Probabilistic Inference of Simulation Parameters via Parallel Differentiable Simulation

Eric Heiden, Christopher E. Denniston, David Millard,

Fabio Ramos, Gaurav S. Sukhatme

Real Robot Measurements

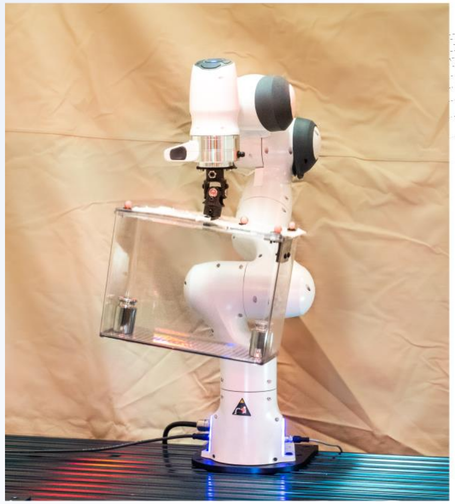

Parallel Differentiable Simulations

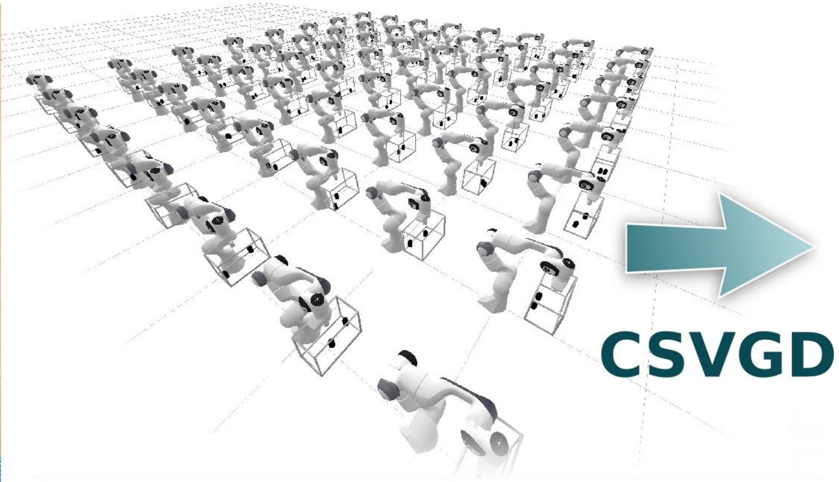

Distribution over Simulation Parameters

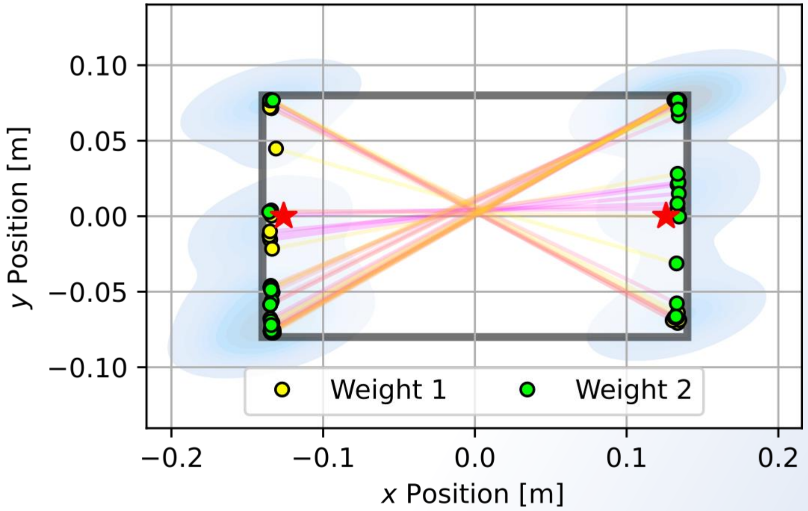

https://uscresl.github.io/prob-diff-sim

ROBOTIC EMBEDDED SYSTEMS LABORATORY
USC University of Southern California
NVIDIA

Probabilistic Parameter Inference

Infer posterior $p(\theta|D_{\mathcal{X}})$ over simulation parameters $\theta \in \mathbb{R}^M$ and a set of trajectories $D_{\mathcal{X}}$ via Bayes' rule:

