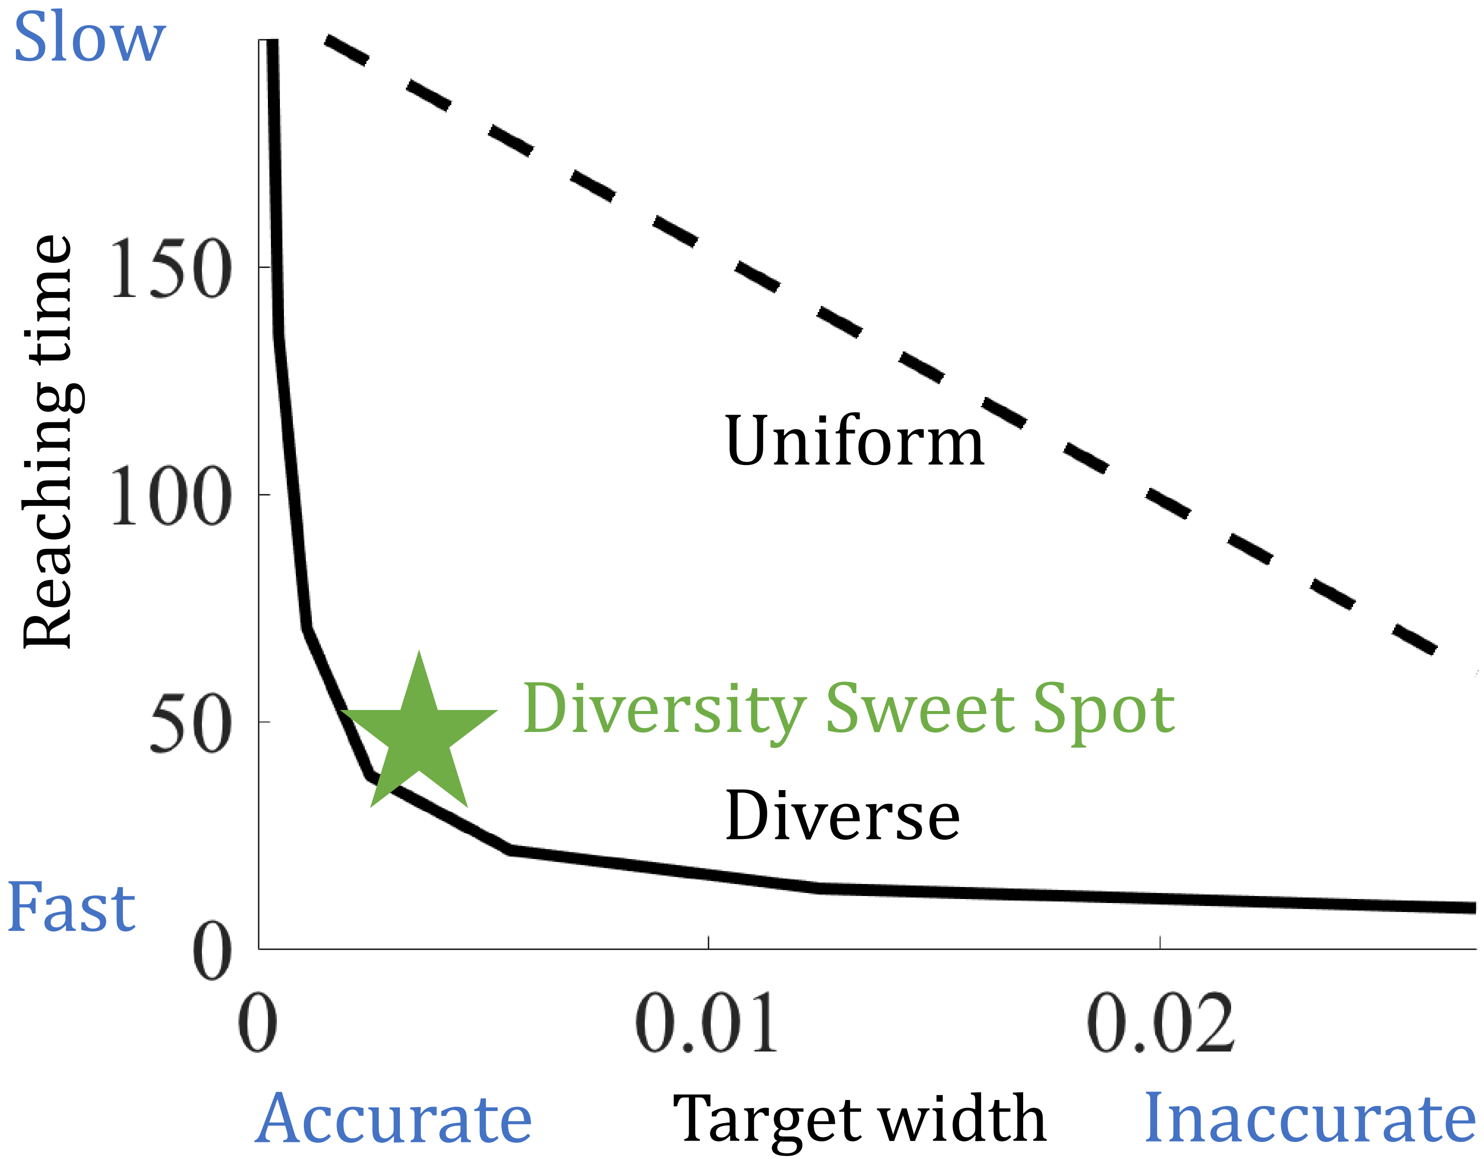

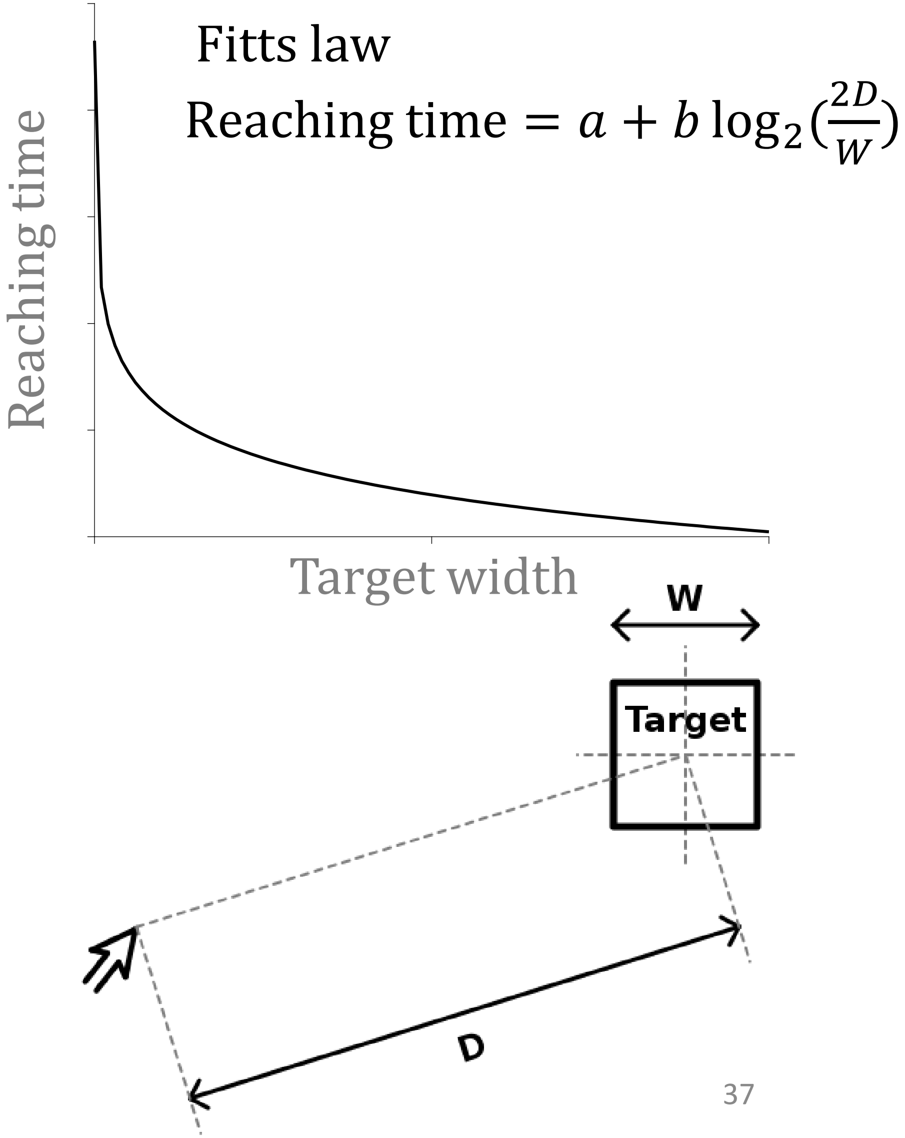

Diversity-enabled sweet spots in biology

Efficient learning

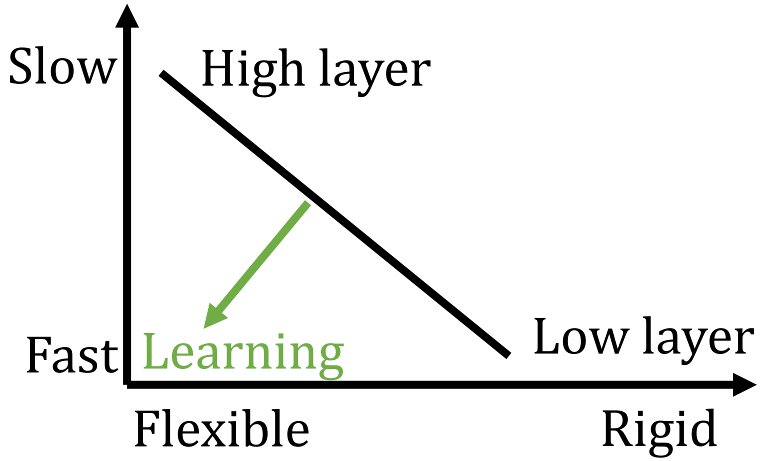

Biking

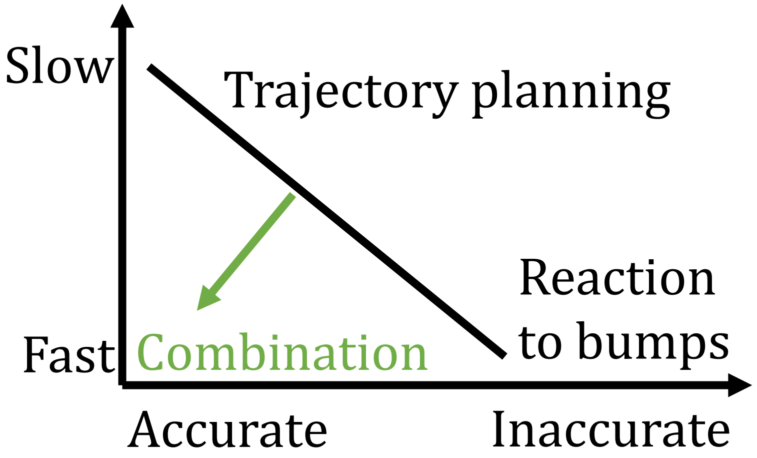

Walk/balance/react

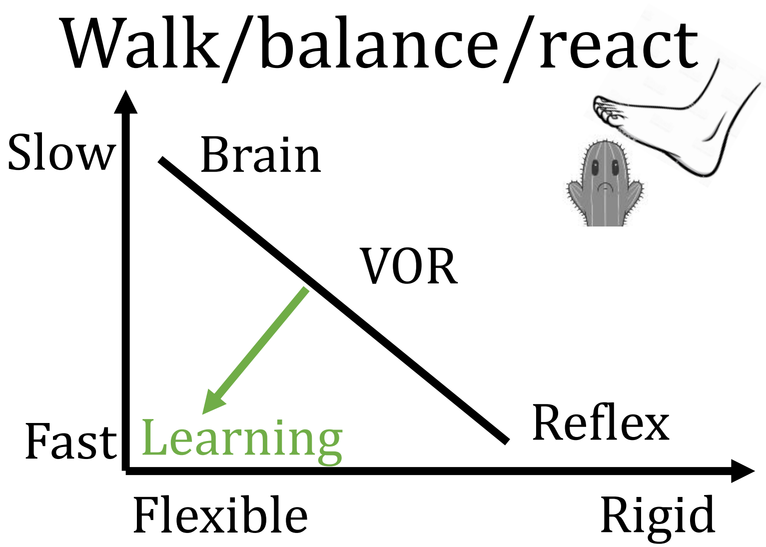

