Convergent Planning and Control of Legged Robots

James Zhu, J. Joe Payne, Nathan J. Kong, Aaron M. Johnson

Motivation

Planning trackable closed-loop trajectories for systems that are underactuated and/or undergo hybrid events is difficult due to system dynamics that can dominate feedback control effort.

Underactuated systems have at least one direciton in state space that can not be controlled while hybrid events induce discrete, unbounded disturbances.

In this work, we present a method to quantify the worst-case divergence of initial errors for a closed-loop trajectory with an LQR tracking controller, and plan a trajectory that optimizes over this worst-case value.

Convergence Measure

Our goal is to characterize the evolution of errors over the entirety of a trajectory.

Within the continuous domains of the trajectory, the variation equation that captures the error evolution is the familiar: $A - BK$

For legged robots, it is key to also establish the error evolution across a hybrid event.

This behavior can be discribed with a similar linear operator called the saltation matrix: Ξ

Integrating these operators over a trajectory gives the linear approximation of the trajectory's error evolution, called the fundamental solution matrix:

$$\Phi = \bar{A}_N \ldots \Xi_{(2,3)} \bar{A}_2 \Xi_{(1,2)} \bar{A}_1$$

$$\text{where} \quad \bar{A}_i = (A_i - B_i K_i)$$

The two norm of the fundamental solution matrix is the convergence measure, which captures the worst case error divergence direction.

$$\chi = ||\Phi||_2$$

This material is based upon work supported by the National Science Foundation under grant #CMMI-1943900.

Convergent iLQR

This work extends the hybrid iLQR algorithm [1], which is a contact-implicit trajectory optimization method for hybrid systems like leged robots (Fig. 1).

We compare the standard (vanilla) trajectory with a convergent trajectory where the convergence measure is included in the cost function.

This method explicitly plans to reduce the worst-case outcome of the closed-loop trajectory.

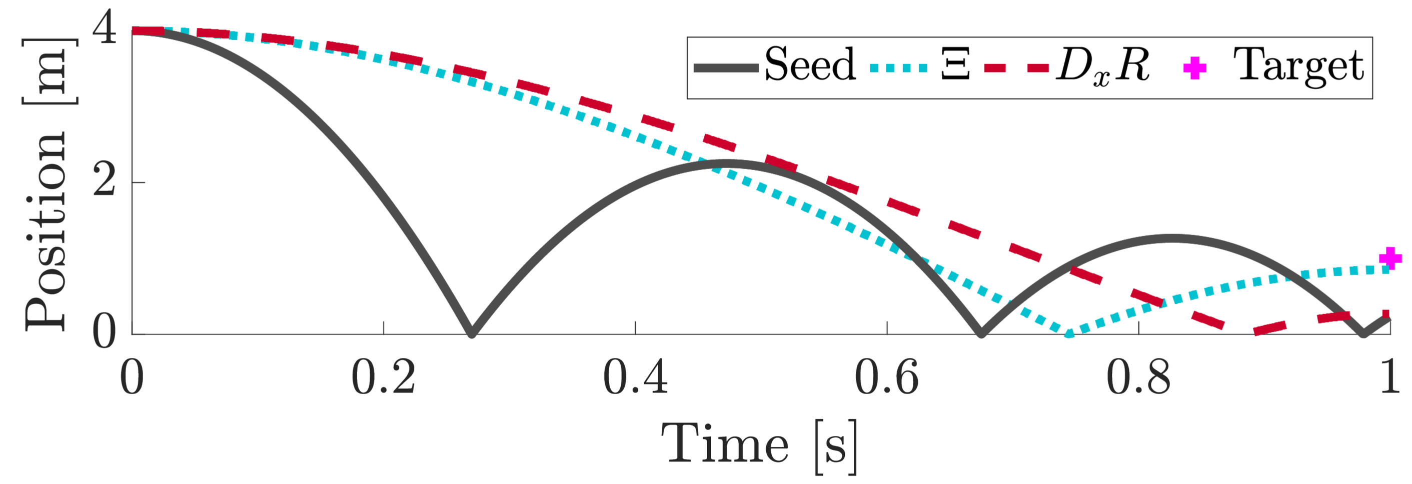

Fig. 1: The hybrid iLQR algorithm is able to optimize the trajectory of a bouncing ball where the initial seed trajectory consists of 3 bounces. The algorithm is able to find the optimal number of bounces to reach the target state.