$$p(\theta|D_{\mathcal{X}}) \propto p(D_{\mathcal{X}}|\theta)\, p(\theta)$$

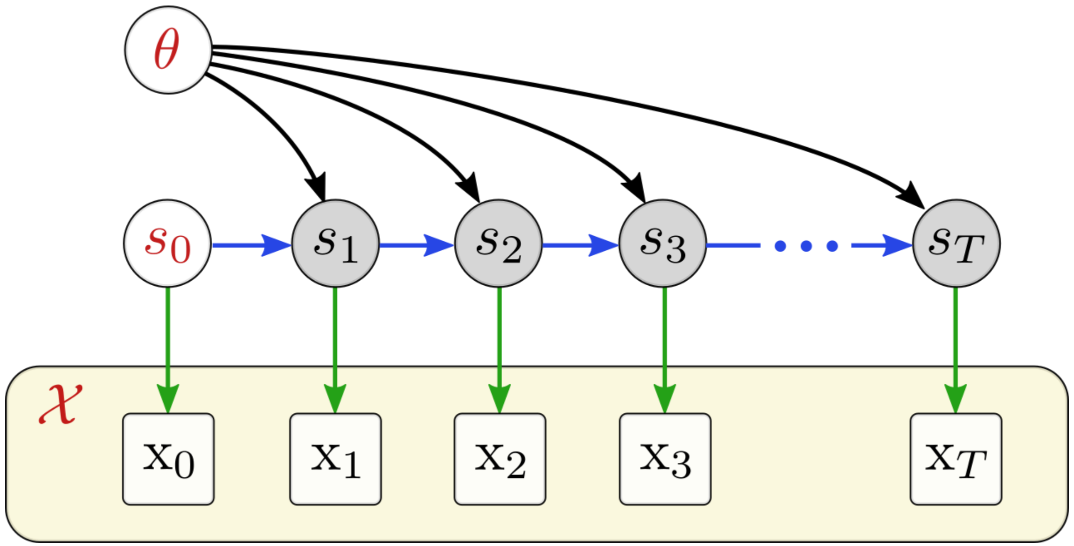

Hidden Markov Model:

initial state s_0,
observation function f_{obs} ,
simulation function f_{sim}

$$\mathcal{X} = f_{\text{obs}}([s]_{t=1}^T)$$

$$f_{\text{sim}}(\theta, s_0) = [s]_{t=1}^T$$

$$D_{\mathcal{X}}^{\text{sim}} = \left[f_{\text{obs}}\left(f_{\text{sim}}(\theta, s_0^{\text{real}})\right)\right]$$

Objective: minimize KL divergence

$$d_{\text{KL}}\left[p\left(D_{\mathcal{X}}^{\text{sim}}|\theta^{\text{sim}}\right) p\left(\theta^{\text{sim}}\right) \parallel p\left(D_{\mathcal{X}}^{\text{real}}|\theta^{\text{real}}\right) p\left(\theta^{\text{real}}\right)\right]$$

simulation reality

Stein Variational Gradient Descent

- Approximates probability distributions through particles $q(\theta|D_{\mathcal{X}})$
- Particles are moved to steepest descent direction to reduce KL divergence between $q(\theta|D_{\mathcal{X}})$ and $p(\theta|D_{\mathcal{X}})$ via

$$\nabla_\theta \log p(\theta|D_{\mathcal{X}}) = \frac{\nabla_\theta p(\theta|D_{\mathcal{X}})}{p(\theta|D_{\mathcal{X}})}$$

  in reproducing kernel Hilbert space
- Parallelizable, efficient at high dimensional parameter distributions

