Deep Reinforcement Learning for Six Degree-of-Freedom Planetary Powered Descent and Landing

Brian Gaudet[*]

University of Arizona, Tucson Arizona, 85721, USA

Richard Linares[†]

Massachusetts Institute of Technology, Cambridge, MA, 02139, USA

Roberto Furfaro[‡]

University of Arizona, Tucson Arizona, 85721, USA

Future Mars missions will require advanced guidance, navigation, and control algorithms for the powered descent phase to target specific surface locations and achieve pinpoint accuracy (landing error ellipse < 5 m radius). The latter requires both a navigation system capable of estimating the lander's state in real-time and a guidance and control system that can map the estimated lander state to a commanded thrust for each lander engine. In this paper, we present a novel integrated guidance and control algorithm designed by applying the principles of reinforcement learning theory. The latter is used to learn a policy mapping the lander's estimated state directly to a commanded thrust for each engine, with the policy resulting in accurate and fuel-efficient trajectories. Specifically, we use proximal policy optimization, a policy gradient method, to learn the policy. Another contribution of this paper is the use of different discount rates for terminal and shaping rewards, which significantly enhances optimization performance.

[*]Research Engineer, Department of Systems and Industrial Engineering, Department of Aerospace and Mechanical Engineering

[†]Charles Stark Draper Assistant Professor, Department of Aeronautics and Astronautics, Email: linaresr@mit.edu.

[‡]Professor, Department of Systems and Industrial Engineering, Department of Aerospace and Mechanical Engineering, E-mail: robertof@email.arizona.edu

We present simulation results demonstrating the guidance and control system's performance in a 6-DOF simulation environment and demonstrate robustness to noise and system parameter uncertainty.

I.   Introduction

Future unconstrained, science-driven, robotic and human missions to large and small planetary bodies will require a high degree of landing accuracy. Indeed, the next generation of planetary landers will need more advanced guidance and control capabilities to satisfy the increasingly stringent accuracy requirements. The latter is driven by the desire to explore regions on planets (e.g. Mars) and satellites (e.g. Moon) that have the potential to yield the highest scientific return. In the case of Mars, the most demanding mission segment is the probably the powered descent phase, where the goal is to achieve a soft pinpoint landing, which we will define as the norm of the position error less than 5 m and the magnitude of the landing velocity below 2 m/s, with the velocity vector directed primarily downward, and negligible deviation from an upright attitude and zero rotational velocity. During a typical Entry, Descent and Landing (EDL) as implemented in past robotic missions to Mars [1, 2], the lander's sensors (radar altimeters) are effectively blind until the heat shield is jettisoned, at which point the lander's guidance, navigation, and control system must quickly estimate the lander's state from a prior distribution extending several km downrange and crossrange, and then use the sequence of state estimates to achieve a soft landing at the target position, typically within one minute of the heat shield being jettisoned.

There are two necessary capabilities to enable a lander to achieve a pinpoint landing. First, the navigation system must be capable of accurately estimating the lander's state during the powered descent phase. Second, the lander's guidance and control system must be capable of using these state estimates to achieve a pinpoint landing. In previous work [3], we made progress towards the development of a guidance, navigation, and control system enabling pinpoint landing for the Mars powered descent phase by developing a navigation system using a Rao-Blackwellized particle filter, a digital terrain map, and the radar altimeters that are currently used on Mars landers, and demonstrated pinpoint landing accuracy in 3-DOF Monte Carlo simulations by coupling the navigation system with the energy-optimal closed-loop guidance algorithm as developed by Battin [4, page 558] and D'Souza [5] (here referred to as DR/DV algorithm). However, as the simulations were run in a 3-DOF environment, this work neglected one of the more challenging portions of the powered descent problem, mapping the guidance system's commanded acceleration to actuator commands. In this work we focus on developing an integrated guidance and control system suitable for the powered descent phase.

For the case of the robotic exploration of Mars, landing accuracy is important for several reasons. First, given accurate high-resolution maps of the Mars surface, the pinpoint-landing problem subsumes the hazard avoidance problem, as a hazard-free landing site can be targeted. Second, delivering a rover closer to a location of scientific interest reduces the risk of the rover malfunctioning before it reaches the desired site. In fact, some sites might be inaccessible to a rover unless the lander can deliver the rover with pinpoint accuracy. Moreover, reducing the distance the rover is required to travel can relax the design requirements for the rover, with a potential reduction in rover mass. Clearly, landing accuracy on the order of several meters is desirable as it would both reduce mission risk and extend the scope of feasible missions.

