Lyapunov Density Models: Constraining Distribution Shift in Learning-Based Control

Katie Kang, Paula Gradu, Jason Choi, Michael Janner, Claire Tomlin, Sergey Levine

University of California, Berkeley

Kalashnikov, et al.

Fu, et al.

Smith, et al.

Ebert, et al.

Ebert, et. al

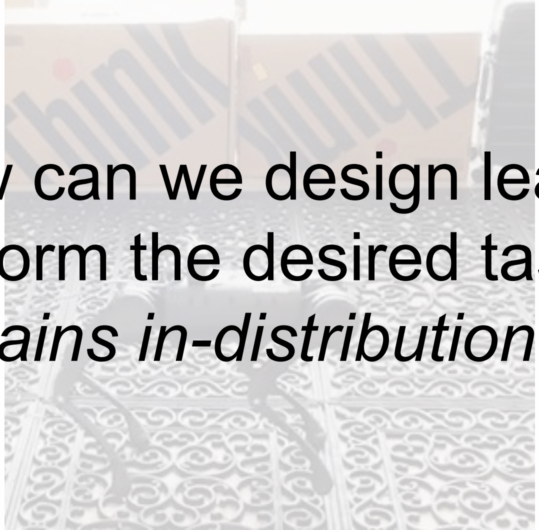

Smith, et. al

How can we design learning-based controllers that perform the desired task, *while ensuring the system remains in-distribution*?