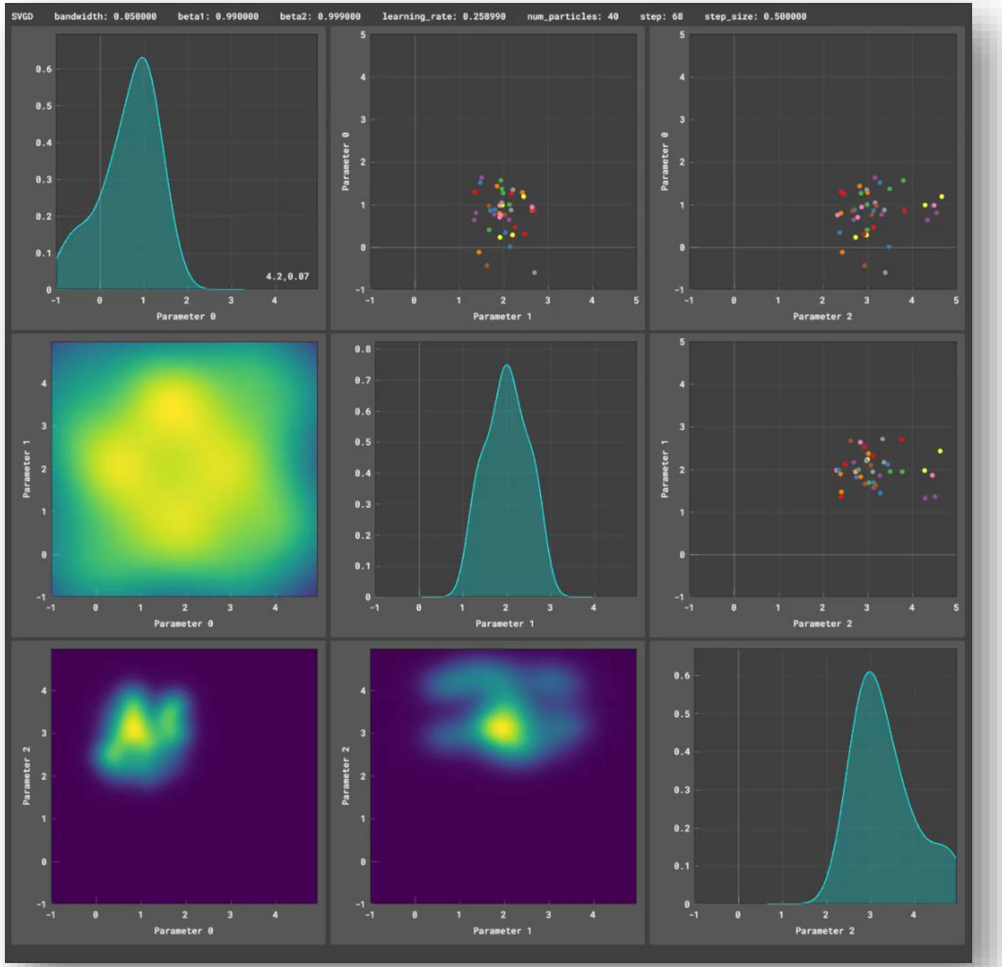

Liu, Wang. Stein Variational Gradient Descent: A General Purpose Bayesian Inference Algorithm. NeurIPS 2016.