Considering guidance and control systems that have been deployed for an actual Mars landing, the current benchmark for landing accuracy is the Mars Science Laboratory (MSL), which used a velocity trigger for chute deployment and had an expected 3-sigma landing error ellipse of 19 km by 7 km [6]. This landing ellipse is significantly reduced as compared to previous missions (e.g. [1]), largely due to active bank angle control during the hypersonic re-entry phase based off of the lander's position as estimated by the lander's inertial measurement unit. Using a range trigger for chute deployment has the potential to improve parachute deployment accuracy, which could potentially give landing accuracy on the order of 4-6 km [7]. While this is a significant improvement, it is still a long way from pinpoint accuracy. The MSL system used a trajectory commander to generate a multi-phase trajectory at the start of the powered descent phase. This trajectory is designed so that it can be tracked by the trajectory follower system, which employs six feedback control loops, taking position and attitude estimates from the navigation filter. Note that the trajectory was not optimized for minimal fuel usage [8]. One reason for the poor landing accuracy is that the lander is only capable of estimating altitude and velocity during the powered descent phase using the lander's radar altimeters. Moreover, even if the lander could accurately estimate its position and velocity in the target-centered reference frame in real time, the guidance law is only capable of relatively small divert maneuvers [7].

More recently, a method of loss-less convexification of non-convex control problems has been applied to the planetary landing problem [9,10]. This work proposes a method to convexify the constraints of minimum thrust and thrust vector direction in such a way that the optimal solution of the convexified problem is optimal for the original problem. Convexification allows the use of interior point solution methods that guarantee a solution in bounded polynomial time, making it suitable for real-time powered descent. A cone constraint is imposed to enforce glide slope constraints and avoid terrain features. The trajectory is optimized to minimize fuel usage, and can target a desired landing site from any state in the deployment ellipse, thereby allowing pinpoint landing if coupled with a

suitable navigation system. Importantly, the optimization problem is defined with three degrees of freedom, with the lander modeled as a point mass. Consequently, the trajectory would need to be mapped to lander engine thrust commands using a separately optimized controller. Although this is a significant improvement over prior methods, this approach may not result in the most fuel-efficient trajectories, as the optimization does not take into account atmospheric drag and the aerodynamics of the lander. Moreover, the approach may not be suitable for missions where a cone constraint on the lander's position is incompatible with mission goals. One example of such a mission is landing at the bottom of a deep crater, where a cone cannot be defined that encloses the lander's position at the start of the powered descent phase, has an apex a the desired landing site, and avoids the crater walls. Another limitation is that future missions might require constraints where loss-less convexification is not possible.

In Ref. 11, the loss-less convexification approach was extended to generate 6-DOF trajectories through successive convexifications. Starting with an initial guess for the trajectory, the dynamics are linearized around the trajectory using trust-regions to ensure that the problem remains bounded and feasible at each successive convexification. The successive convexifications continue until convergence, which is defined using tolerance thresholds. To make the problem tractable, the authors assume a constant inertia tensor, although in general, the inertia tensor will change as fuel is consumed. Importantly, this approach will require either a separate controller to track the open loop 6-DOF trajectory or the trajectory will have to be recomputed at each time step as in model predictive control. The paper presents simulation results using a lander with a single gimbaled thruster in an environment which may not be representative of a specific mission.

In Ref. 12 the author uses an optimal control formulation to solve a 3-DOF Mars powered descent problem. The method, named Universal Powered Guidance (UPG), finds a closed form solution for the thrust amplitude and number of switches for the thrust profile and uses a numerically robust zero-finding algorithm to determine the initial costates as well as the optimal burn times. When compared to the Guidance for Fuel-Optimal Large Divert (G-FOLD, [9, 13, 14]) algorithm which is based on the above-mentioned convexification methodology,the UPG code is generally simpler and requires no special customization of the optimization algorithm. Additionally, it is capable of accommodating a variety of problem formulations which may not be possible via a second-order coning approach. The approach has the additional advantage of computing an optimal state at which to start the powered descent phase such that divert requirements are minimized. A disadvantage of this approach is that constraints on glideslope and thrust vector direction are difficult to enforce. Recently, UPG has been recently employed to evaluate the interplay occurring between entry trajectory and fuel usage during a human mission to Mars [15].

To summarize, current practice and the majority of proposed powered decent phase guidance and control systems use two separate and independently optimized systems for guidance and control. The navigation system estimates the lander's state from sensor measurements and passes this estimate to a guidance system. This guidance system generates a trajectory that in general specifies the lander's target state as a function of time. The trajectory is then passed to a control system that tracks the trajectory by determining which thrusters to fire and at what thrust level. There are variations on this approach that, to date, are primarily used for missile guidance. For example, instead of generating a trajectory in real time based off of the lander's state at the beginning of the powered descent phase, a global guidance law such as DR/DV [5] can be developed that maps the lander's state in the target-centered reference frame to a commanded acceleration in the inertial frame. This commanded acceleration is then passed to an actuator control system (sometimes referred to as the autopilot), which must map the commanded acceleration in the inertial frame to a command to fire one or more of the lander's thrusters. Regardless of whether the approach uses trajectory generation coupled with a tracking controller or a guidance law coupled with an autopilot, the basic structure of three separately optimized systems (guidance, navigation, and control) is currently the state-of-the-art.