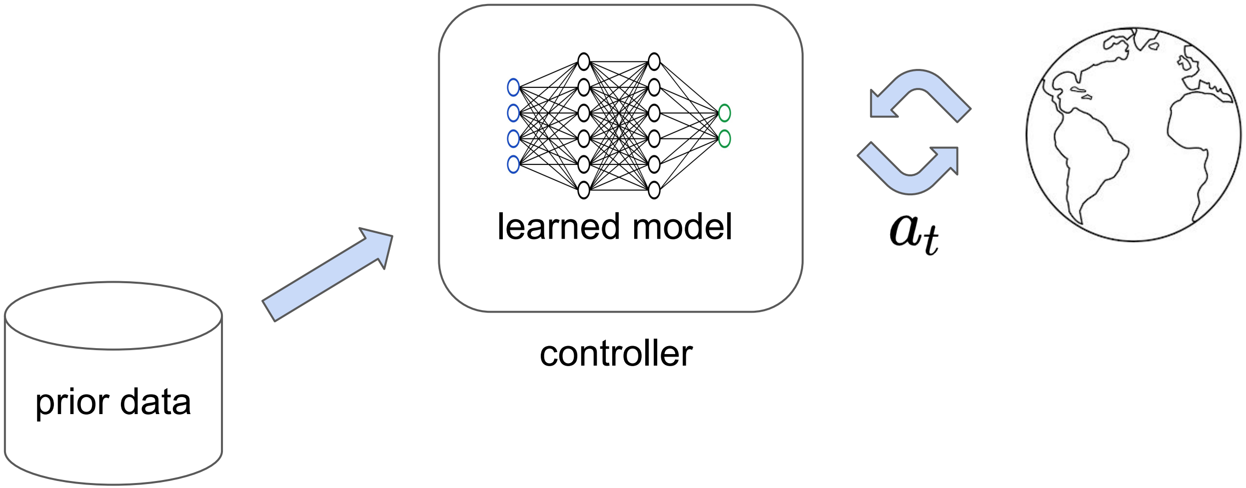
learned model
controller
prior data
a_t

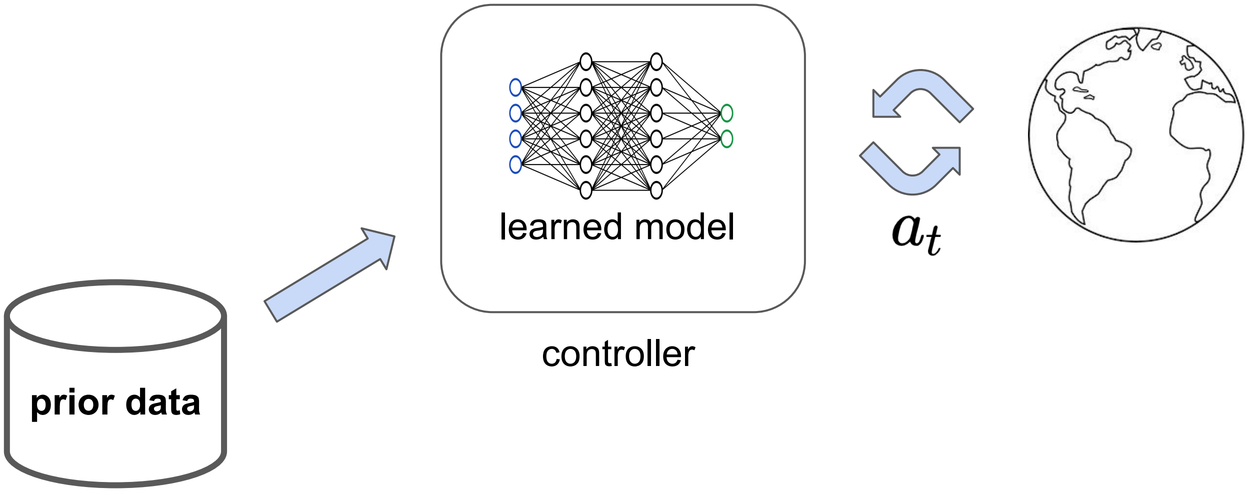
learned model
controller
a_t
prior data

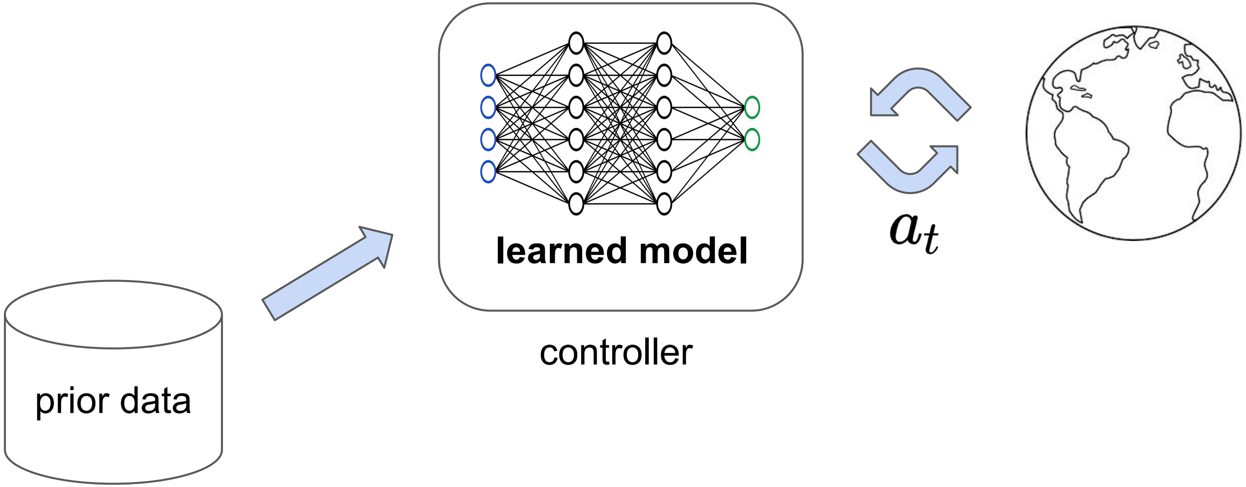
learned model
controller
a_t
prior data

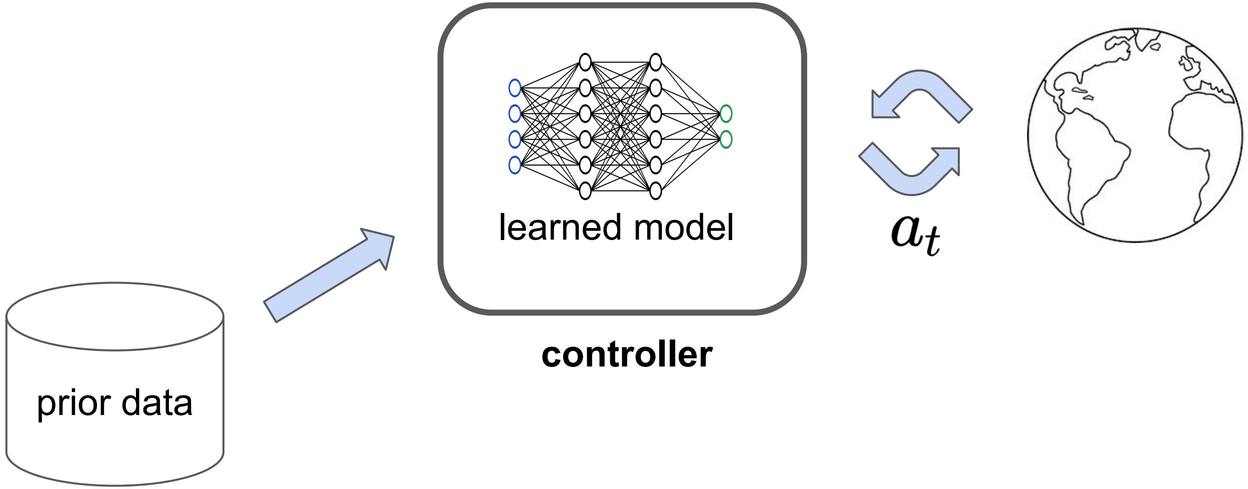
learned model
controller
a_t
prior data

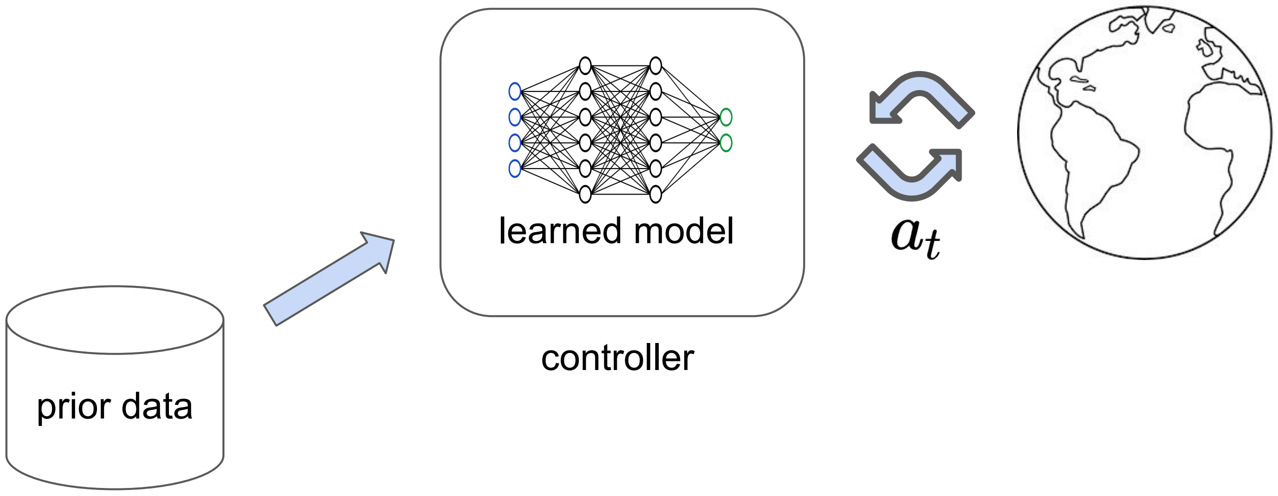
learned model
controller
a_t
prior data

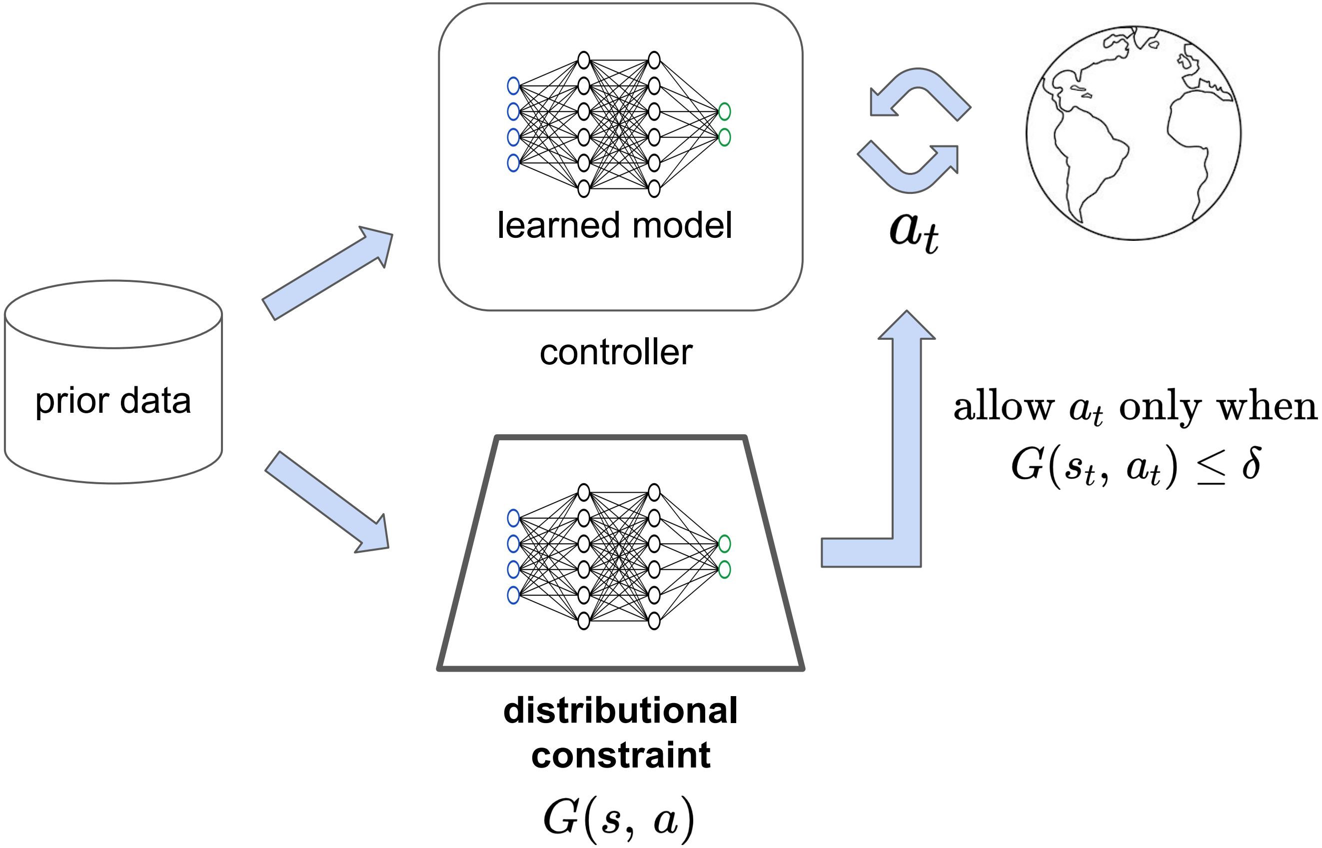
learned model
controller
prior data
distributional constraint
$G(s, a)$
a_t
allow a_t only when $G(s_t, a_t) \leq \delta$

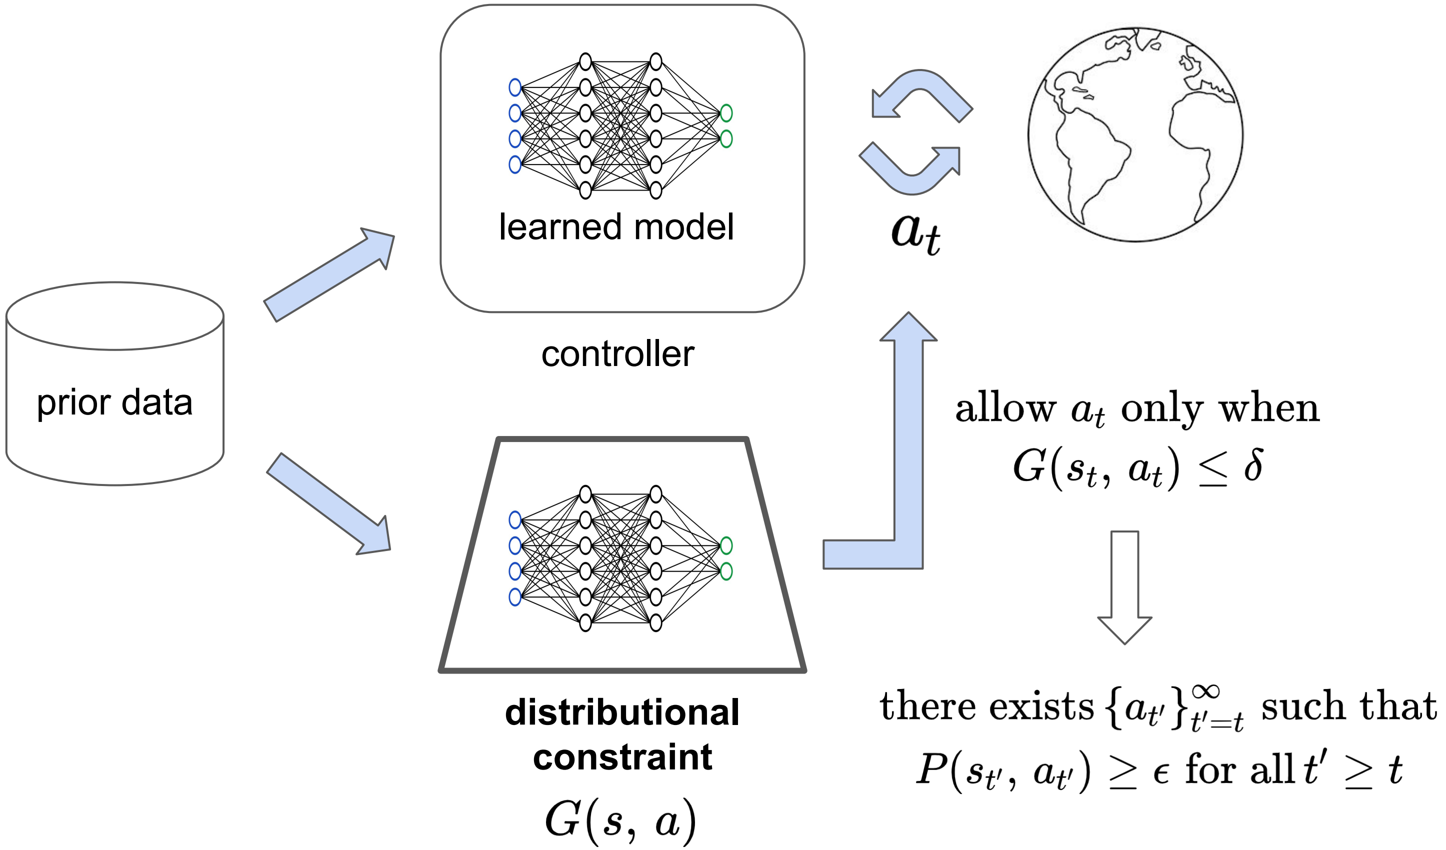

learned model
controller
a_t
prior data
allow a_t only when
$G(s_t, a_t) \le \delta$
distributional
constraint
$G(s, a)$
there exists $\{a_{t'}\}_{t'=t}^{\infty}$ such that
$P(s_{t'}, a_{t'}) \ge \epsilon$ for all $t' \ge t$

Ideal Distributional Constraint

?

$$G(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$
$$P(s_{t'}, a_{t'}) \geq \epsilon \, \forall t' \geq t$$

Ideal Distributional Constraint

?

$$G(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$
$$P(s_{t'}, a_{t'}) \geq \epsilon \, \forall t' \geq t$$

Density Model

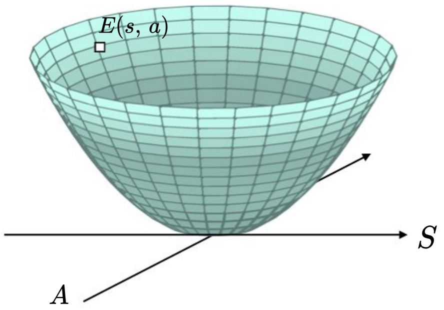

Lyapunov Function

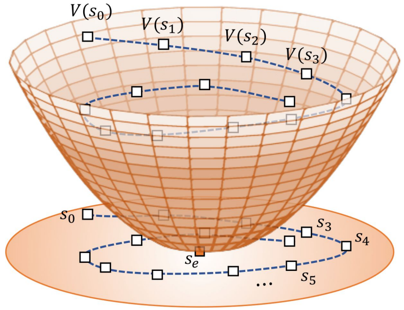

Ideal Distributional Constraint

?

$$G(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$
$$P(s_{t'}, a_{t'}) \geq \epsilon \, \forall t' \geq t$$

Density Model

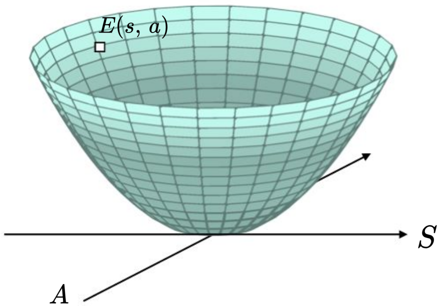

$E(s,\, a)$ approximates $-\log(P(s,\, a))$

Lyapunov Function

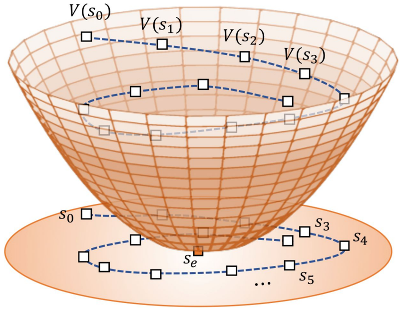

Ideal Distributional Constraint	Density Model	Lyapunov Function
?	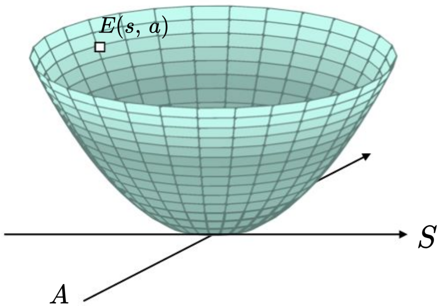 $E(s,\, a)$ approximates $-\log(P(s,\, a))$	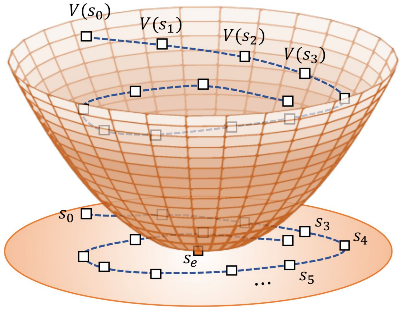
$G(s_t,\, a_t) \leq \delta$ $\Downarrow$ $\exists \{a_{t'}\}_{t'=t}^{\infty}$ s.t. $P(s_{t'}, a_{t'}) \geq \epsilon \, \forall t' \geq t$	$E(s_t,\, a_t) \leq \delta$ $\Downarrow$ $\exists\, a_t$ s.t. $P(s_t,\, a_t) \geq \epsilon$	

Ideal Distributional Constraint

?

$$G(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$

$$P(s_{t'}, a_{t'}) \geq \epsilon \, \forall t' \geq t$$

Density Model

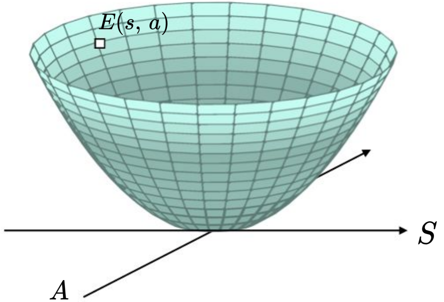

$E(s,\, a)$ approximates $-\log(P(s,\, a))$

$$E(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists\, a_t \text{ s.t.}$$

$$P(s_t,\, a_t) \geq \epsilon$$

Lyapunov Function

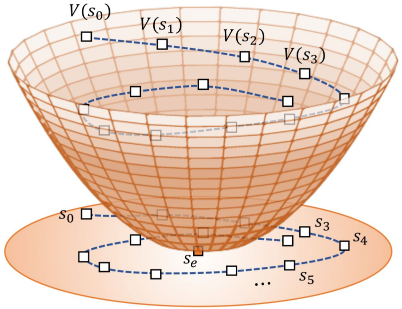

Ideal Distributional Constraint

?