Example and Results

In this example, we use a half-quadruped model that is symmetrical about the sagittal plane that has a parallel torsion spring at the knee joint.

At 6 diferent initial error magnitudes, 100 paired simulations were run for the vanilla and convergent trajectories.

We found that the convergent trajectory consistently displayed superior tracking performance and robustness, while actually using less feedback effort than the vanilla trajectory.

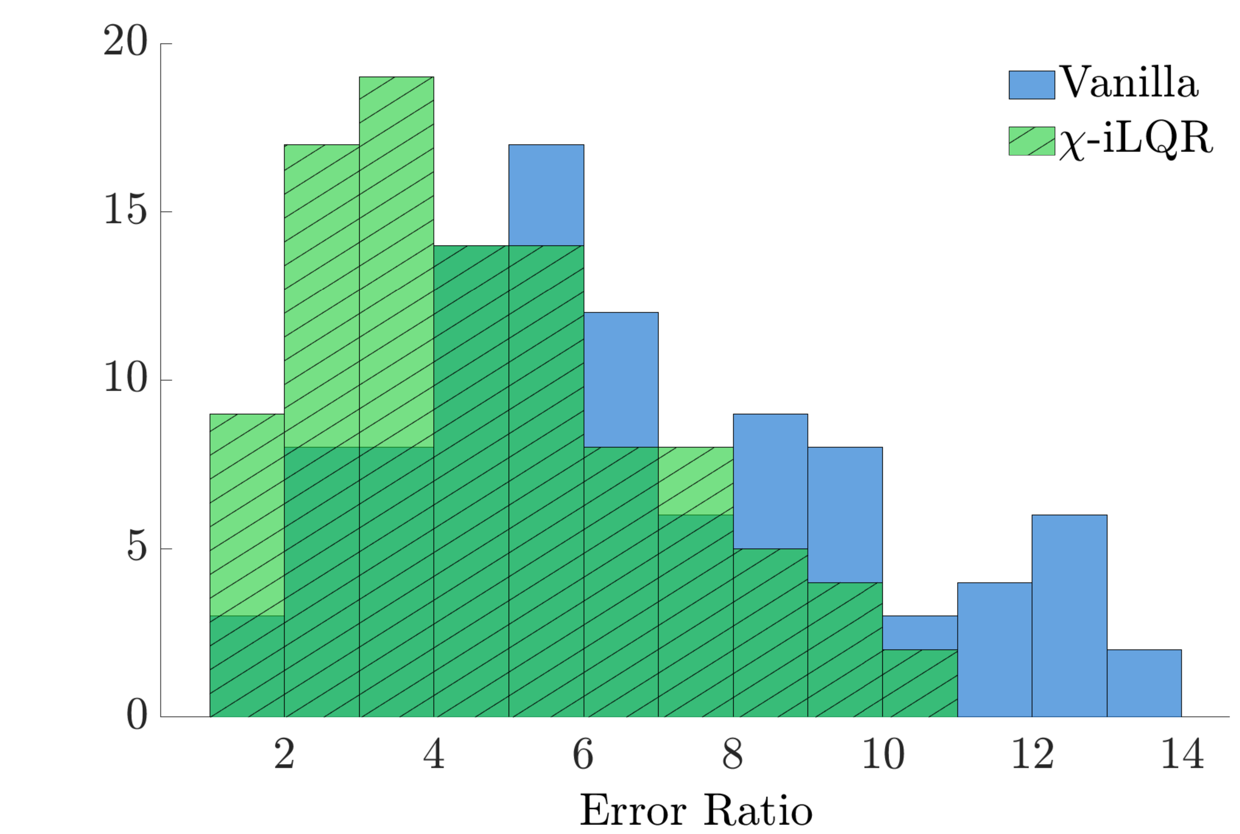