Multiple Shooting

Break up trajectory into shooting windows

Augment parameter vector by start states s_t^s of shooting windows

Impose defect constraints to ensure continuity

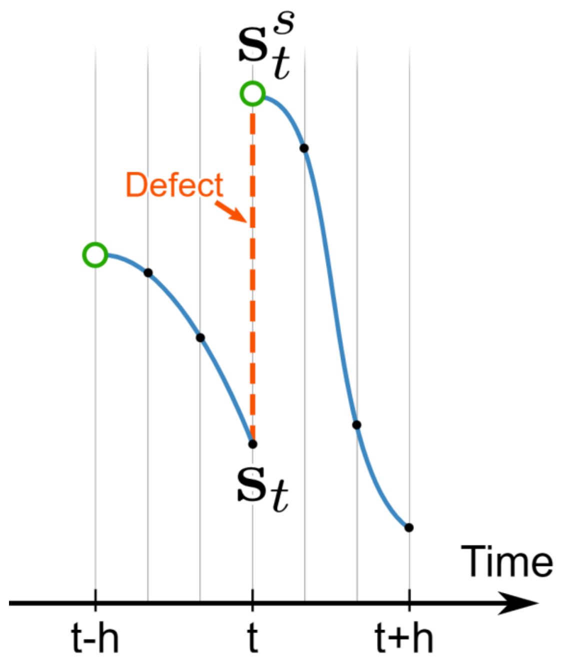

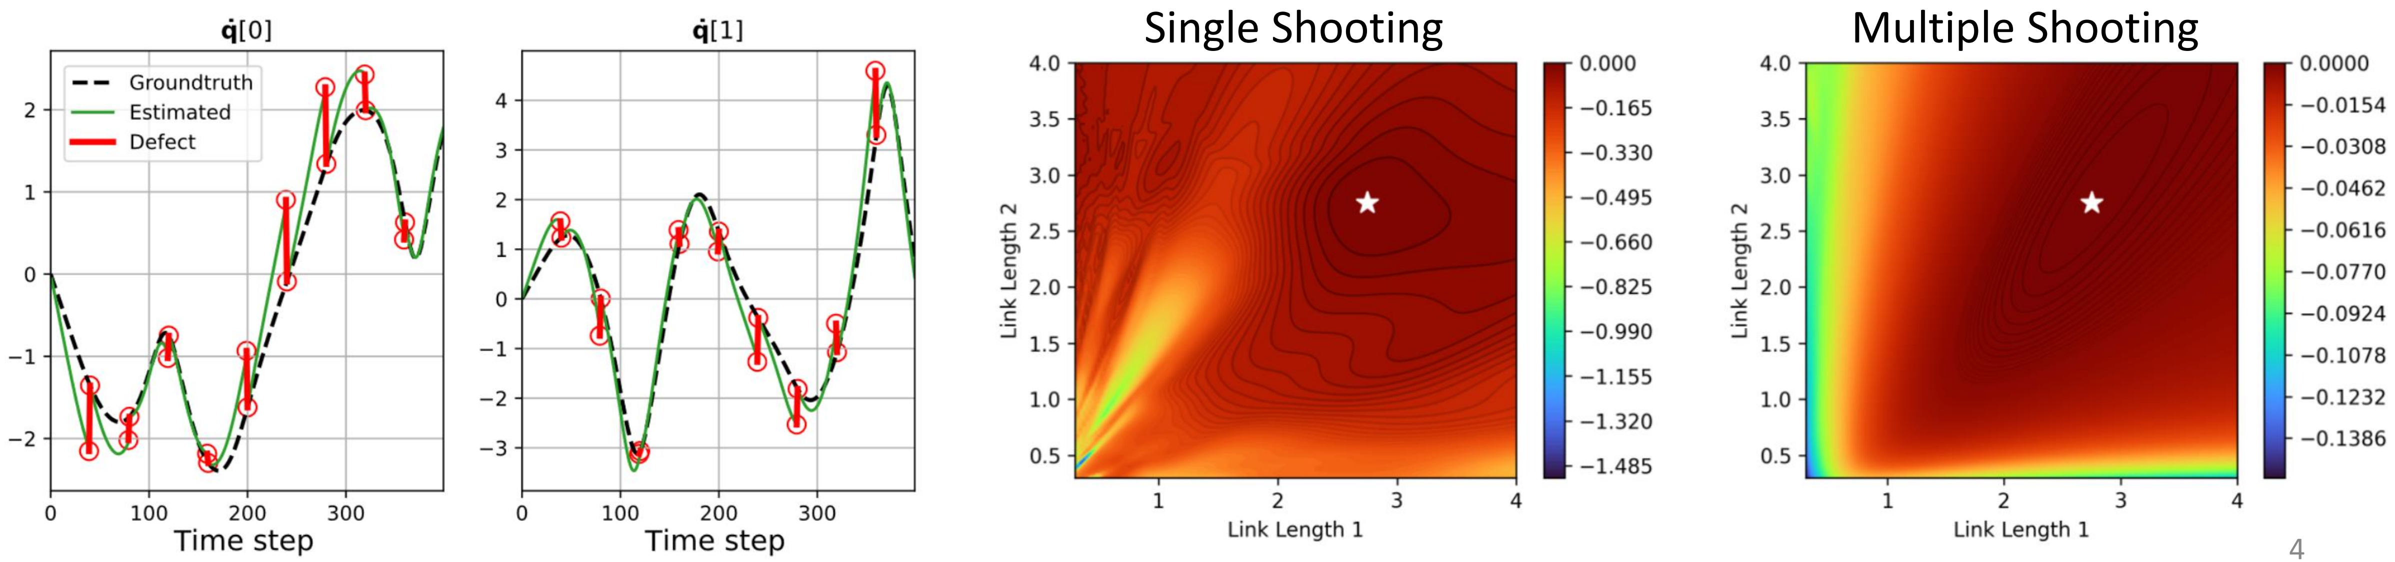

Constrained SVGD (CSVGD)

Combine SVGD and the Modified Differential Method of Multipliers (MDMM) to consider hard constraints $g(\theta)$ in the estimation:

- Parameter limits: $g_{\text{lim}}(\theta) = \text{clamp}(\theta, \theta_{\text{min}}, \theta_{\text{max}}) - \theta$
- Defect constraints: $g_{\text{def}}(\theta) = \|\mathrm{s}_t^s - \mathrm{s}_t\|^2 / \sigma_{\text{def}}^2$

Leverages parallel differentiable simulation on the GPU to evaluate $\nabla_\theta \log p(\theta | D_\mathcal{X})$ for all particles

$$\dot{\theta} = \phi(\theta) - \lambda \frac{\partial g(\theta)}{\partial \theta} - c g(\theta) \frac{\partial g(\theta)}{\partial \theta}, \qquad \dot{\lambda} = g(\theta)$$

CSVGD

$$\phi(\cdot) = \frac{1}{N} \sum_{j=1}^{N} \left[k(\theta_j, \theta) \nabla_{\theta_j} \log p(D_\mathcal{X} | \theta)\, p(\theta) + \nabla_{\theta_j} k(\theta_j, \theta) \right]$$