$$G(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$
$$P(s_{t'}, a_{t'}) \geq \epsilon \; \forall t' \geq t$$

Density Model

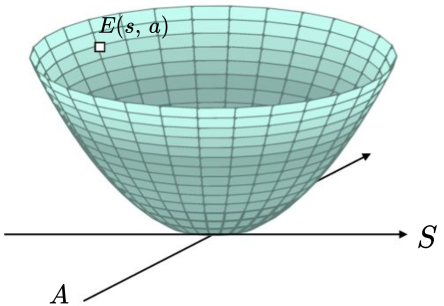

$E(s,\, a)$ approximates $-\log(P(s,\, a))$

$$E(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists\, a_t \text{ s.t.}$$
$$P(s_t,\, a_t) \geq \epsilon$$

Lyapunov Function

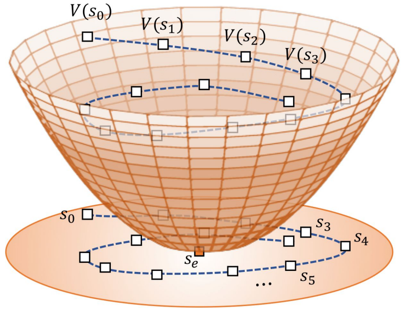

Ideal Distributional Constraint

?

$$G(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$
$$P(s_{t'}, a_{t'}) \geq \epsilon \,\forall t' \geq t$$

Density Model

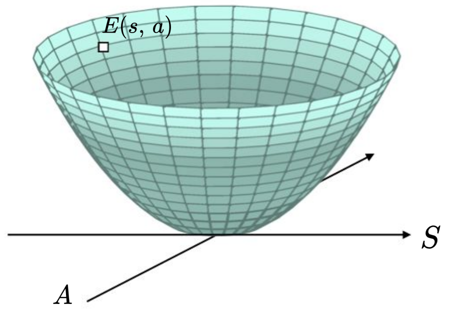

$E(s,\, a)$ approximates $-\log(P(s,\, a))$

$$E(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists\, a_t \text{ s.t.}$$
$$P(s_t,\, a_t) \geq \epsilon$$

Lyapunov Function

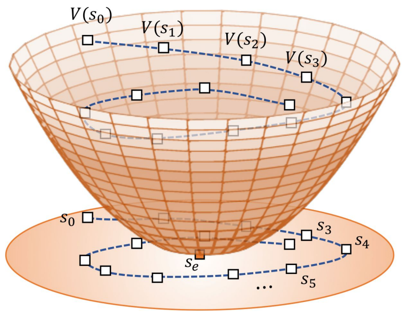

$$V(s_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$
$$||s_{t'} - s_e|| \leq \epsilon \,\forall t' \geq t$$

Ideal Distributional Constraint	Density Model	Lyapunov Function
?	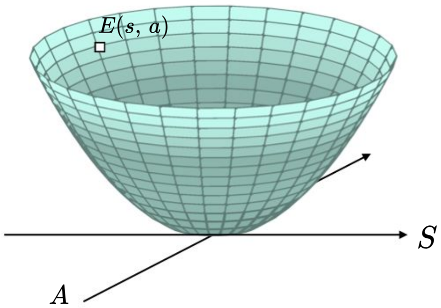 $E(s,\, a)$ approximates $-\log(P(s,\, a))$	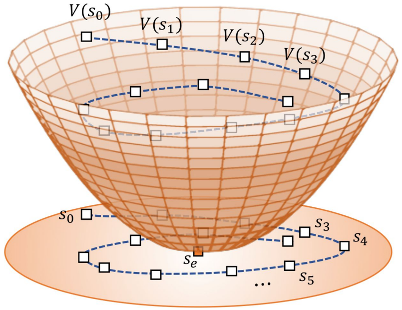
$G(s_t,\, a_t) \leq \delta$	$E(s_t,\, a_t) \leq \delta$	$V(s_t) \leq \delta$
$\Downarrow$	$\Downarrow$	$\Downarrow$
$\exists \{a_{t'}\}_{t'=t}^{\infty}$ s.t.	$\exists\, a_t$ s.t.	$\exists \{a_{t'}\}_{t'=t}^{\infty}$ s.t.
$P(s_{t'}, a_{t'}) \geq \epsilon\, \forall t' \geq t$	$P(s_t,\, a_t) \geq \epsilon$	$\lVert s_{t'} - s_e \rVert \leq \epsilon\, \forall t' \geq t$

Ideal Distributional Constraint

?

$$G(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$

$$P(s_{t'}, a_{t'}) \geq \epsilon \; \forall t' \geq t$$

Density Model

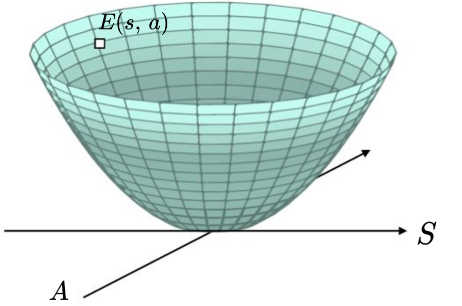

$E(s,\, a)$ approximates $-\log(P(s,\, a))$

$$E(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists\, a_t \text{ s.t.}$$

$$P(s_t,\, a_t) \geq \epsilon$$

Lyapunov Function

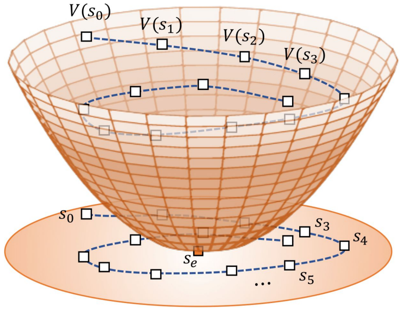

$$V(s_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$

$$||s_{t'} - s_e|| \leq \epsilon \; \forall t' \geq t$$

Lyapunov Density Model (LDM)

Lyapunov Density Model (LDM)

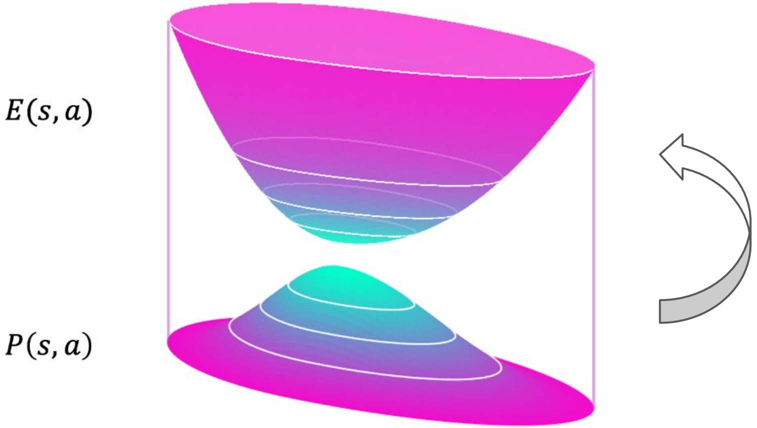

Any density modeling technique

Lyapunov Density Model (LDM)

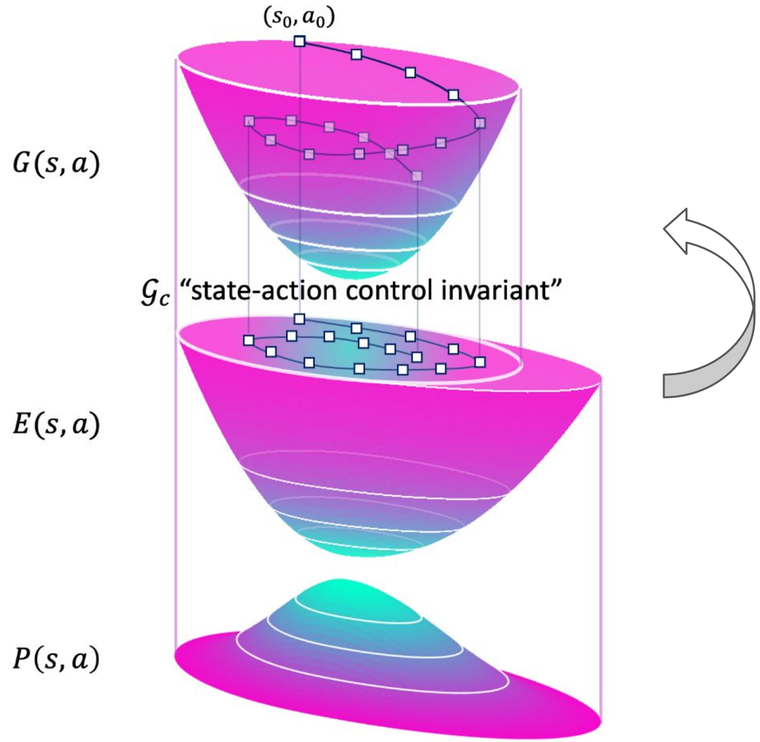

Lyapunov Density Model (LDM)

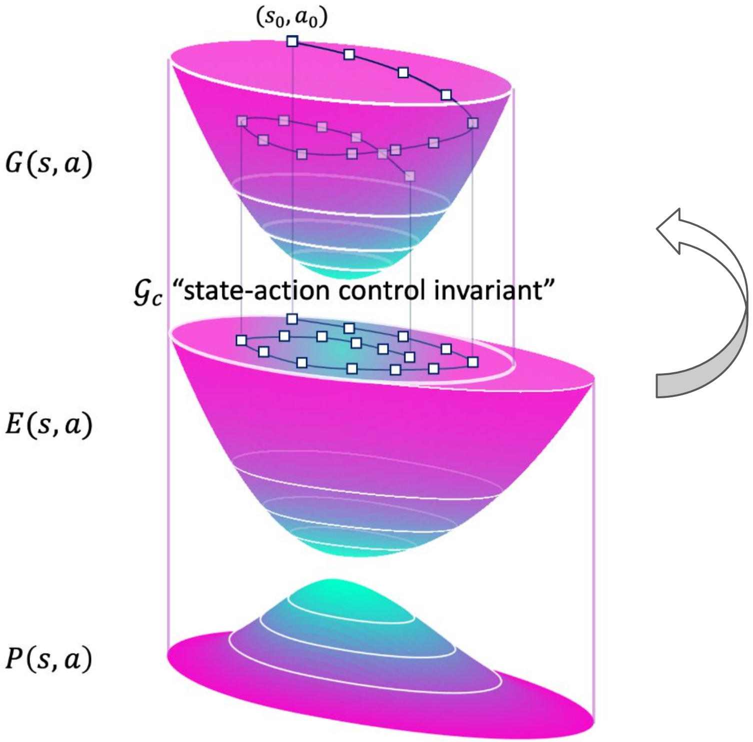

$$G(s_t,\ a_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$
$$P(s_{t'}, a_{t'}) \geq \epsilon \, \forall t' \geq t$$

Lyapunov Density Model (LDM)