There are multiple issues that can arise with the separate optimization of the guidance and control systems. First, traditional autopilot design is typically simplified by having three separate controllers, one for roll stabilization, and the other two adjusting the actuators to modify yaw and pitch to attempt to match the commanded acceleration [16]. When an autopilot is split into multiple controllers, it is possible that the combined commanded thrust results in a thrust vector with components that partially cancel, thus reducing fuel efficiency. Autopilot controllers that decouple different types of motion can also limit the types of maneuvers a vehicle can make, which in turn slows the flight control response time. Regardless of the autopilot architecture, a common problem is actuator saturation, where the guidance system's optimal commanded acceleration cannot be implemented given actuator limitations. Although a well-designed control system can compensate for actuator saturation, performance is typically degraded. A complication of separate guidance and control system design is that the reference trajectory must ensure that the sequence of acceleration commands in $\mathbb{R}^3$ is implementable given the thruster configuration of the lander. As an extreme example, if the commanded thrust were to shift by a large angle in a short amount of time, the lander may not be able to change its attitude fast enough to track the commanded thrust. Separate optimization of guidance and control systems can also complicate the design process, as performance issues may only be revealed when the two systems are combined in simulation, requiring a redesign of one or more of the separate systems.

Integrated Guidance and Control (IGC) is an alternative to separate guidance and control systems that have the potential to improve system performance [17]. With IGC, the guidance and control systems are treated as a single system for design purposes. Recent work on IGC has focused primarily on the problem of missile homing phase guidance for both endoatmospheric and exoatmospheric intercepts [18,19,17,20]. Recently, an IGC system for Mars landing using a disturbance observer and multiple sliding surface techniques has been proposed [21], although only a single trajectory was simulated to verify the approach. Therefore it is not clear if the chosen parameters in the gain matrix would generalize to different initial conditions in the deployment ellipse. Moreover, the dynamics used to derive the guidance law do not take into account the lander's changing mass during the powered descent phase. Finally, the system is not truly integrated as the authors assume the ability to directly set the torque and the ability to generate translational thrust in arbitrary directions in $\mathbb{R}^3$ (not possible for a realistic thruster configuration). This essentially assumes decoupling of translational and rotational motion; consequently, the system does not map observations directly to commanded thrust for the lander's individual engines, and a separate controller would be required to do so. Wibben and Furfaro use a multiple sliding surface approach to develop an IGC system for lunar landing [22]. However, the system was not truly integrated, as the 3-DOF thrust vector is first computed and mapped to the body frame, where the algorithm then works out the required attitude and the required torque to target that attitude. Consequently, the system requires a separate controller to map commanded thrust in $\mathbb{R}^3$ and commanded torque to individual engine thrust commands. Moreover, the method was only demonstrated for a single trajectory (there were no Monte Carlo simulations), and the gain matrix had to be tuned to ensure the lander touched down in an upright position. Consequently, it is not clear if that gain matrix would work over a full deployment ellipse.

In this work, we develop an integrated guidance and control system that learns a global policy over the region of state space defined by the deployment ellipse and potential landing sites. This global policy maps the navigation system's estimate of the lander's state directly to commands specifying thrust levels for each engine. The global policy is learned using reinforcement learning (RL). Learning involves simulated interaction between an agent instantiating the policy and the environment over many episodes with randomly generated initial conditions that cover the deployment ellipse. Constraints such as minimum and maximum thrust, glide slope, maximum attitude and rotational velocity, and terrain feature avoidance (such as targeting the bottom of a deep crater) can be included in the reward function, and will be accounted for when the policy is optimized. Note that there is no guarantee on the optimality of trajectories induced by the policy, although in practice it is possible to get close to optimal performance by tuning the reward function.