Immune system

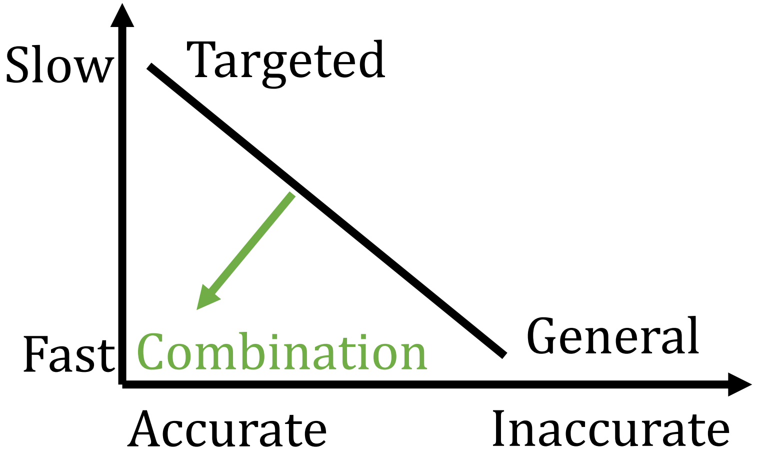

Logarithmic laws in nature

- Fitts law
- Weber-Fechner Law
- Ricco Law
- Accot-Zhai
- Power law of practice

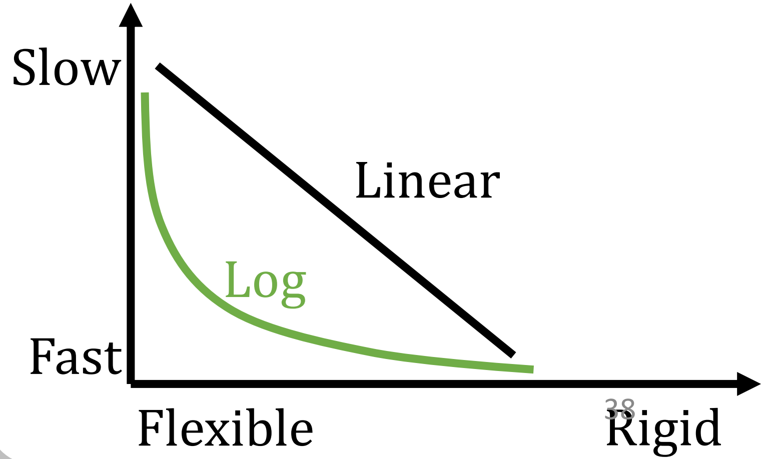

Diversity-enabled sweet spots in engineering

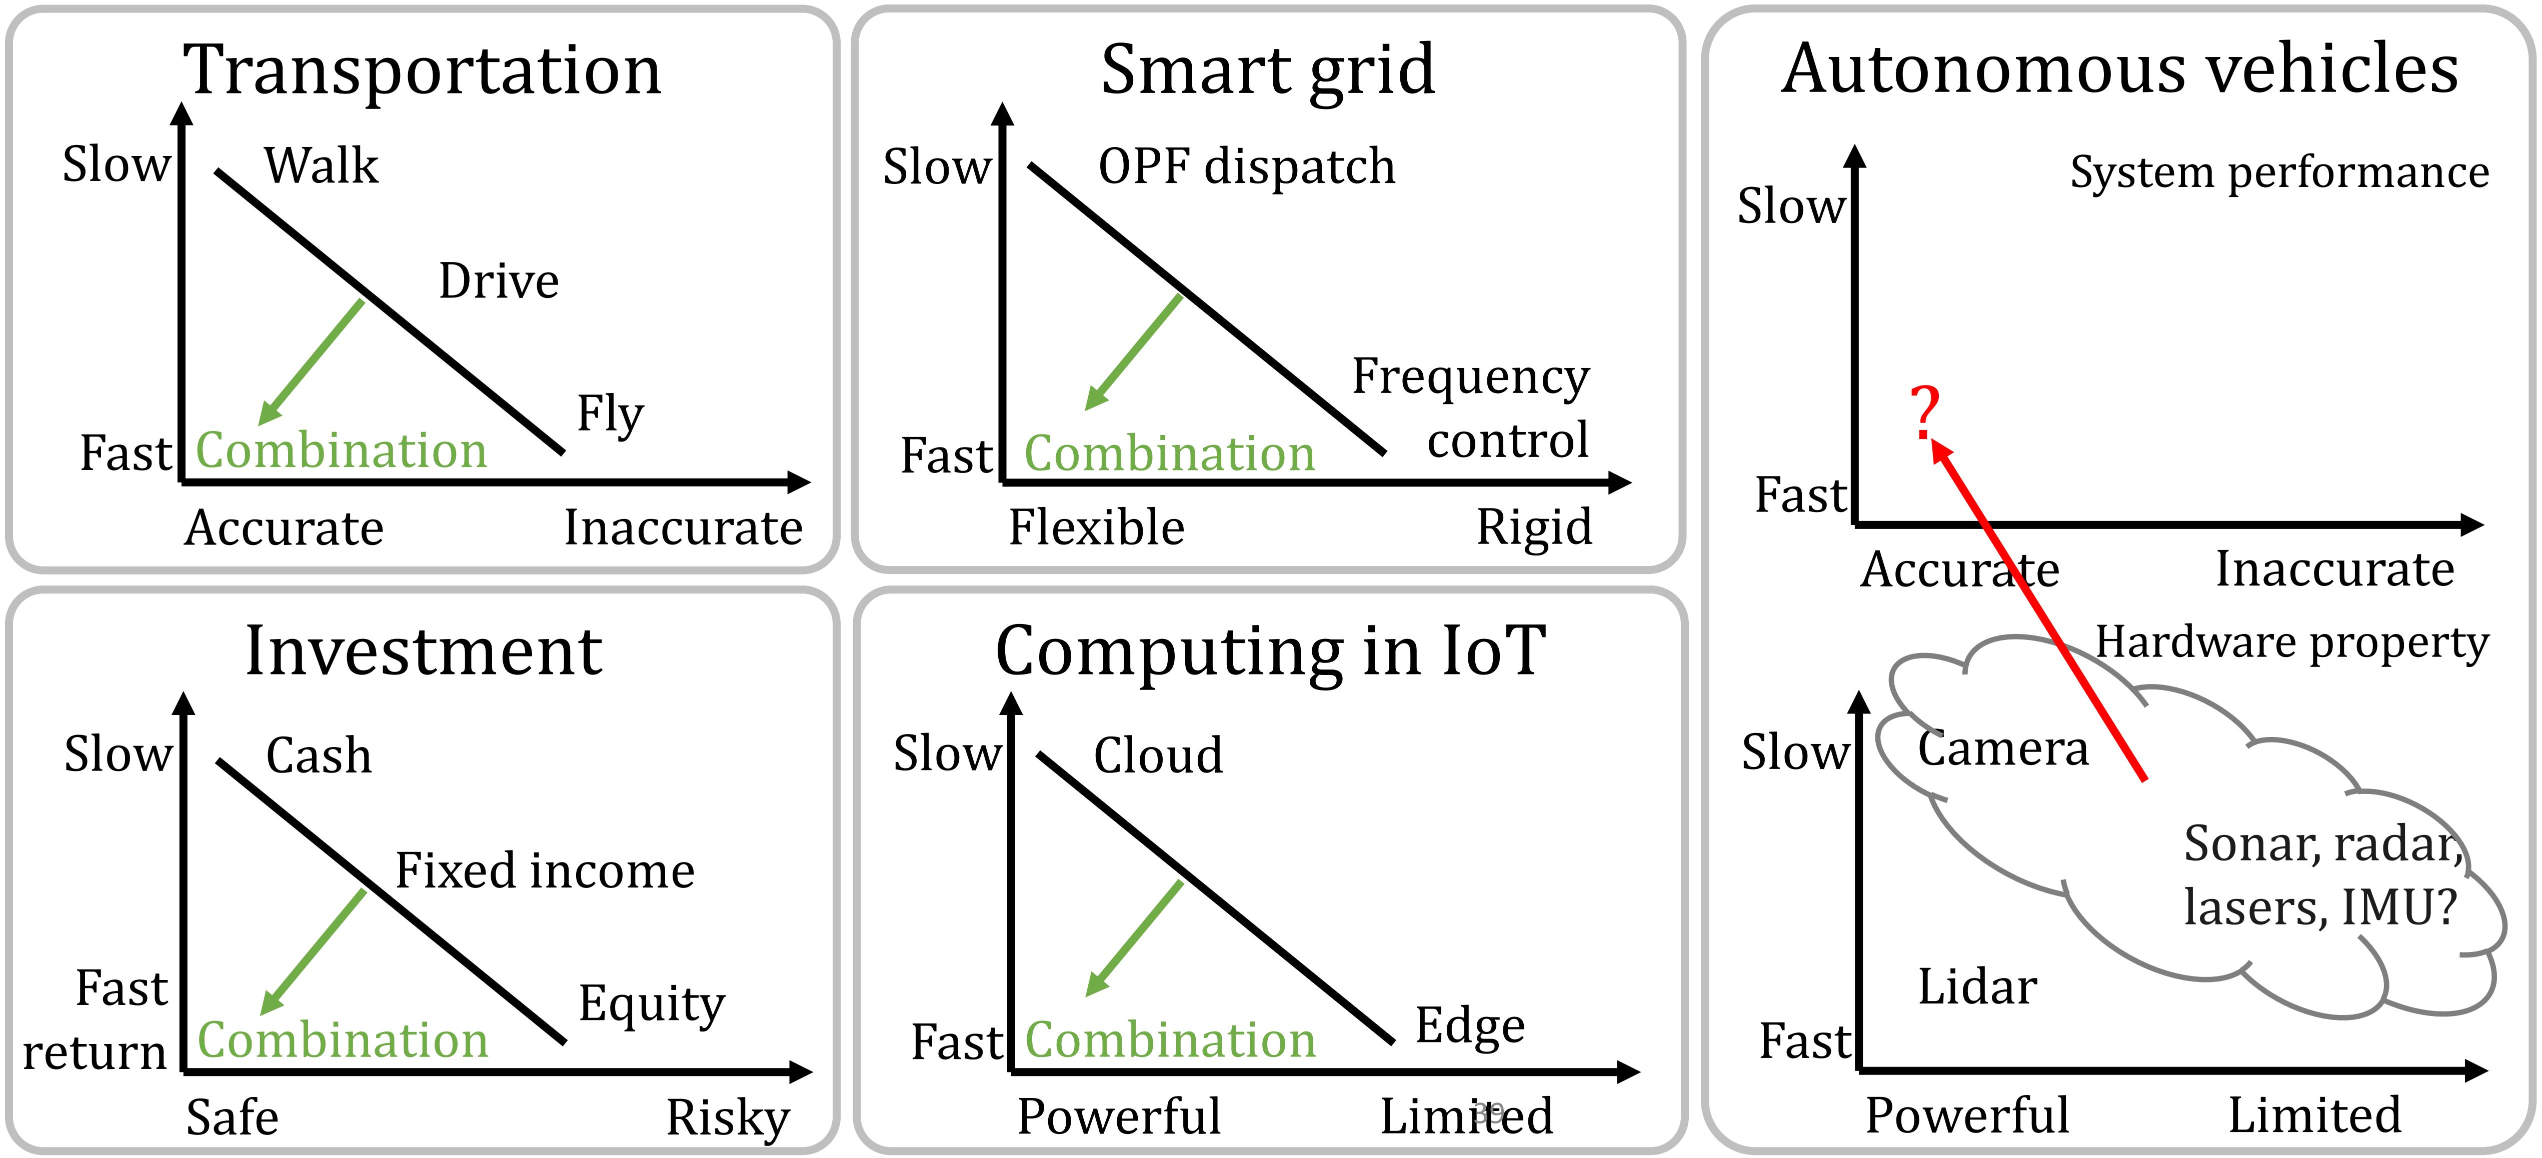

Today's talk

理学

Stochastic safe control
Robust control
Optimization
Information theory …

工学

Motivation

Microfinance is a category of financial services targeting individuals and small businesses who lack access to conventional banking and related services.

Microfinance services are designed to reach excluded customers, usually poorer population segments, possibly socially marginalized, or geographically more isolated, and to help them become self-sufficient.

Microfinance in developing areas has been proven to improve the local economy significantly.

Microfinance from a control perspective

Challenges in microfinance:

1. Complexity in understanding default process
2. Asymmetry, heterogeneity, and incomplete information of individual applications
3. The scarcity of available past data
4. The dynamically evolving social and economic conditions