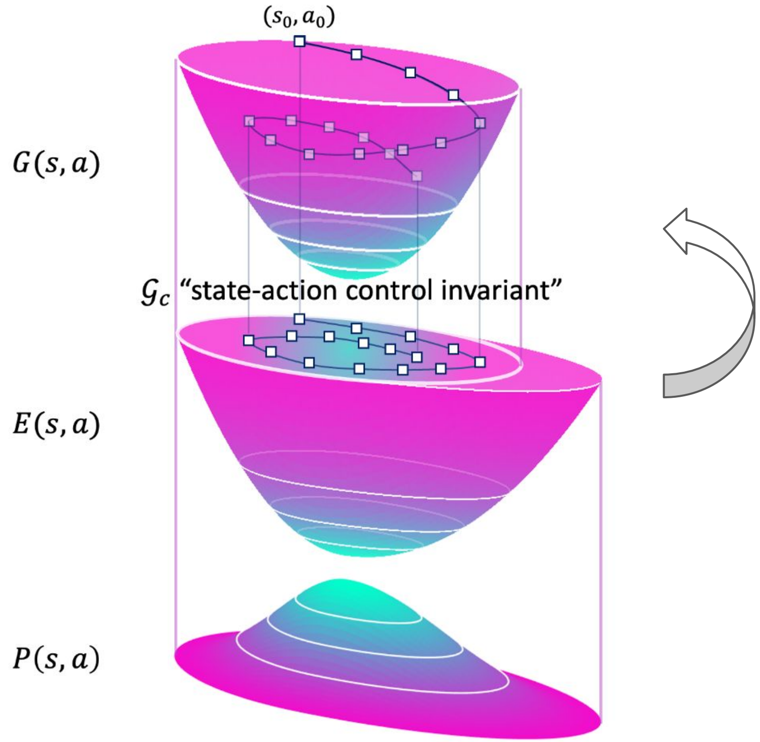

$$G(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$
$$P(s_{t'}, a_{t'}) \geq \epsilon \, \forall t' \geq t$$

$$G(s_t,\, a_t) \; = \; \min_{\{a_{t'}\}_{t'=t}^{\infty}} \max_{t' \geq t} -\log(P(s_{t'},\, a_{t'}))$$

Lyapunov Density Model (LDM)

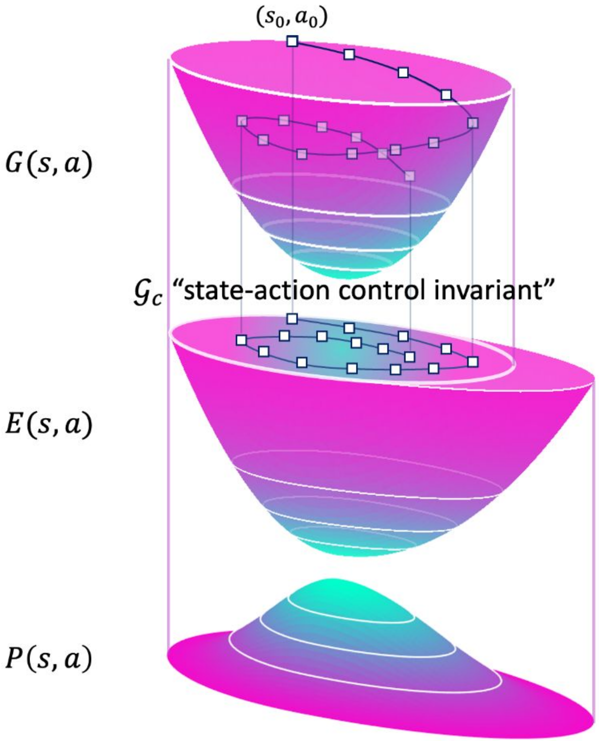

$$G(s_t,\ a_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$
$$P(s_{t'}, a_{t'}) \geq \epsilon \, \forall t' \geq t$$

$$G(s_t,\ a_t) \ = \ \min_{\{a_{t'}\}_{t'=t}^{\infty}} \ \max_{t' \geq t} -\log(P(s_{t'},\ a_{t'}))$$

Can be solved recursively via

$$TG(s_t,\ a_t) \ = \max\left\{-\log(P(s_t,\ a_t)),\ \min_{a_{t+1}} G(s_{t+1},\ a_{t+1})\right\}$$

LDM operator

Lyapunov Density Model (LDM)

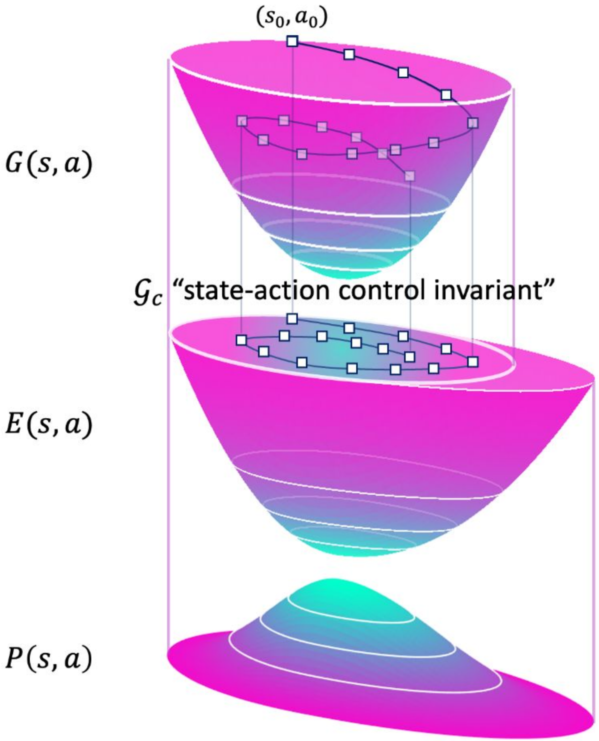

1. initialize $G = E$
2. iteratively update $G = TG$

$$G(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$
$$P(s_{t'}, a_{t'}) \geq \epsilon \, \forall t' \geq t$$

$$G(s_t,\, a_t) = \min_{\{a_{t'}\}_{t'=t}^{\infty}} \max_{t' \geq t} -\log(P(s_{t'},\, a_{t'}))$$

Can be solved recursively via

$$TG(s_t,\, a_t) = \max\left\{ -\log(P(s_t,\, a_t)),\ \min_{a_{t+1}} G(s_{t+1},\, a_{t+1}) \right\}$$

LDM operator, analogous to Bellman operator

Lyapunov Density Model (LDM)

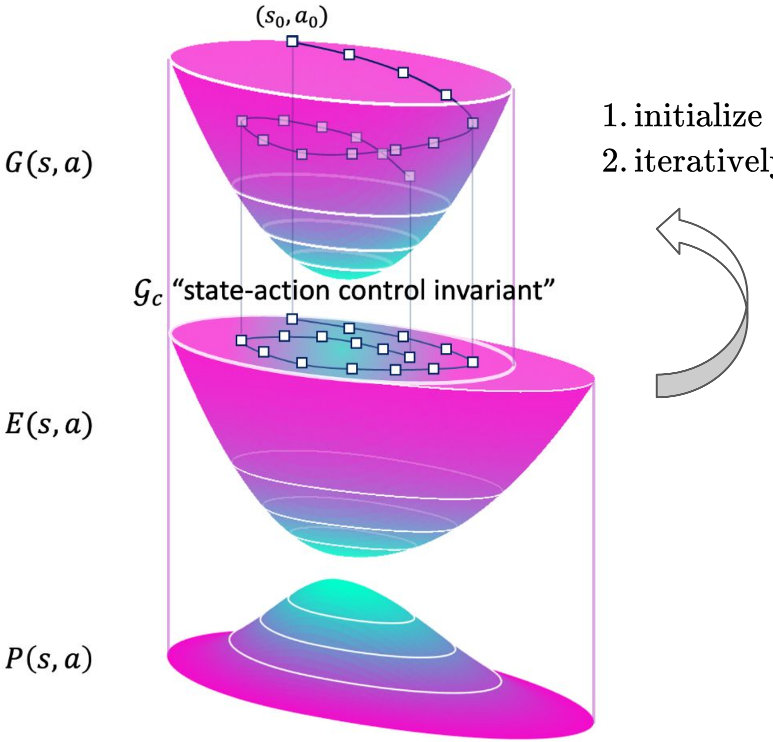

1. initialize $G = E$
2. iteratively update $G = TG$

$$G(s_t,\, a_t) \leq \delta$$

$$\Downarrow$$

$$\exists \{a_{t'}\}_{t'=t}^{\infty} \text{ s.t.}$$

$$P(s_{t'}, a_{t'}) \geq \epsilon \, \forall t' \geq t$$

$$G(s_t,\, a_t) \; = \; \min_{\{a_{t'}\}_{t'=t}^{\infty}} \; \max_{t' \geq t} -\log(P(s_{t'},\, a_{t'}))$$

Can be solved recursively via

$$TG(s_t,\, a_t) \; = \max\left\{ -\log(P(s_t,\, a_t)),\; \min_{a_{t+1}} G(s_{t+1},\, a_{t+1}) \right\}$$

LDM operator, analogous to Bellman operator

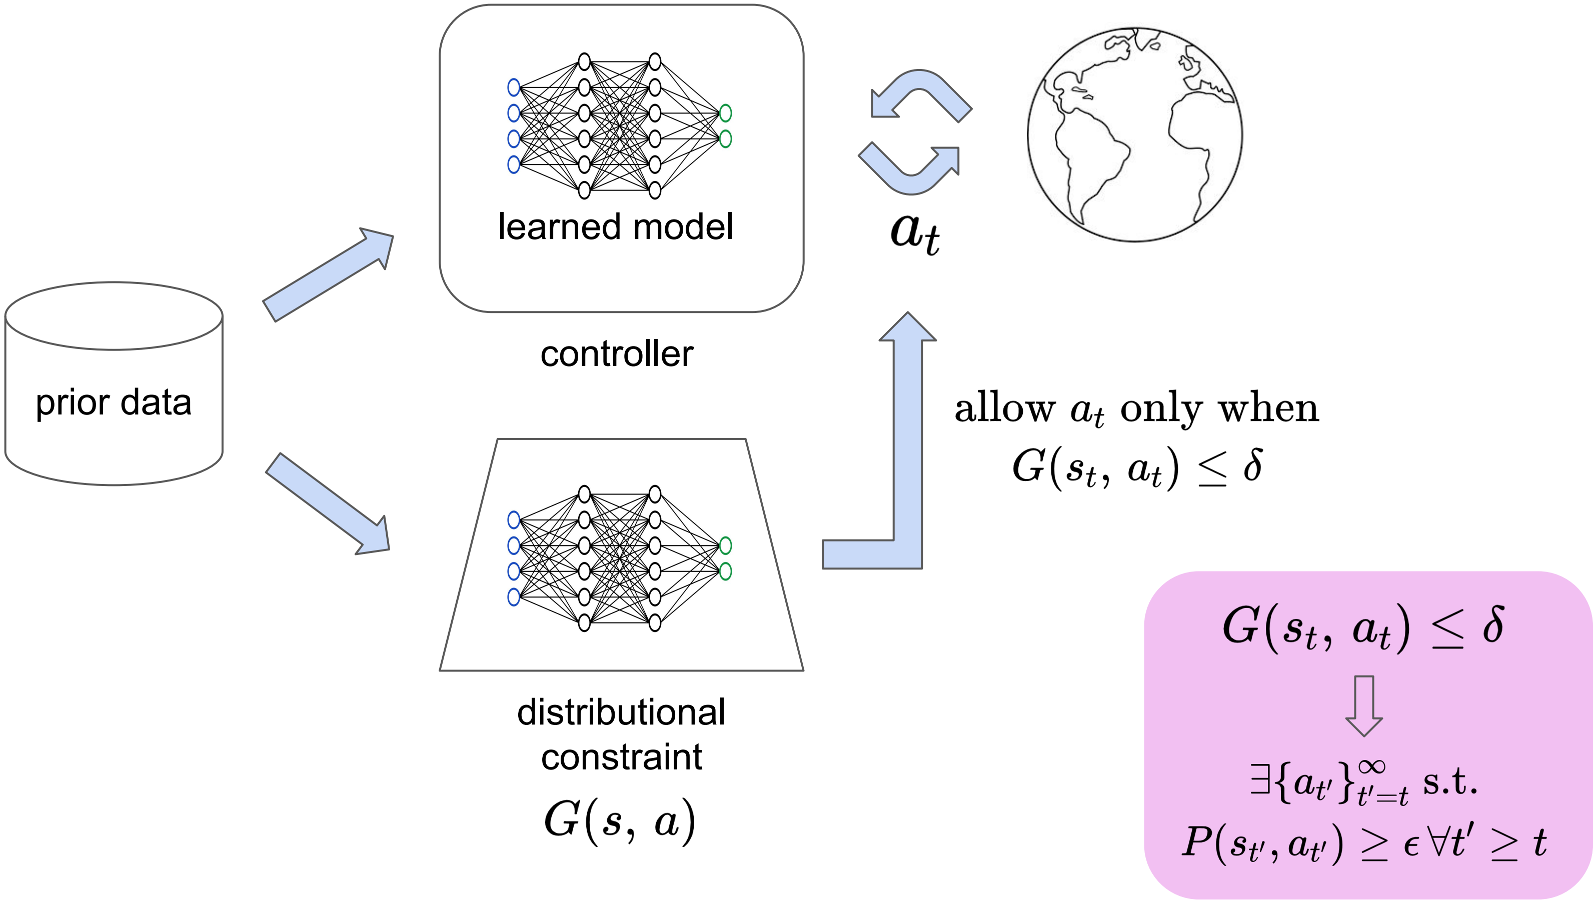
learned model
controller
a_t
prior data
allow a_t only when
$G(s_t,\, a_t) \leq \delta$
distributional
constraint
$G(s,\, a)$
$G(s_t,\, a_t) \leq \delta$
$\Downarrow$
$\exists \{a_{t'}\}_{t'=t}^{\infty}$ s.t.
$P(s_{t'}, a_{t'}) \geq \epsilon \, \forall t' \geq t$