SVGD

where $k(\cdot,\cdot)$ is positive definite kernel (RBF), N is number of particles

https://uscresl.github.io/prob-diff-sim

Analytical Parameter Distribution

- Infer parameters drawn from a known 2D Gaussian
- Double pendulum system where the 2 link lengths are inferred

Posterior

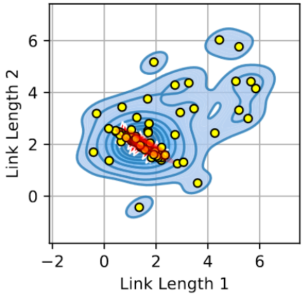

(a) Emcee

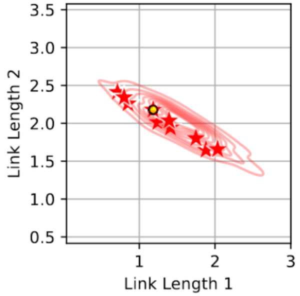

(b) CEM

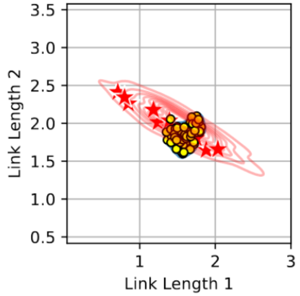

(c) SGLD

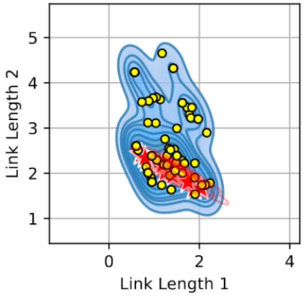

(d) NUTS

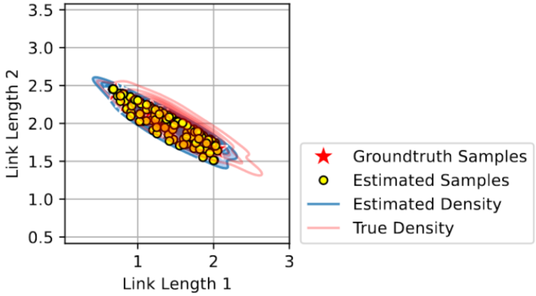

(e) SVGD

