Differentiable Physics Simulations with Contacts: Do They Have Correct Gradients w.r.t. Position, Velocity and Control?

Yaofeng Desmond Zhong[1]

Jiequn Han[2]

Georgia Olympia Brikis[1]

[1]Siemens Technology

[2]Flatiron Institute

SIEMENS
Ingenuity for life

FLATIRON INSTITUTE

ICML
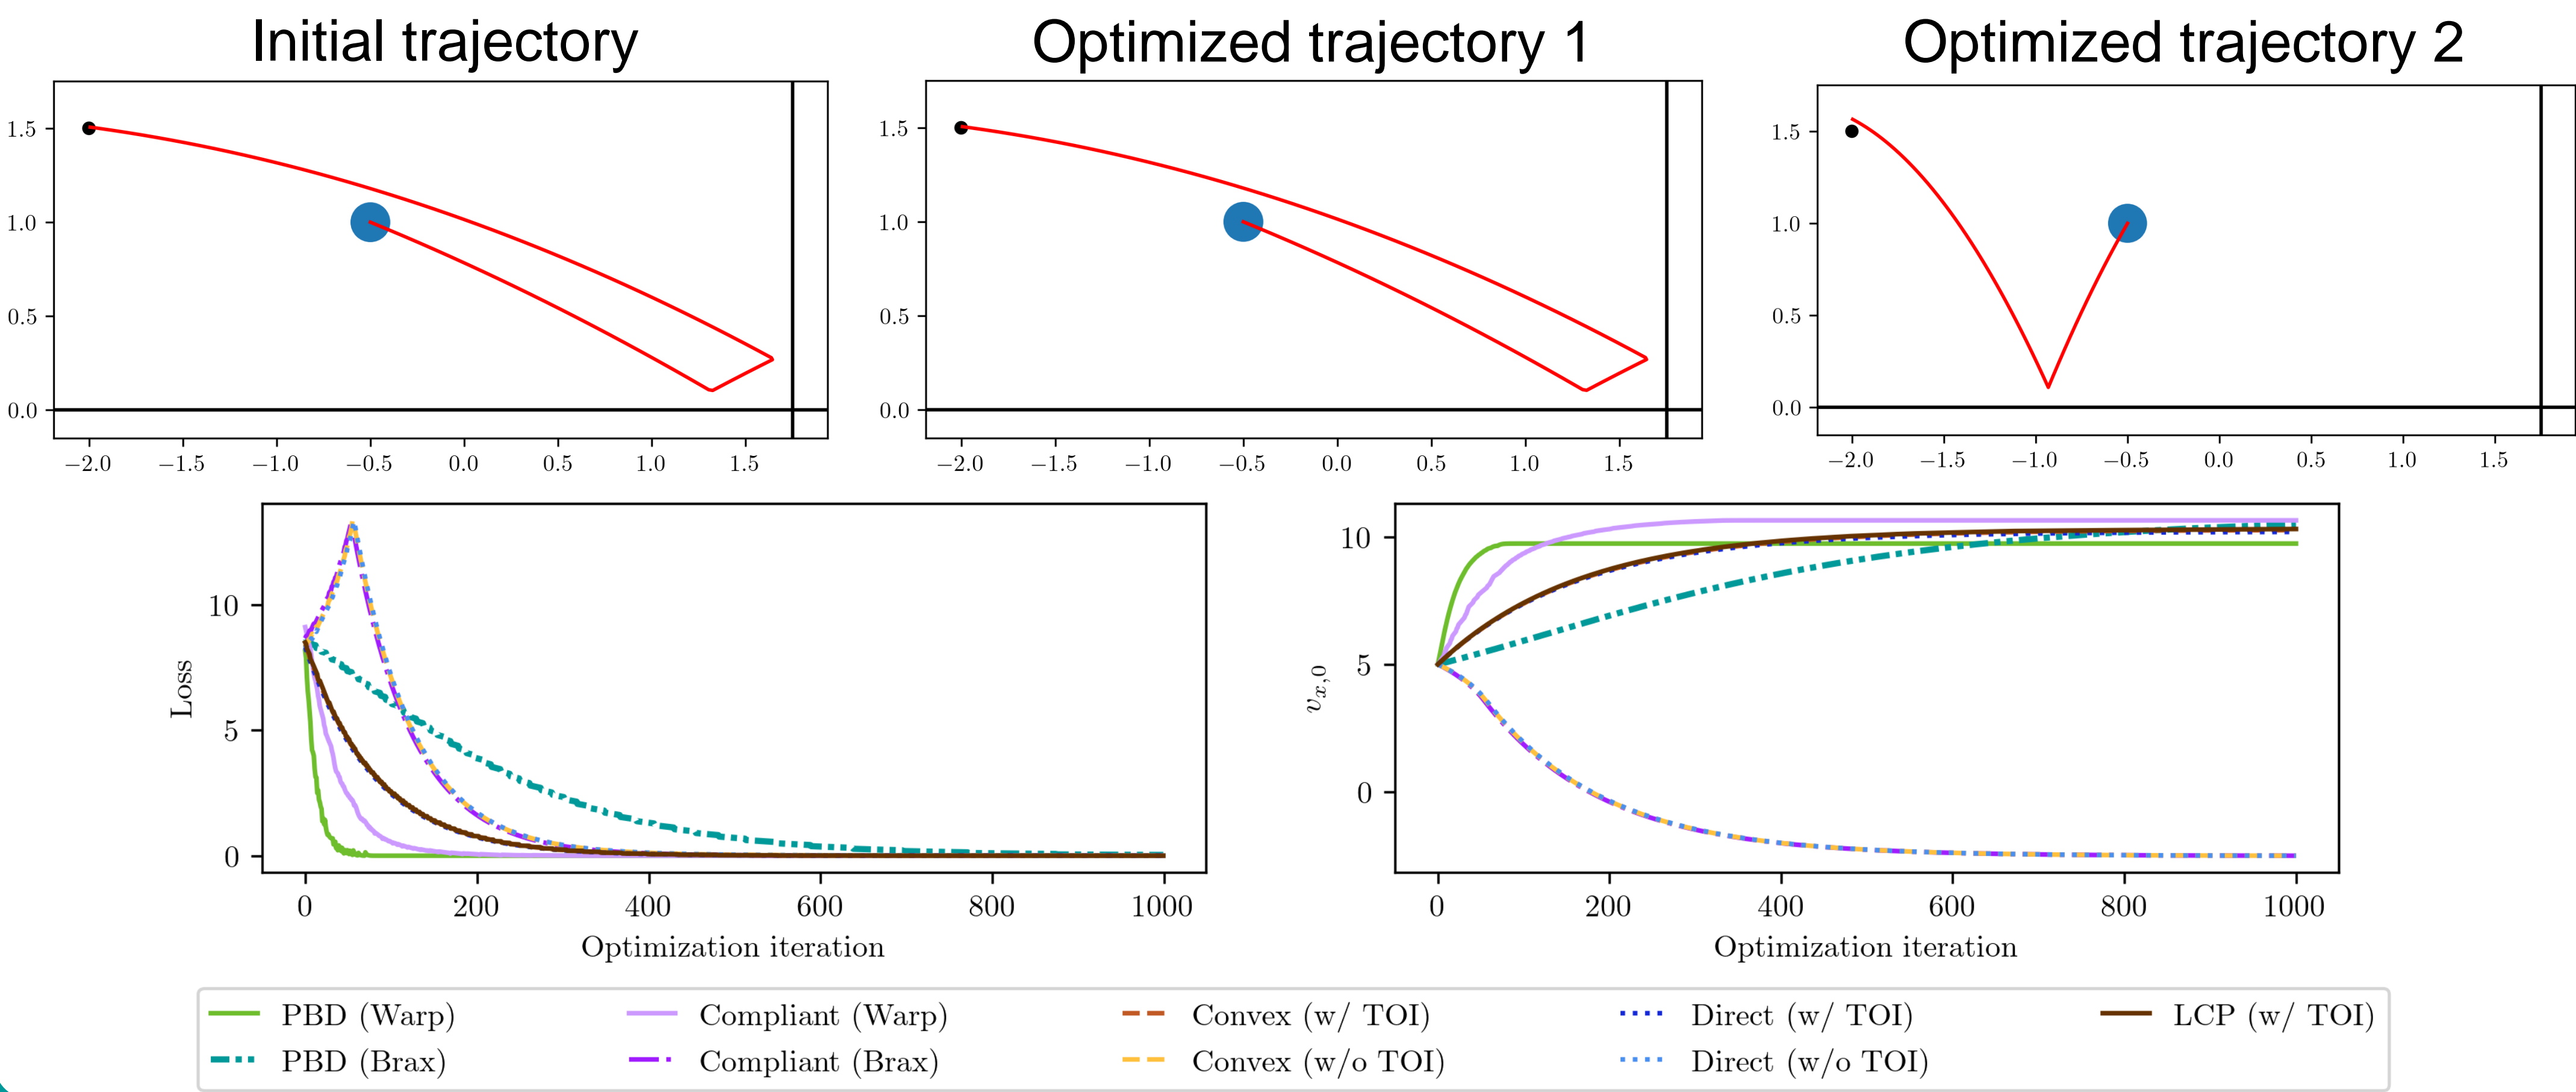International Conference On Machine Learning