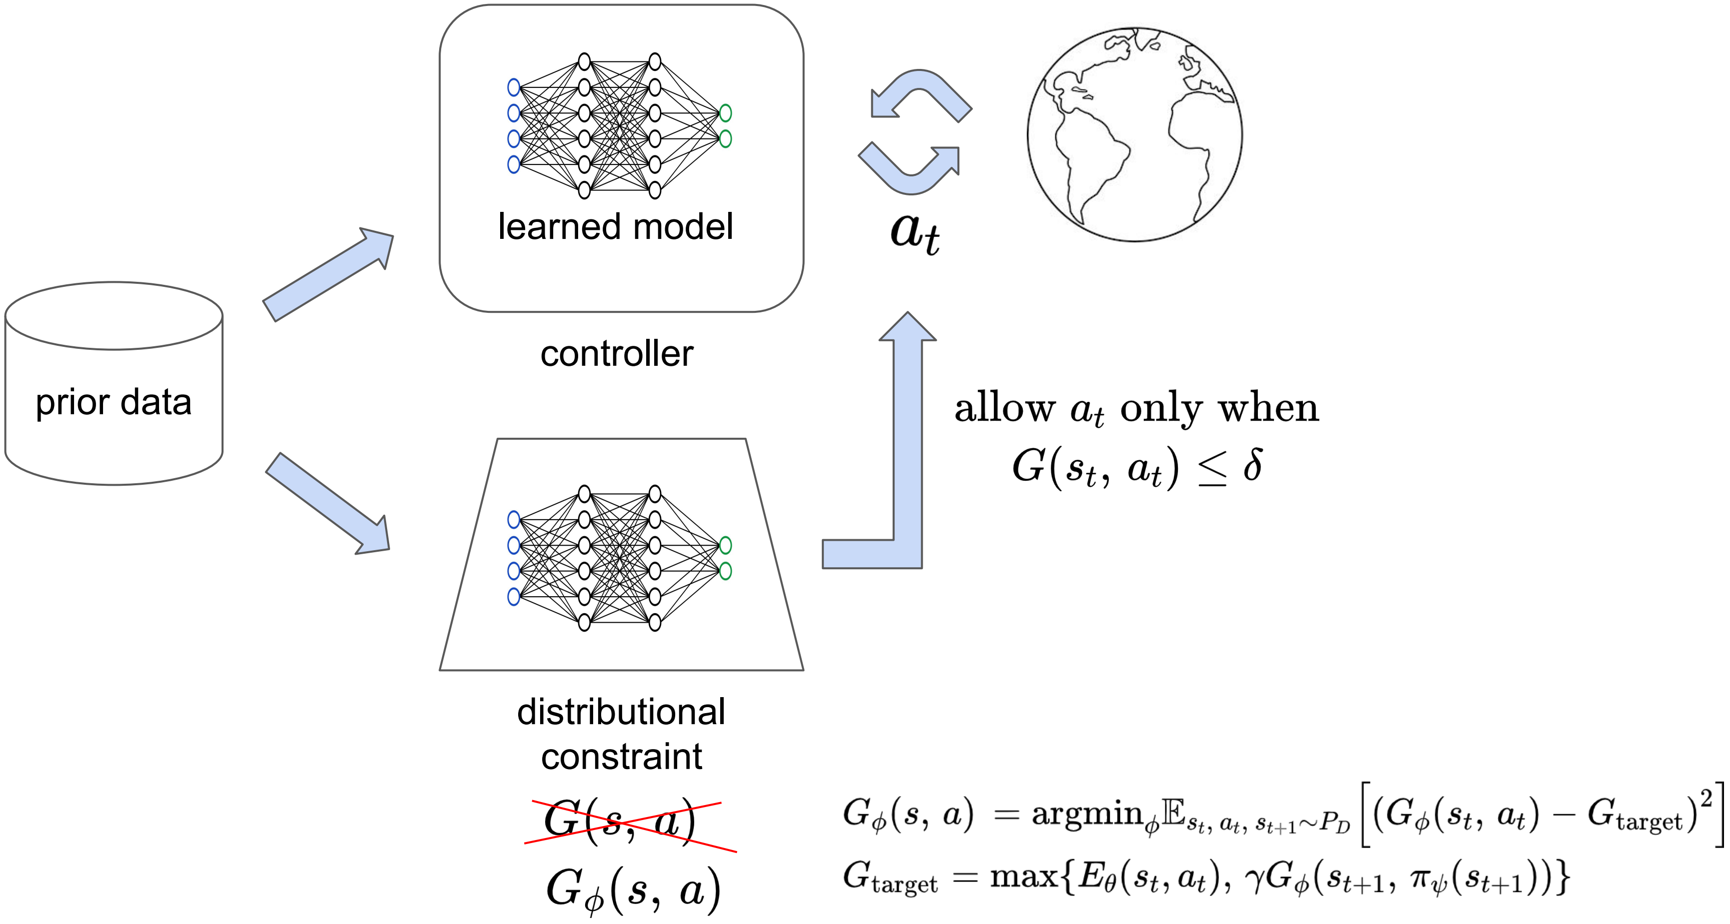

prior data
learned model
controller
a_t
allow a_t only when
$G(s_t, a_t) \leq \delta$
distributional
constraint
$G(s, a)$
$G_\phi(s, a)$
$G_\phi(s, a) = \text{argmin}_\phi \mathbb{E}_{s_t, a_t, s_{t+1} \sim P_D}\left[\left(G_\phi(s_t, a_t) - G_{\text{target}}\right)^2\right]$
$G_{\text{target}} = \max\{E_\theta(s_t, a_t), \gamma G_\phi(s_{t+1}, \pi_\psi(s_{t+1}))\}$

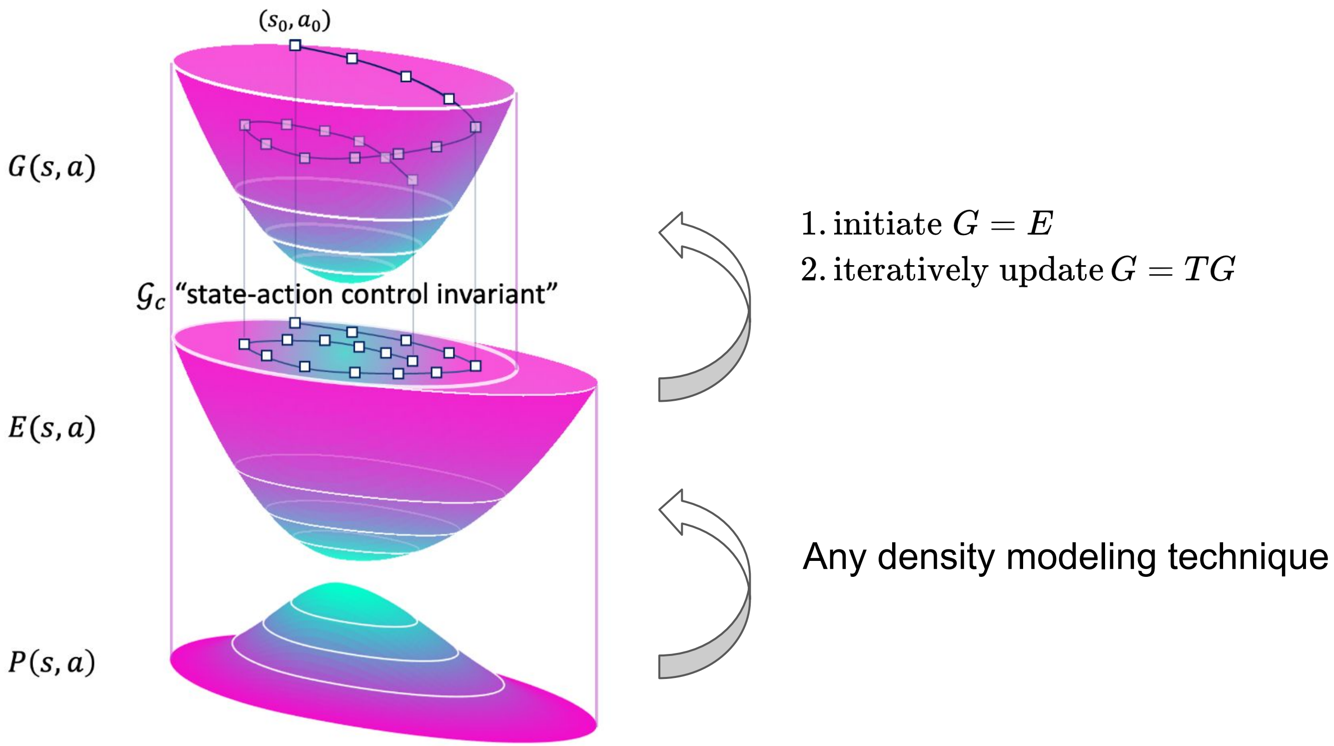
(s_0, a_0)
$G(s,a)$
$\mathcal{G}_c$ "state-action control invariant"
$E(s,a)$
$P(s,a)$
1. initiate $G = E$
2. iteratively update $G = TG$
Any density modeling technique

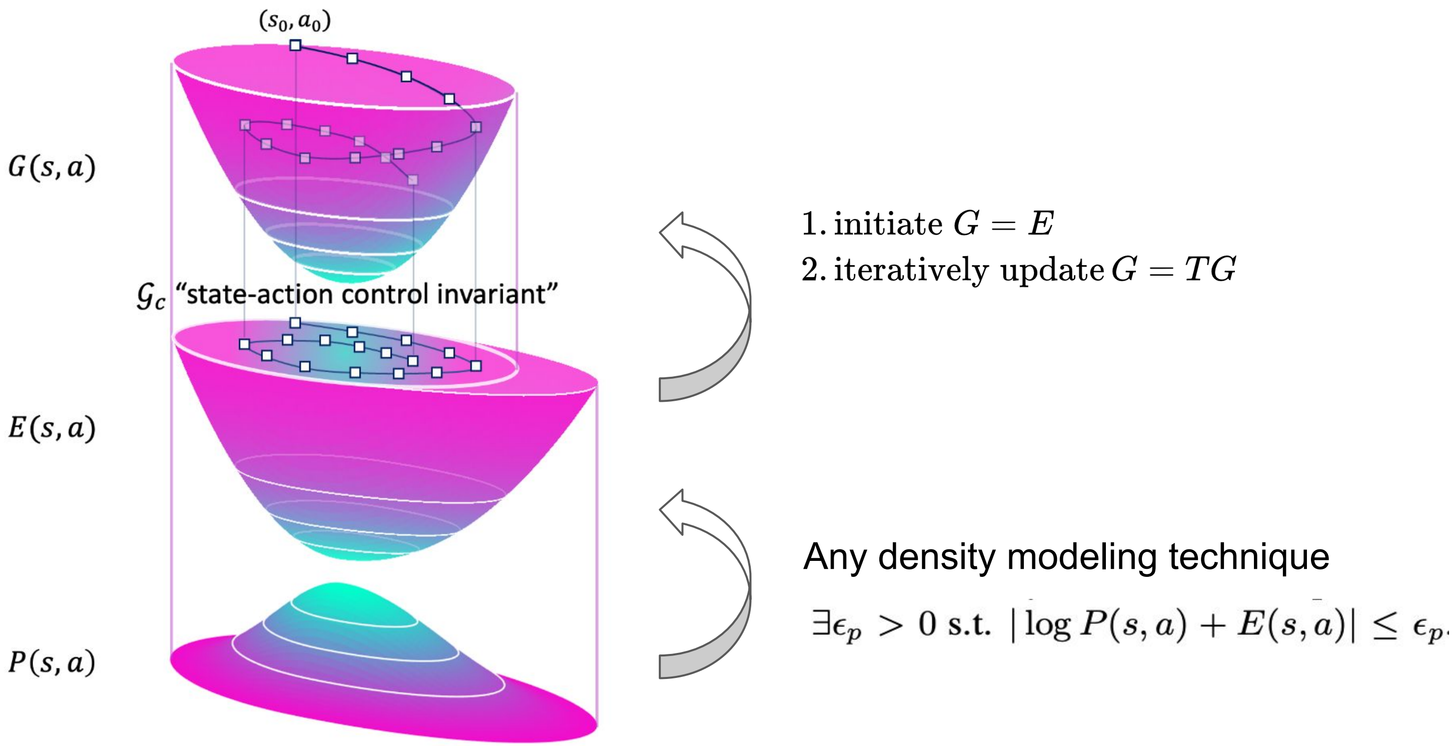

(s_0, a_0)
$G(s, a)$
$\mathcal{G}_c$ “state-action control invariant”
$E(s, a)$
$P(s, a)$
1. initiate $G = E$
2. iteratively update $G = TG$
Any density modeling technique
$\exists \epsilon_p > 0$ s.t. $|\log P(s, a) + E(s, a)| \leq \epsilon_p$.