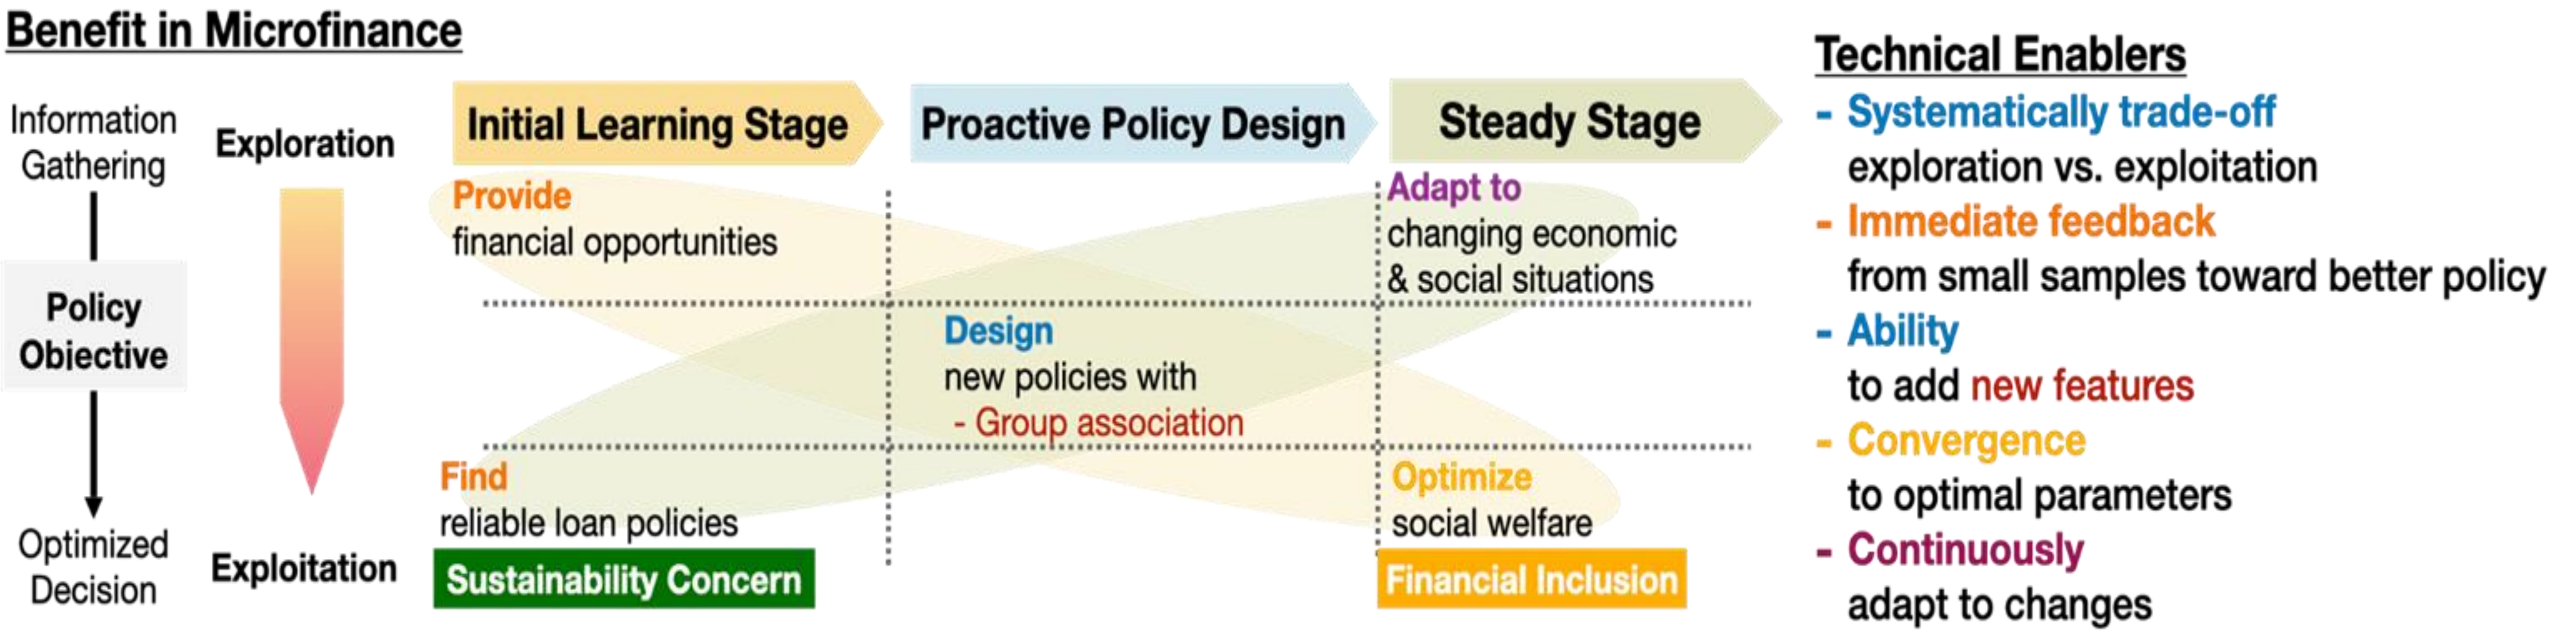

Microfinance from a control perspective

1. **Robustness against missing data**

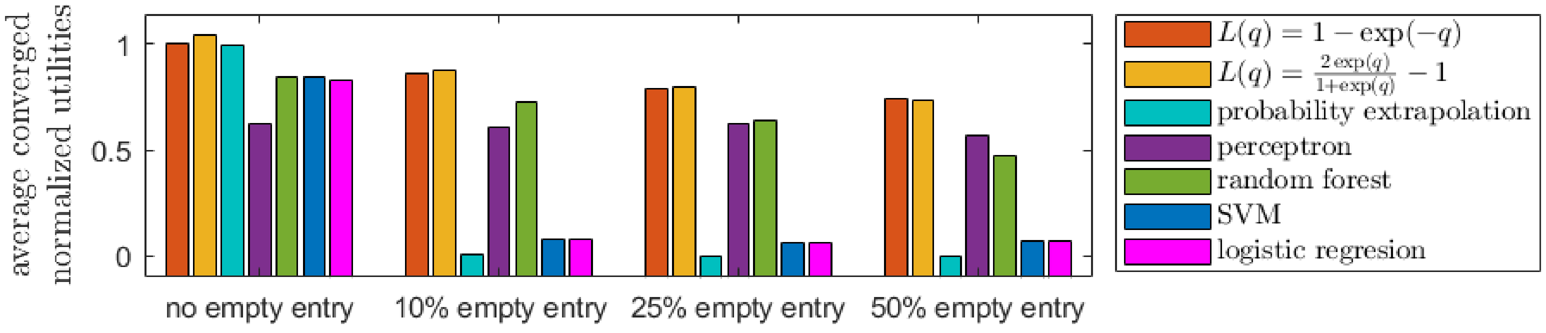

2. **Ability to deal with diverse microfinance distributions**

