Locomotion: An Emergent Risk-mitigating Behavior in Uncertain Environments

Jacob Hackett*, Dylan Epstein-Gross
Mechanical Engineering
FAMU-FSU College of Engineering
Tallahassee, FL, USA
*jeh15@my.fsu.edu

Monica Daley
Ecology & Evolutionary Biology
University of California Irvine
Irvine, CA, USA
madaley@uci.edu

Christian Hubicki
Mechanical Engineering
FAMU-FSU College of Engineering
Tallahassee, FL, USA
hubicki@eng.famu.fsu.edu

I. Background

Energetics is a powerful framework for explaining locomotion in animals [1]. However, the practicalities of real-world locomotion force us to reconcile energy minimization with other conflicting pressures (*e.g.* injury or predation). We seek a unifying mathematical framework that can include all pressures by encoding them as probabilistic sources of risk. We hypothesize that behaviors of real-life locomotors can emerge from a mathematical agent that can both learn and optimally mitigate risks in its environment.

II. Methods

Our proposed method creates a trajectory optimization problem that encodes failure probabilities, that are state and input dependent, as constraints [2]. We transcribe the optimization problem into a nonlinear program using direct collocation. This formulation can be solved quickly with existing large scale NLP solvers (*i.e.* IPOPT). The result of this optimization is a motion plan that either constrains or minimizes risk of failure (*e.g.* motors overheating, critical contact points slipping away, or running out of battery).

Further, we create a learning framework that correlates experienced failures to candidate system variables; effectively learning sources of risk. Specifically, we regress a convex linear piecewise function to approximate these probabilities. By looping our planning and learning models we have created an agent that learns and acts in its environment on-the-fly (*i.e.* a single loop taking 2.77 seconds); a framework analogous to model predictive control.

III. Results & Discussion

We created two double-integrator toy problems. The first of which simulates an agent learning to traverse a danger zone to complete a task with an efficiency objective, and the second simulates an agent learning to track a moving energy source by minimizing its chance of failure (*i.e.* running out of energy).

In the first toy problem, we found individualistic behaviors emerging (*i.e.* risk-averse, risk-neutral, and risk-loving) as a result of its history of experienced failures. In the second toy problem, we found with certain parameters the agent successfully learns to track the energy source despite not being told it needs to. This gives us early evidence that energy economical behaviors and individualistic habits emerge from this risk learning and planning framework. We noticed these behaviors are parameter specific; if locomotion cost is made too expensive, instead of actively seeking the energy source, the agent minimizes its energy expenditure by not moving (exhibiting a "hibernation"-like strategy).

This work is funded by the FAMU-FSU College of Engineering, Department of Mechanical Engineering.

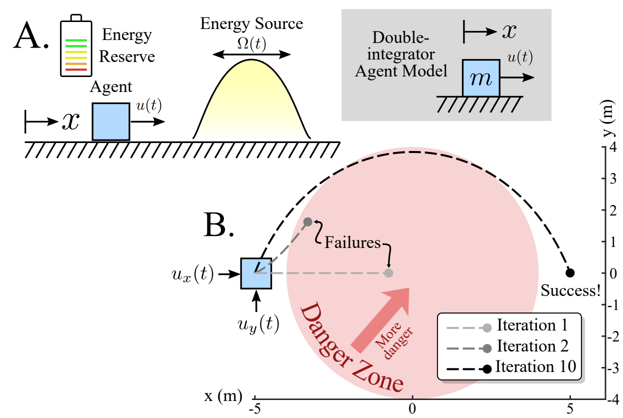

Fig. 1. Toy scenarios for risk-constrained and risk-minimizing locomotion.

IV. Conclusion

Preliminary results of our rapid iterative risk learning and planning framework appear to show emergent energy minimization and individualistic behaviors. Our aim is to apply this method to analyze biological locomoting species in their ecosystems, and enable robots to learn how to perform novel tasks on-the-fly.

References

[1] M. Srinivasan and A. Ruina, "Computer optimization of a minimal biped model discovers walking and running," *Nature*, vol. 439, no. 7072, pp. 72–75, Jan 2006. [Online]. Available: https://doi.org/10.1038/nature04113

[2] J. Hackett, W. Gao, M. Daley, J. Clark, and C. Hubicki, "Risk-constrained motion planning for robot locomotion: Formulation and running robot demonstration," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 3633–3640.