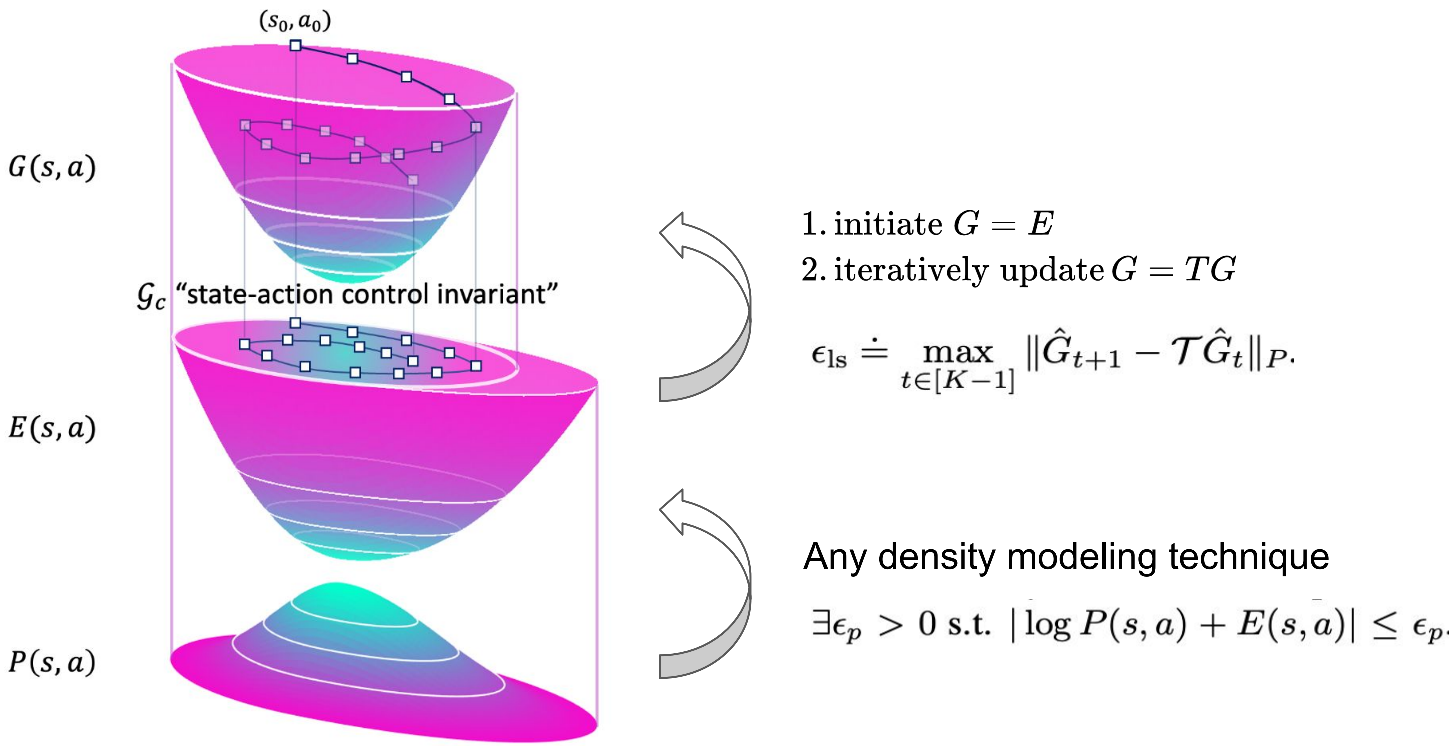

(s_0, a_0)
$G(s,a)$
$\mathcal{G}_c$ "state-action control invariant"
$E(s,a)$
$P(s,a)$
1. initiate $G = E$
2. iteratively update $G = TG$
$\epsilon_{\text{ls}} \doteq \max_{t \in [K-1]} \|\hat{G}_{t+1} - \mathcal{T}\hat{G}_t\|_P.$
Any density modeling technique
$\exists \epsilon_p > 0$ s.t. $|\log P(s,a) + E(s,a)| \leq \epsilon_p.$

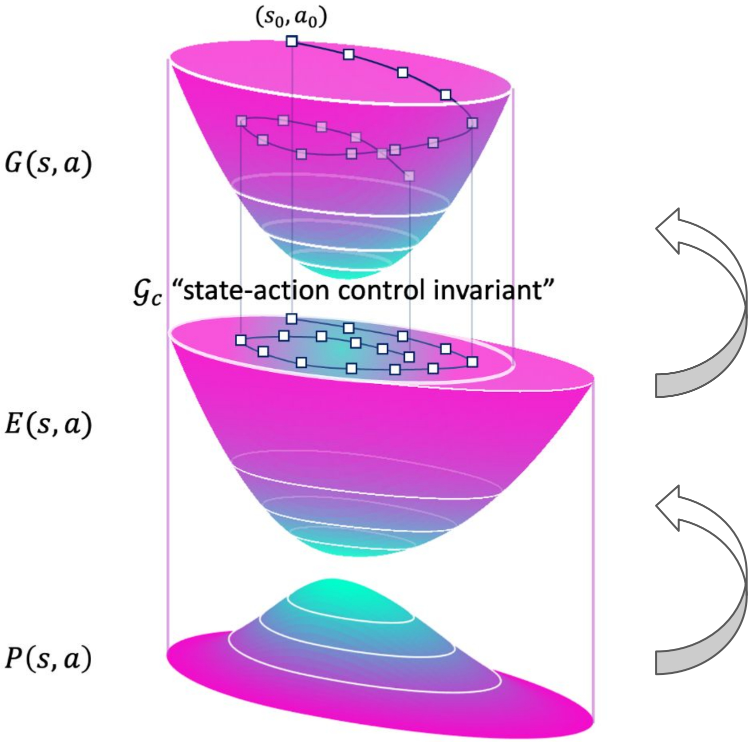

1. initiate $G = E$
2. iteratively update $G = TG$

$$\epsilon_{\text{ls}} \doteq \max_{t \in [K-1]} \|\hat{G}_{t+1} - \mathcal{T}\hat{G}_t\|_P.$$

Any density modeling technique

$$\exists \epsilon_p > 0 \text{ s.t. } |\log P(s,a) + E(s,a)| \leq \epsilon_p.$$

$$\log P(s_t, a_t) \geq \gamma^{-t} \log c - \frac{\gamma^{-t} R \, \epsilon_{\text{ls}} \exp \epsilon_p}{c(1-\gamma)} - \epsilon_P$$