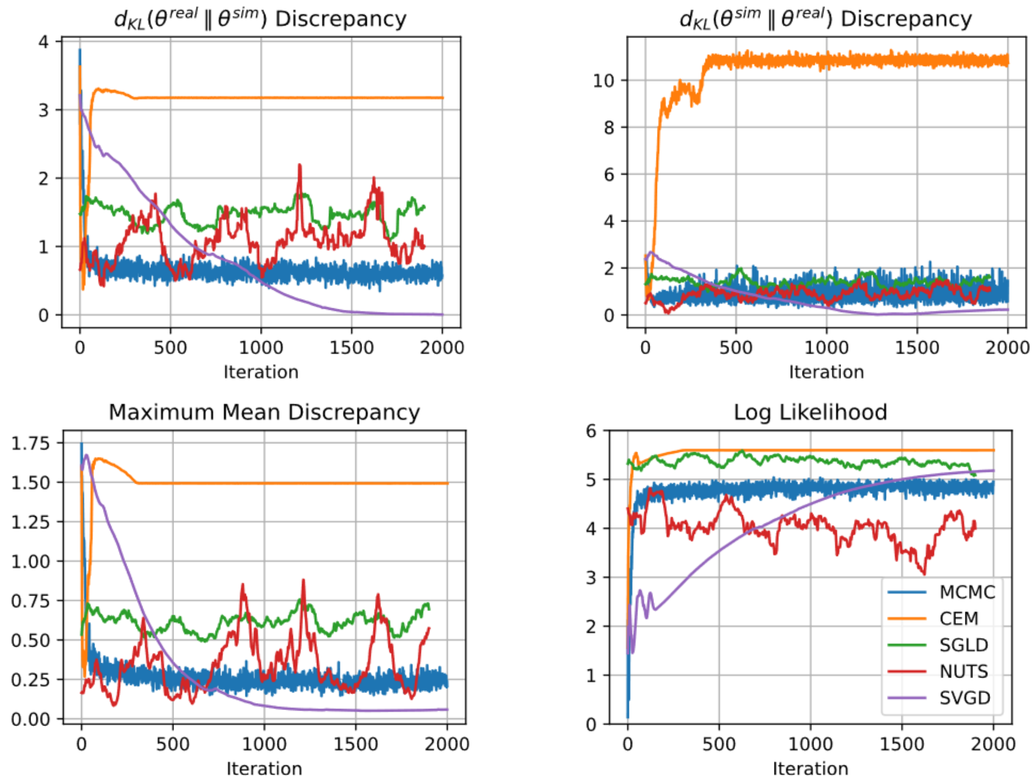

Double Pendulum

Infer 11 parameters from a real double pendulum

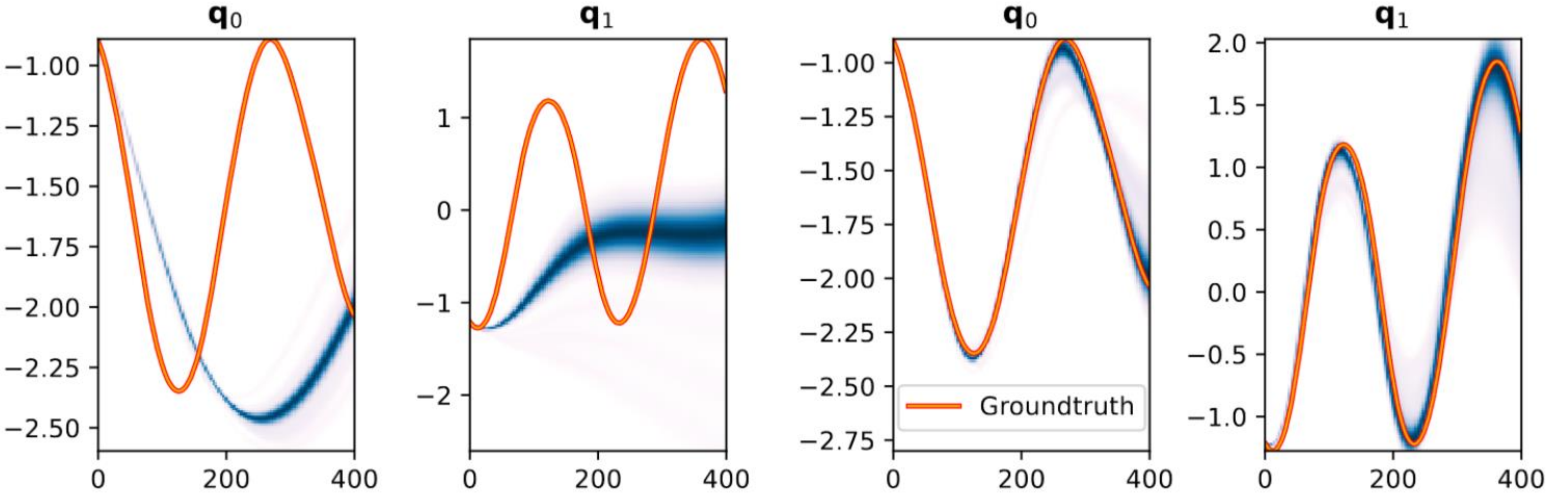

(a) SVGD (b) CSVGD

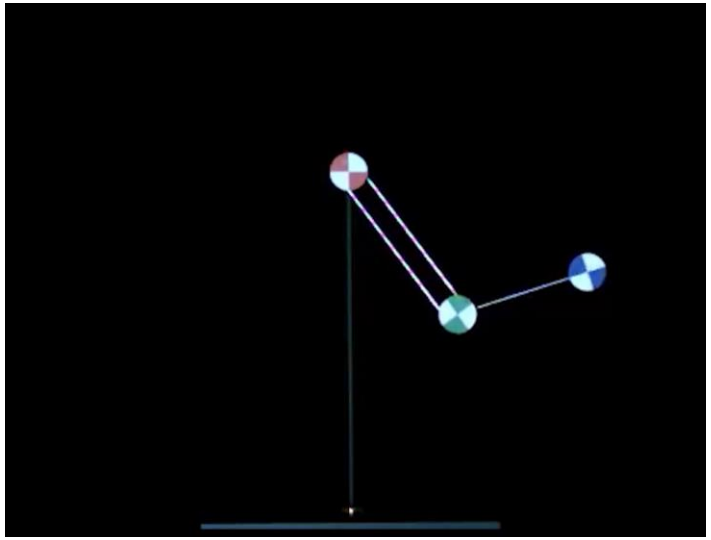

Link	Parameter	Minimum	Maximum
Link 1	Mass	0.05 kg	0.5 kg
	I_{xx}	0.002 $\mathrm{kg\,m^2}$	1.0 $\mathrm{kg\,m^2}$
	COM x	−0.2 m	0.2 m
	COM y	−0.2 m	0.2 m
	Joint friction	0.0	0.5
Link 2	Length	0.08 m	0.3 m
	Mass	0.05 kg	0.5 kg
	I_{xx}	0.002 $\mathrm{kg\,m^2}$	1.0 $\mathrm{kg\,m^2}$
	COM x	−0.2 m	0.2 m
	COM y	−0.2 m	0.2 m
	Joint friction	0.0	0.5