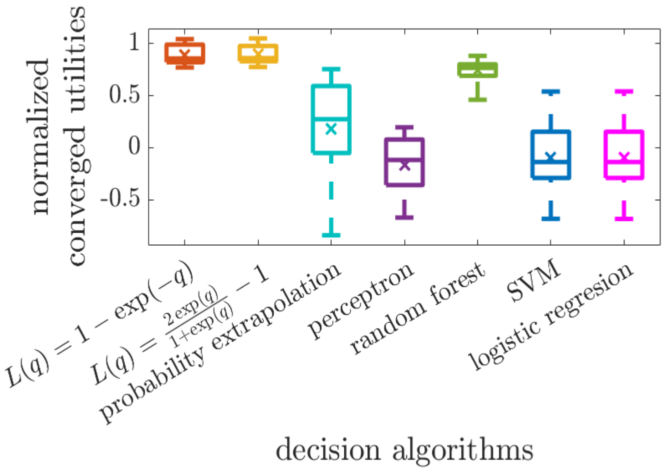

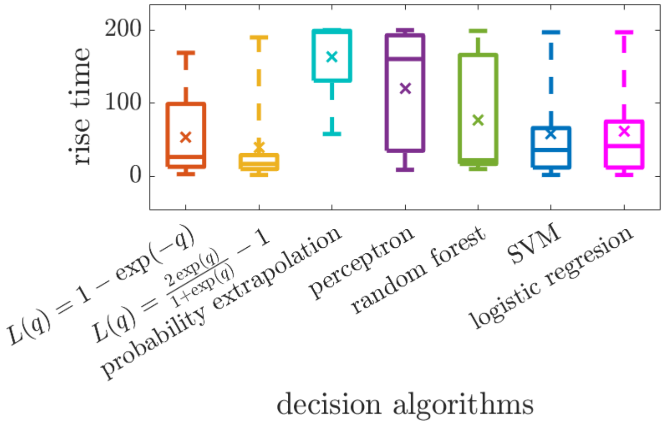

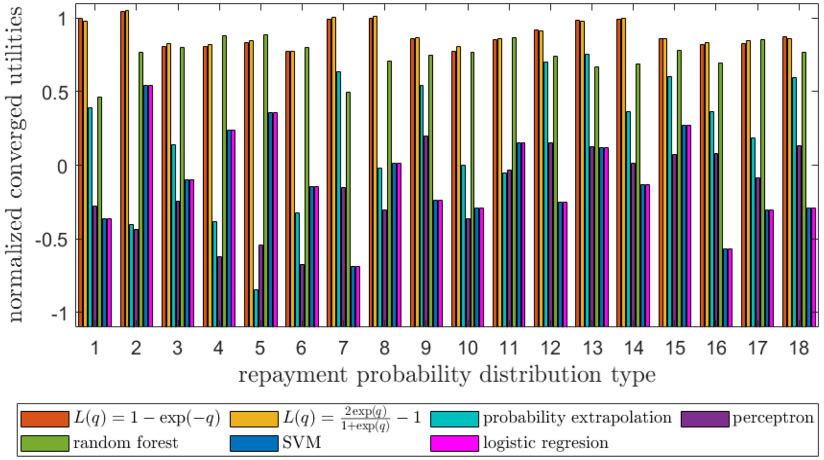

Microfinance from a control perspective