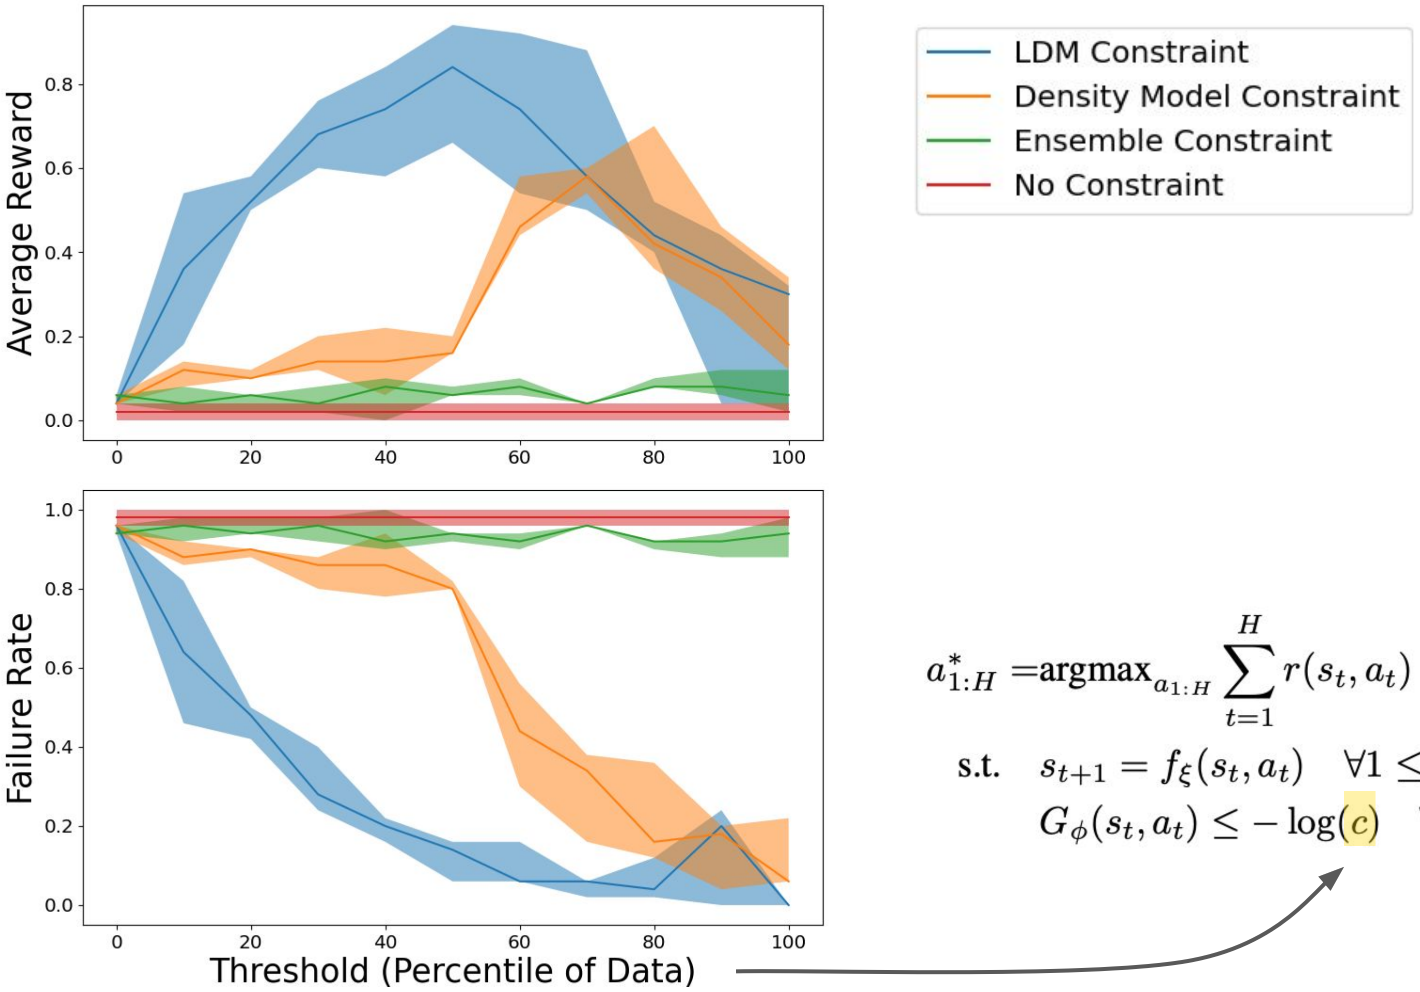

$$a^*_{1:H} = \text{argmax}_{a_{1:H}} \sum_{t=1}^{H} r(s_t, a_t)$$

$$\text{s.t.} \quad s_{t+1} = f_\xi(s_t, a_t) \quad \forall 1 \leq t \leq H-1$$

$$G_\phi(s_t, a_t) \leq -\log(c) \quad \forall 1 \leq t \leq H$$

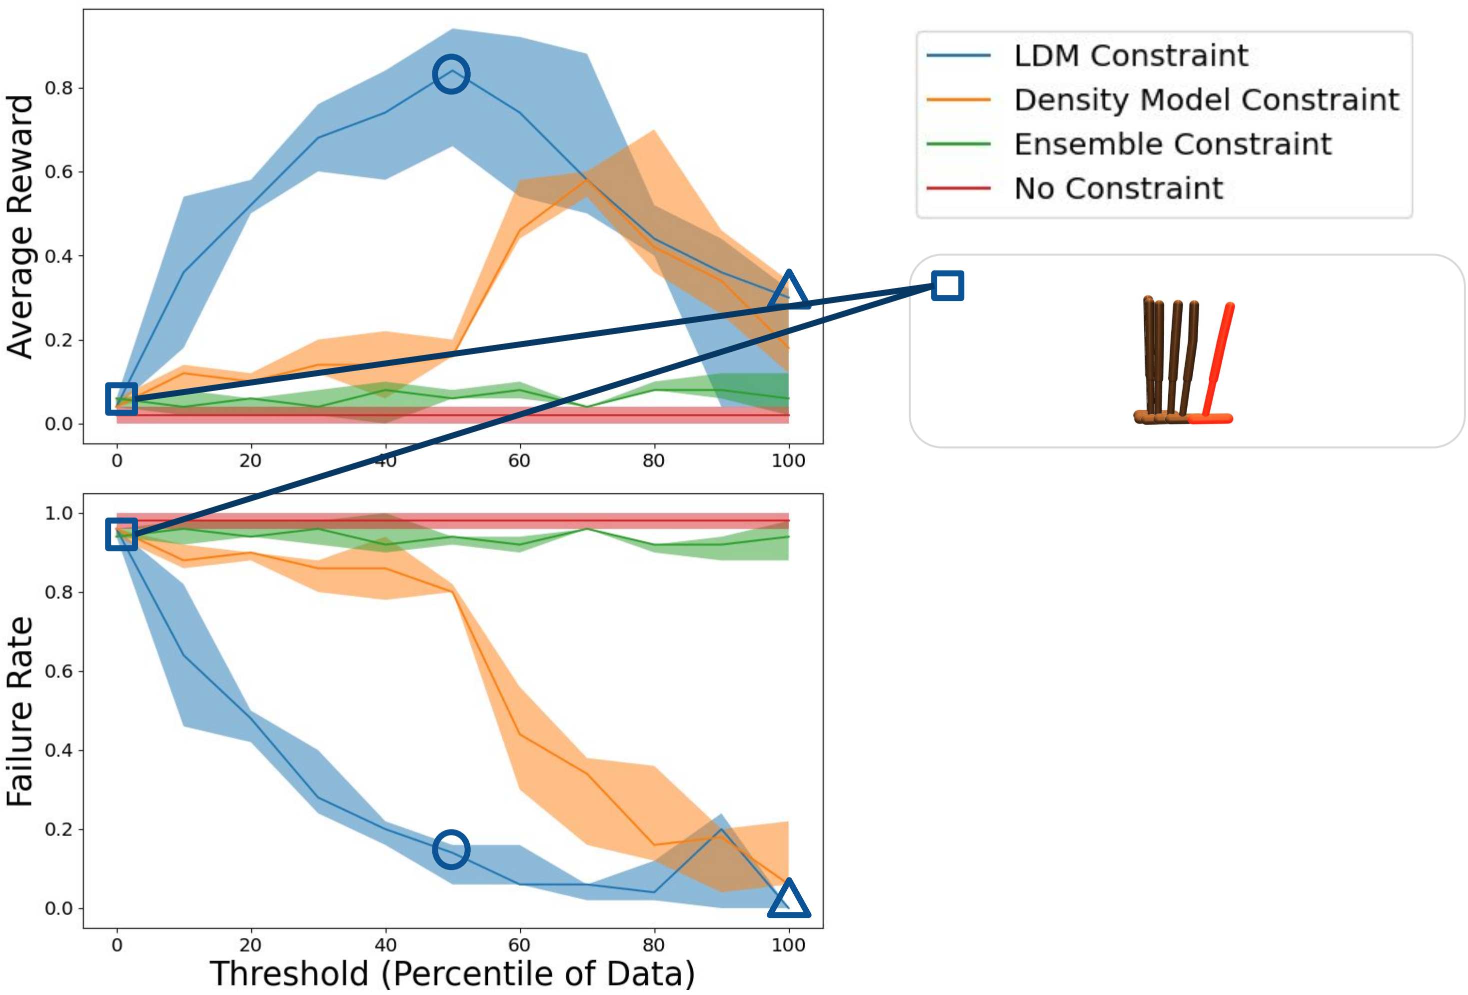

LDM Constraint
Density Model Constraint
Ensemble Constraint
No Constraint
Average Reward
0.0
0.2
0.4
0.6
0.8
0
20
40
60
80
100
Failure Rate
0.0
0.2
0.4
0.6
0.8
1.0
Threshold (Percentile of Data)

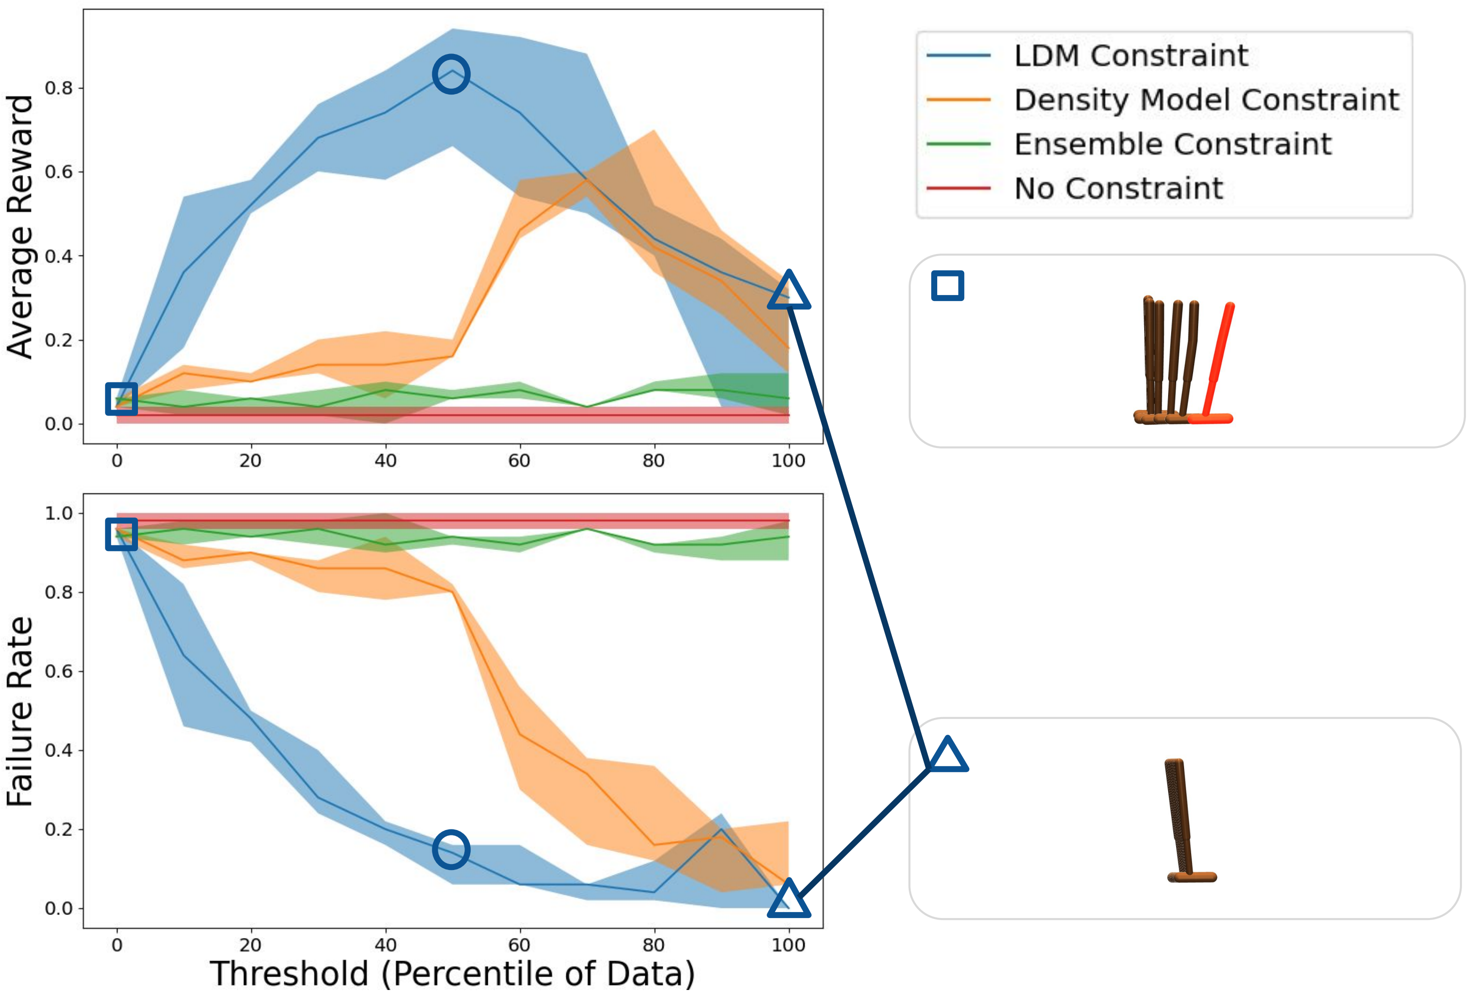

Average Reward
Failure Rate
Threshold (Percentile of Data)
0.0
0.2
0.4
0.6
0.8
1.0
0
20
40
60
80
100
LDM Constraint
Density Model Constraint
Ensemble Constraint
No Constraint