Fig. 2: Results of 100 trials for vanilla and convergent trajectories with an initial error covariance weighting of 1e-4. The error ratio is the ratio between the size of the final error and the size of the initial error. The convergent trajectory reduces the mean error ratio by 28%.

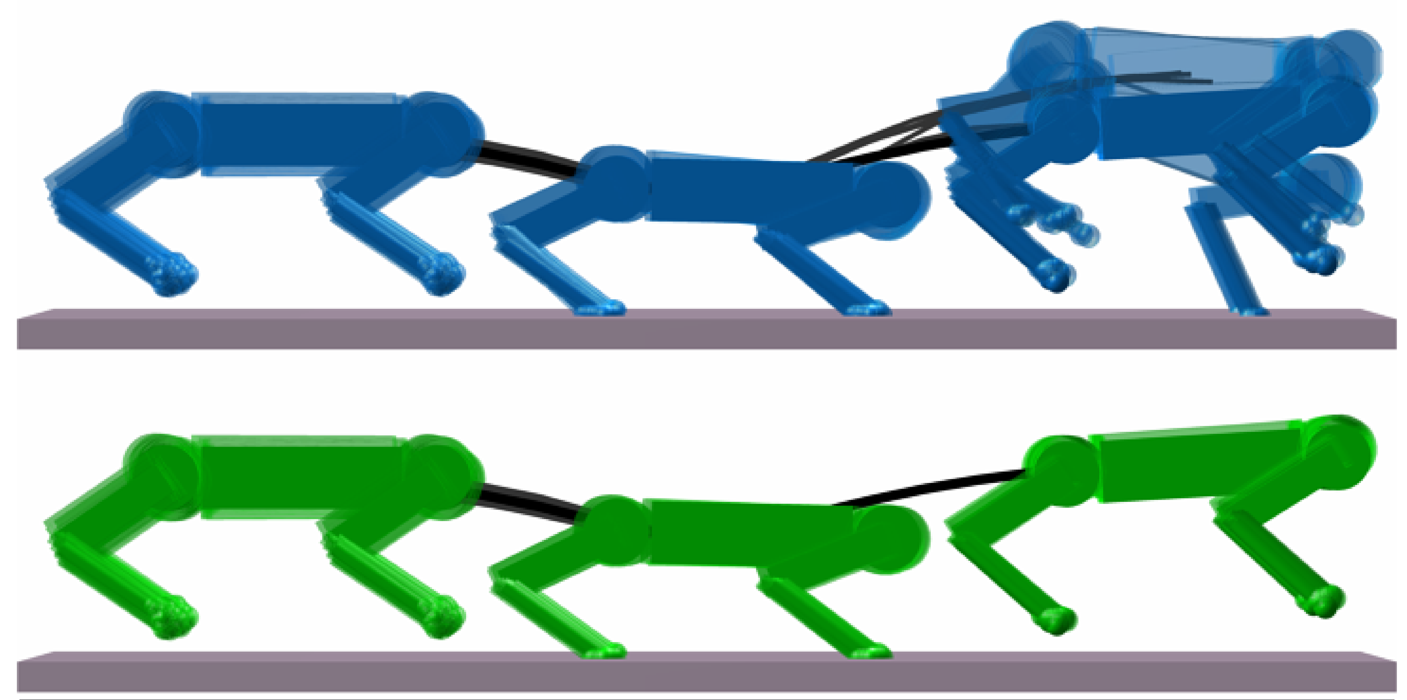

Fig. 3: Over larger initial perturbations, the tracking performance of the vanilla trajectory begins to deteriorate, while the convergent trajectory performs significantly better.