CSVGD yields more accurate predictions than SVGD, Monte-Carlo algorithms, and BayesSim

Underactuated Mechanism

Panda robot arm with an object of unknown mass distribution attached to its end-effector through a universal joint

Infer locations of 2 weights from box motion

CSVGD

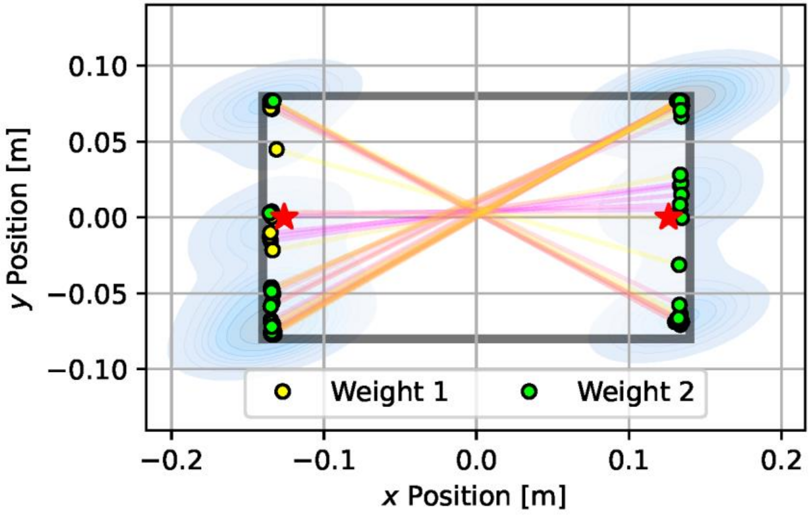

SVGD

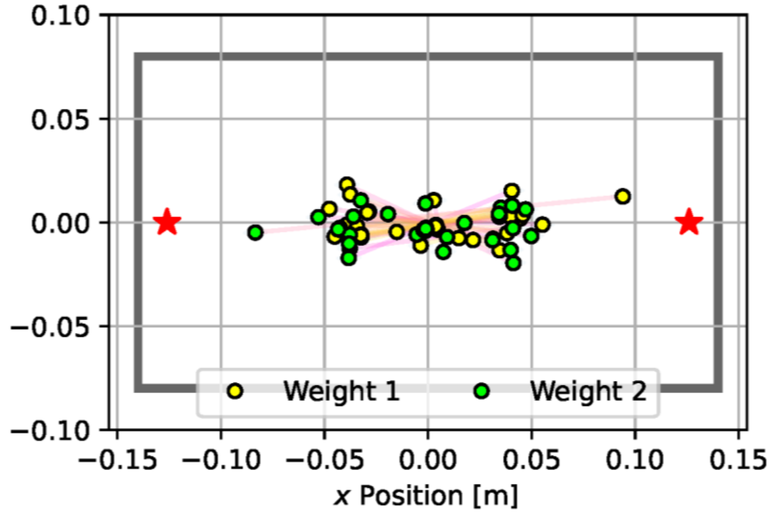

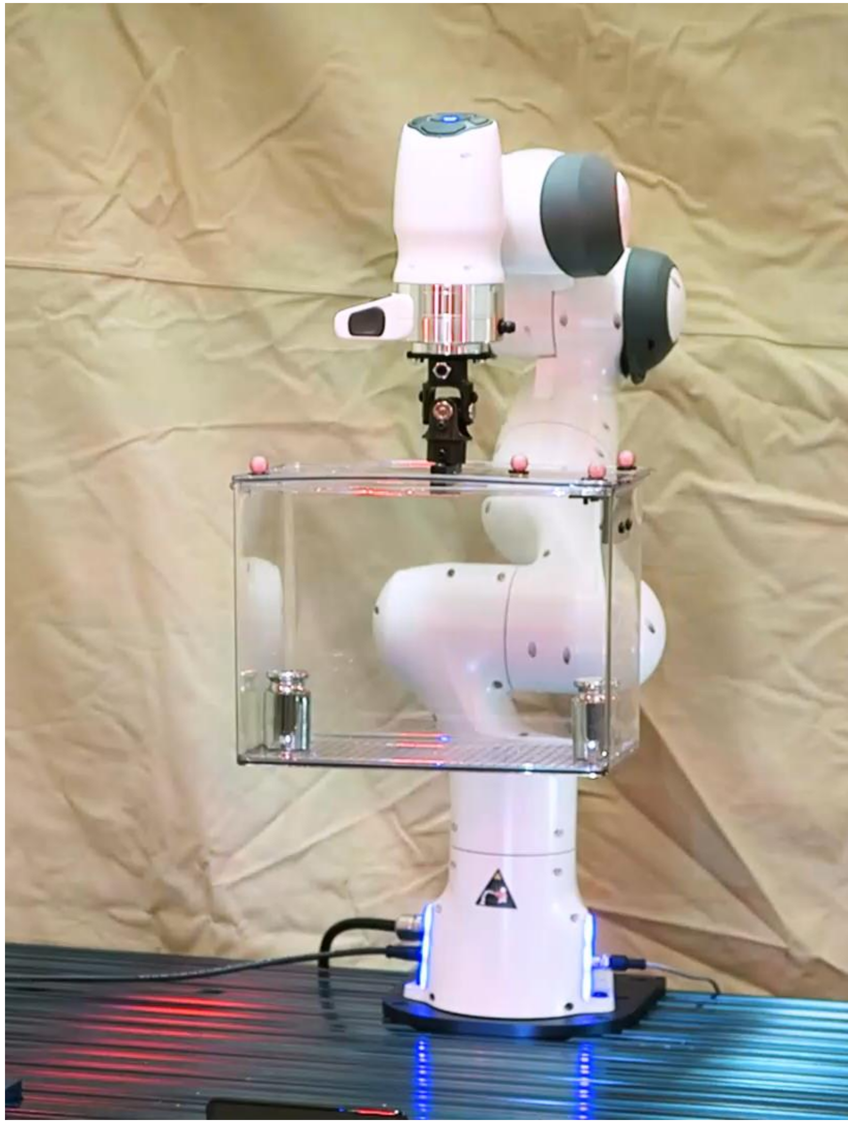