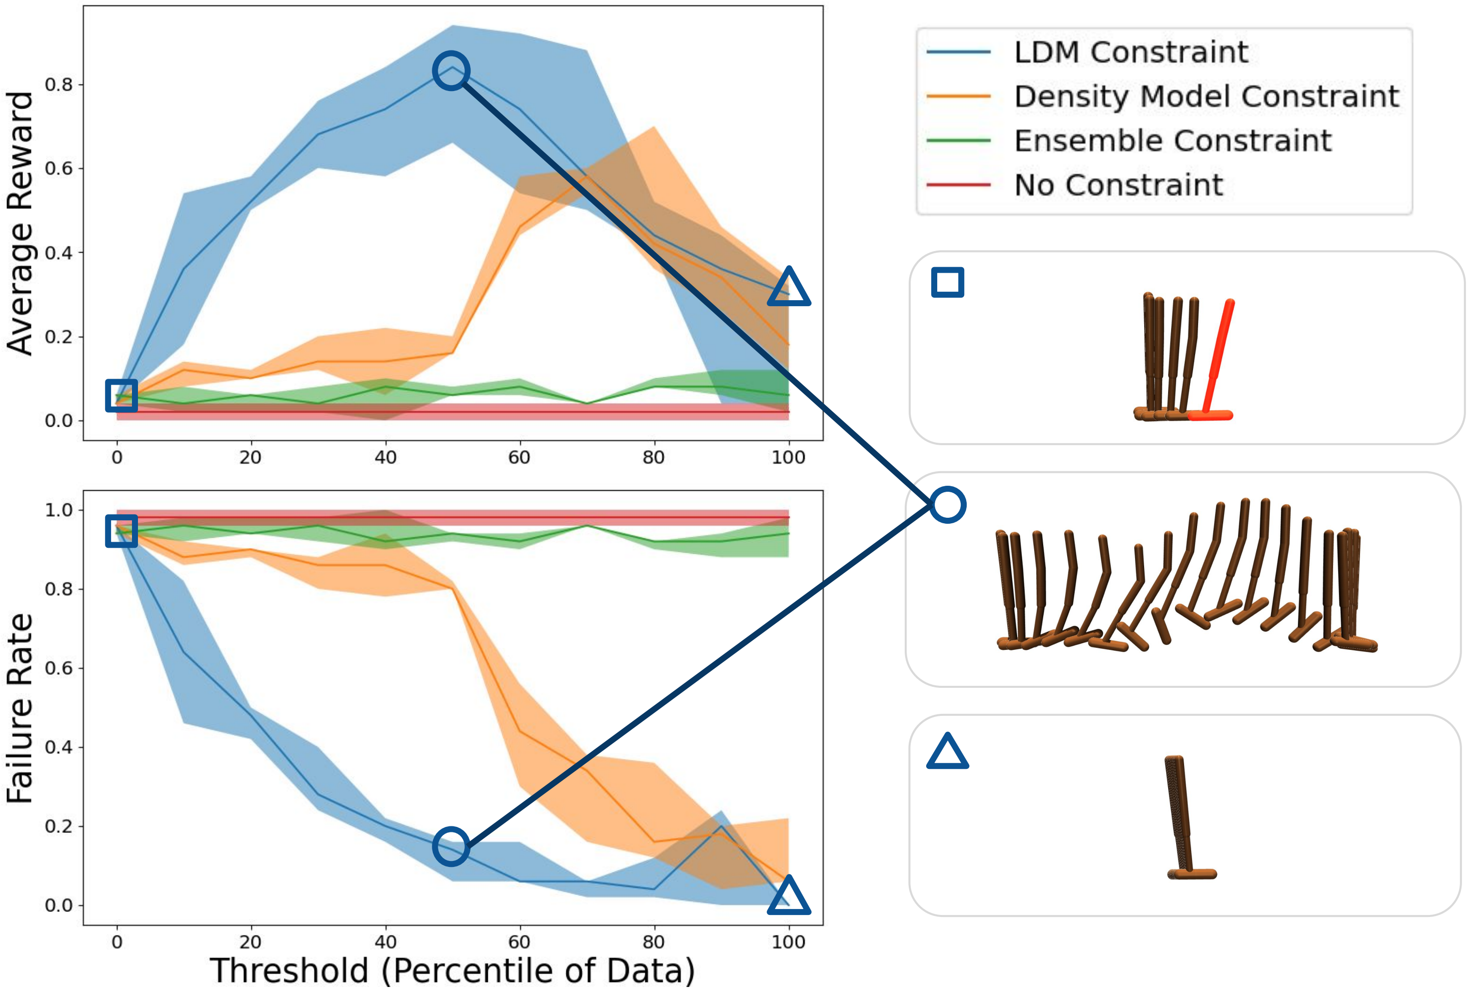

Average Reward
0.0
0.2
0.4
0.6
0.8
0
20
40
60
80
100
Failure Rate
0.0
0.2
0.4
0.6
0.8
1.0
Threshold (Percentile of Data)
LDM Constraint
Density Model Constraint
Ensemble Constraint
No Constraint

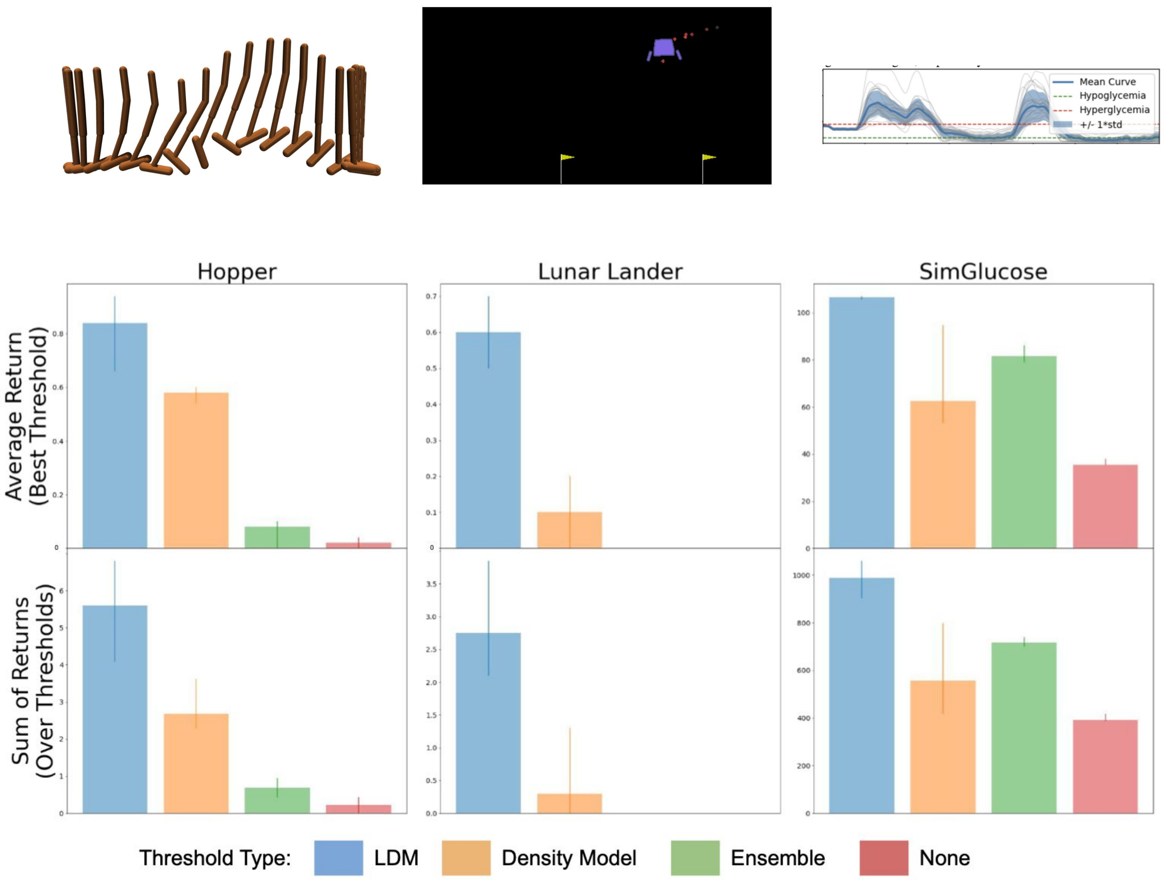
Mean Curve
Hypoglycemia
Hyperglycemia
+/- 1*std
Hopper
Lunar Lander
SimGlucose
Average Return
(Best Threshold)
Sum of Returns
(Over Thresholds)
Threshold Type:
LDM
Density Model
Ensemble
None

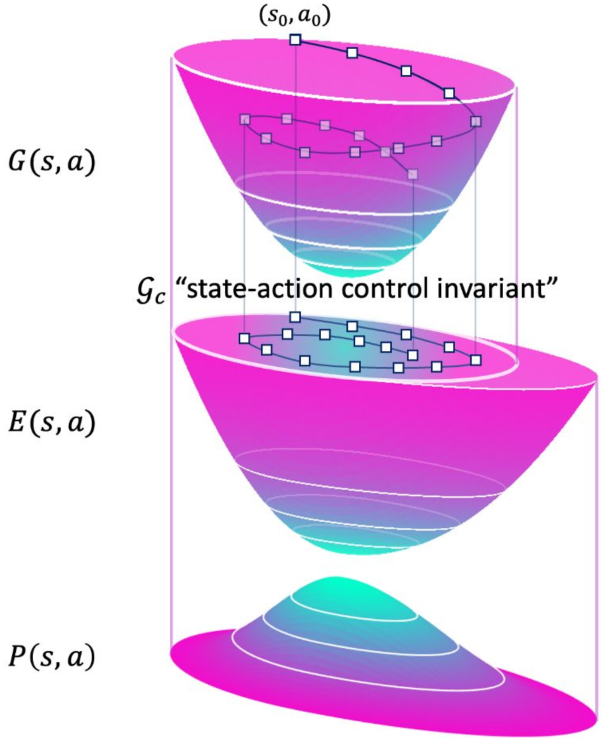
(s_0, a_0)
$G(s,a)$
$\mathcal{G}_c$ “state-action control invariant”
$E(s,a)$
$P(s,a)$

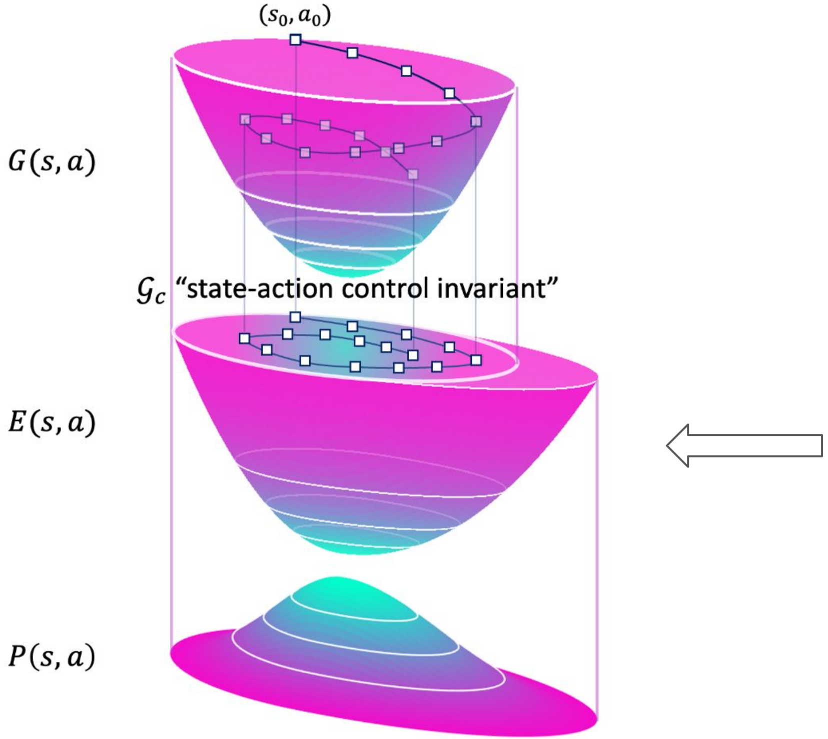

Single-step constraint
Can fail even in simple settings

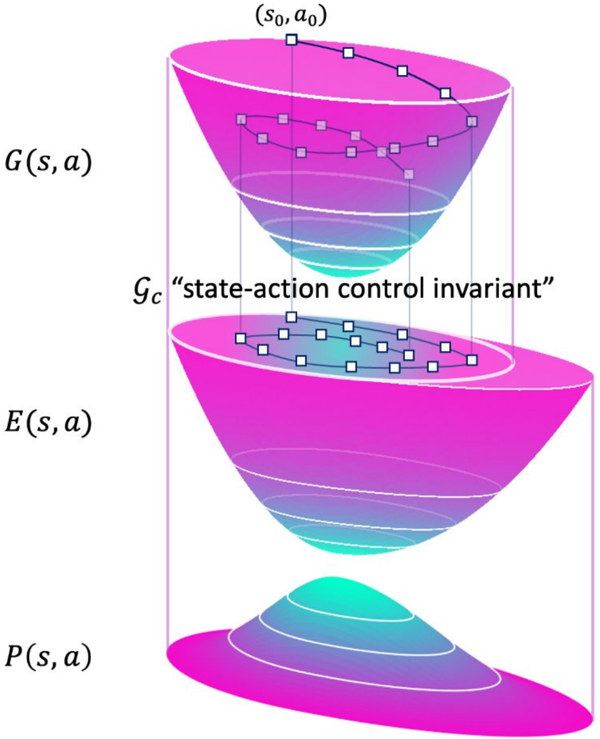

Keeps agent in-distribution over a long horizon

Single-step constraint
Can fail even in simple settings