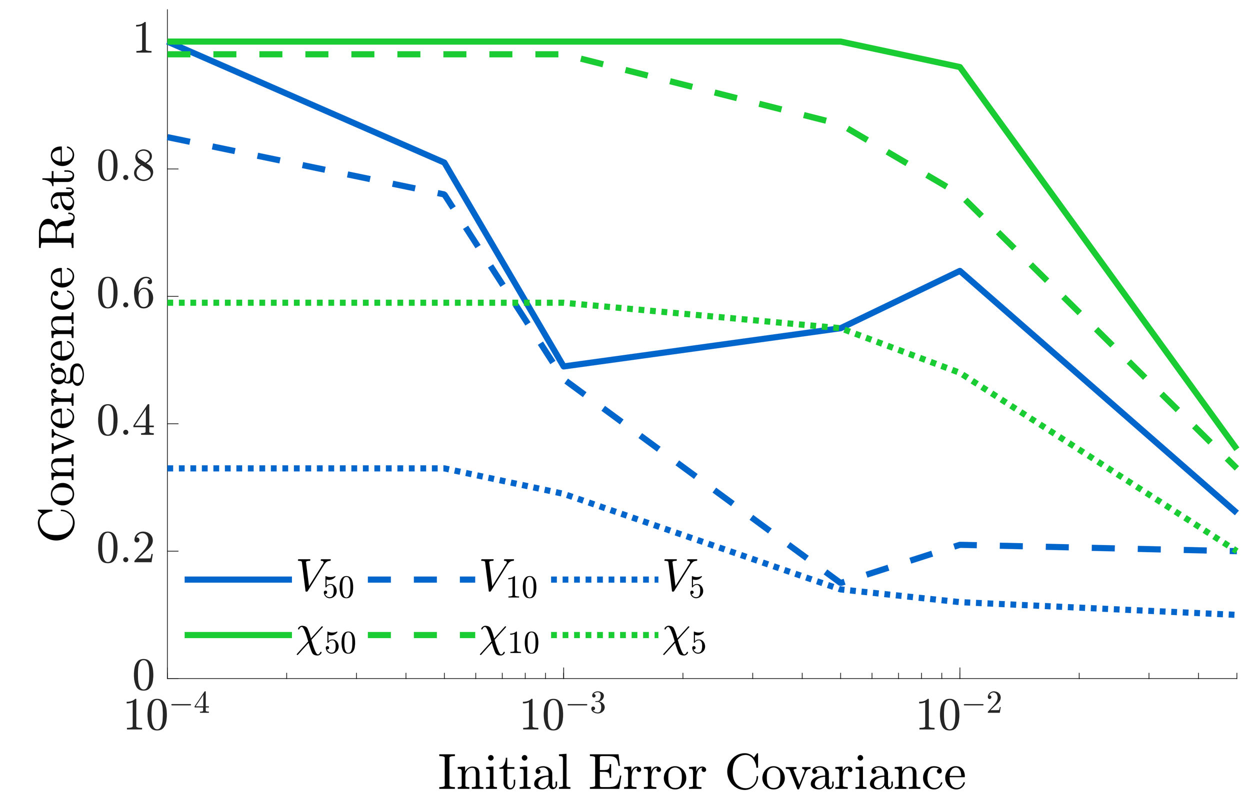

Fig. 4: Across a wide range of initial error magnitudes, the convergent trajectory performs consistently better than the vanilla trajectory. Each pair of plot lines shows each trajectories' rate of error ratios below 50, 10, and 5 respectively.

References. [1] N. J. Kong, G. Council, and A. M. Johnson, "iLQR for piecewise-smooth hybrid dynamical systems," in IEEE Conference on Decision and Control, 2021.

Robomechanics Lab