- Differentiable physics simulations make it easy to use gradient-based method for learning and control tasks and have huge potential in solving control and design problems.
- An increasing number of differentiable simulators are developed and open sourced.
- However, their performance has not been compared and benchmarked.

Differentiable Contact Models

Contact Formulations		impulses	Implementation choices
Linear complementarity problems (LCPs) (w. TOI)			NimblePhysics
Convex Optimizations	w. TOI	velocity	DiffCoSim
	w/o TOI		
Direct Velocity Impulse	w. TOI		DiffTaichi
	w/o TOI		
Position-based dynamics (PBD)		position	Warp/Brax
Compliant models		force	Warp/Brax

- TOI: time of impact, proposed by DiffTaichi to improve gradient calculation.

Task 2: Optimizing the Initial Velocity to Hit a Target

Initial trajectory   Optimized trajectory 1   Optimized trajectory 2

PBD (Warp) — Compliant (Warp) — Convex (w/ TOI) — Direct (w/ TOI) — LCP (w/ TOI)
PBD (Brax) — Compliant (Brax) — Convex (w/o TOI) — Direct (w/o TOI)

Task 3: Learning Optimal Control

- Some implementations fail to learn the optimal control sequence.
- Surprisingly, none of the calculated gradients match the analytical gradients.

PBD (Warp) — Compliant (Warp) — Convex (w/ TOI) — Direct (w/ TOI) — LCP (w/ TOI)
PBD (Brax) — Compliant (Brax) — Convex (w/o TOI) — Direct (w/o TOI) — Analytical Optimal Loss/Control