3. **Tradeoff between default rate vs. approval rate**

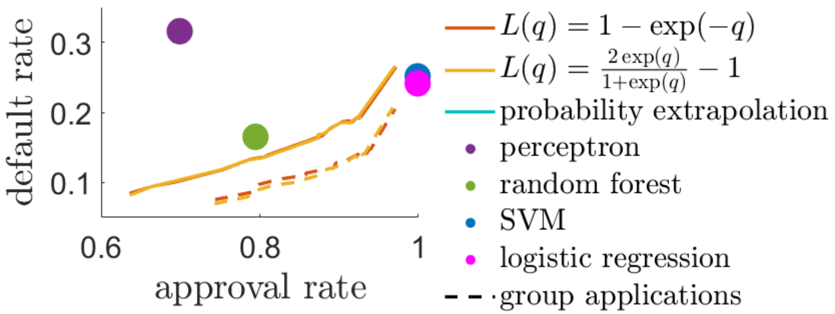

4. **Cheaper computational cost**

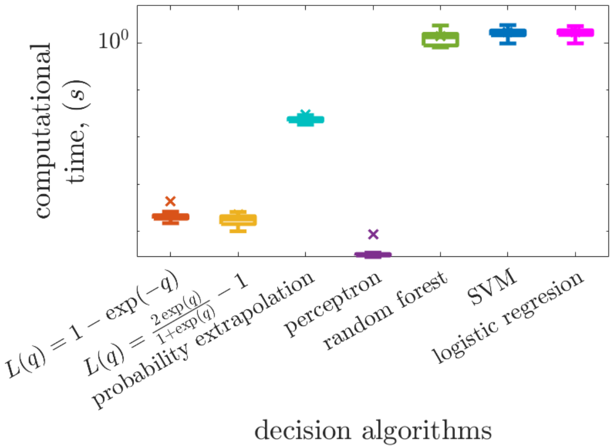

5. **Adaptation to changes**