Probabilistic Inference of Simulation Parameters via Parallel Differentiable Simulation

Eric Heiden, Christopher E. Denniston, David Millard,

Fabio Ramos, Gaurav S. Sukhatme

Real Robot Measurements

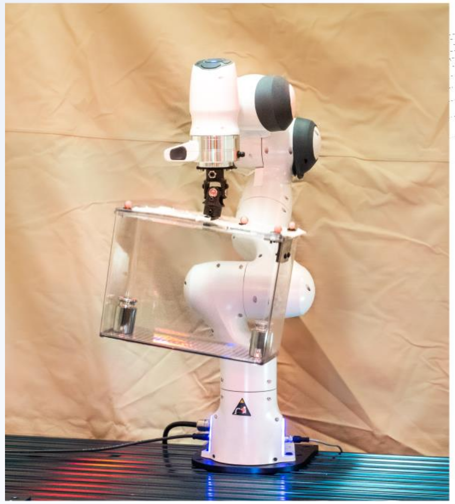

Parallel Differentiable Simulations

Distribution over Simulation Parameters

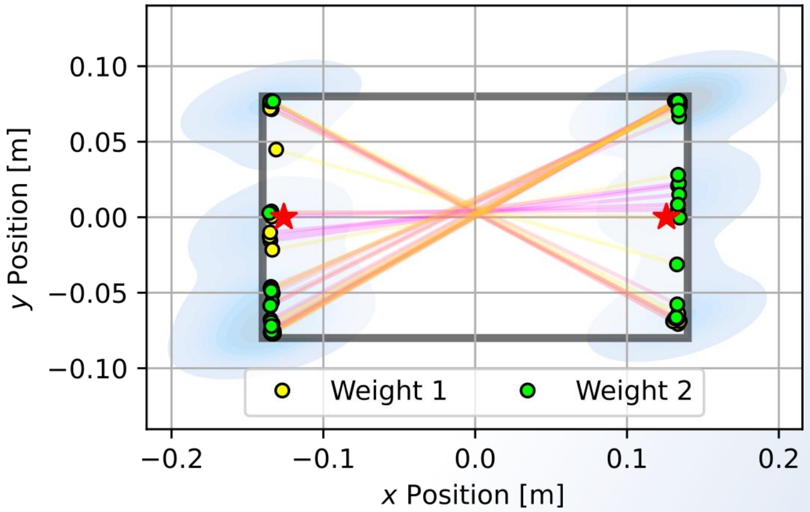

https://uscresl.github.io/prob-diff-sim

ROBOTIC EMBEDDED SYSTEMS LABORATORY

USC University of Southern California

NVIDIA