RL algorithms can be broken down into two major classes, i.e., value function methods

and policy gradient methods. Value function-based algorithms learn a mapping between a state - action tuple and the sum of the future discounted rewards received when starting in that specific state and taking that specific action. Actions with the highest value are then greedily chosen, with limited exploration achieved by choosing a random action with some small probability. Since all possible actions must be considered, these methods only work with discrete action spaces. First proposed by Watkins [23], recent work in value function-based algorithms include deep Q networks [24]. The latter use experience replay to remove temporal correlation between samples and target networks which results in a reduced instability of combining bootstrapping with non-linear function approximators. Deep Q networks have proven effective at control tasks requiring the mapping of pixel level observations directly to control actions, as demonstrated in [24], although they are less effective at control problems where the dimensionality of the dynamics is larger, such as problems in robotic control, and are not suitable for problems requiring continuous action spaces. Deterministic policy gradient algorithms, first proposed by Silver et al., [25], and including the DDPG algorithm [26], can be employed as an alternative to methods based on Deep Q learning. These algorithms also learn a mapping between state-action tuples and values. However, rather than taking actions that globally maximize the value function, the algorithm maintains a separately parameterized policy mapping states to actions, and uses the chain rule to follow the gradient of the value function with respect to the action, given that the action is a function of the state as parameterized by the policy,

Generally, policy gradient algorithms learn a direct stochastic mapping between an observation and action, with the policy network's parameters updated such that the probability of actions leading to higher future rewards is increased. First proposed by R. Williams [27], these algorithms suffer from high variance, which can be substantially reduced by using a state value function baseline. With a baseline, actions that result in future discounted rewards higher than the estimated value of being in that state, as given by the state value function, have their probability increased. In this approach, a state value function mapping a state to the expected sum of future discounted rewards is learned concurrently with a stochastic policy mapping states to actions. For continuous action spaces, it is common to parameterize the policy as a Gaussian distribution with a diagonal covariance matrix, where the policy outputs the mean and variance of the actions conditioned on the state. Recently, Shulman et. al. [28] have proven that policy gradient methods with stochastic policies can have monotonic improvement guarantees, provided that the policy changes during optimization as measured by the KL divergence are constrained to be within certain bounds. The policy optimization problem is posed as a constrained optimization problem that ensures the KL divergence between policy updates remains within specified bounds. The authors used this result to develop the trust region policy optimization algorithm (TRPO), which has

proven effective at solving high-dimensional control problems. Later, Schulman [29] proposed the proximal policy optimization (PPO) method that uses a heuristic to keep KL divergence between policy updates at a level that in practice ensures monotonic improvement during optimization. In practice, PPO works slightly better than TRPO, and has the advantage of not requiring second-order optimization methods. Policy gradient methods are much less sample efficient than methods based on Deep Q learning as they operate on policy. The latter means that policy gradient methods cannot take advantage of sample reuse through techniques such as experience replay [24]. However, policy gradient methods are more stable, and work with minimal hyperparameter tuning, and tend to perform better in systems with high dimensional dynamics.

A potential drawback of the RL approach is the lack of local or global stability guarantees. In our opinion, this is not a major problem, as competing approaches such as optimal control, traditional linear control, and Lyapunov control only have stability guarantees if the system model is an accurate representation of the actual dynamic system. In some cases, this assumption may hold, but in others, such as hypersonic re-entry, it may not be possible to apply these methods using a high fidelity system model. Nevertheless, the optimized guidance and control system could be tested using a high fidelity model. Even for the relatively simple dynamics of the Mars powered descent phase, recent works assume either a static inertia tensor [11], or a static mass [21]. In contrast, RL allows both optimization and verification using a high fidelity dynamics model. Since RL can be model free (as in this work), the complexities of the dynamics do not complicate the problem formulation within the RL framework. Another issue with RL is implementing hard state constraints. It is possible to approximate hard state constraints by terminating the episode and giving the agent a large negative reward upon a constraint violation, but the latter typically requires manual tuning of the reward function, as we discuss in the sequel (see section II.D). Such manual tuning typically requires a certain amount of trial and error and results in the constraints being satisfied. However, there has been recent work on implementing hard constraints in RL using constrained policy optimization [30]. Note that hard constraints on actions are simple; we clip any action returned by the policy such that it satisfies the constraints, and the agent learns to deal with this clipping.

A comparison of RL and optimal control approaches to guidance and control are given below in Table 1. The point of the comparison is not to make the case that one approach should be preferred over the other, but rather to suggest the scenarios where it might make sense to use the RL framework to solve guidance and control problems.

Previous work using RL to solve control problems has focused on applications in robotics, with very few published works addressing problems in aerospace guidance and control. The first application of RL to a problem in the aerospace domain was applied to autonomous

Table 1. A Comparison of Optimal Control and RL.

Optimal Control	Reinforcement Learning
Single trajectory (except for trivial cases where HJB equations can be solved)	Global over theatre of operations
Unbounded run time except for special cases such as convex constraints	Extremely fast run time for trained policy ($<$ 1ms in this work)
Hybrid dynamics requires special treatment	Hybrid Dynamics handled seamlessly
Dynamics need to be represented as ODE, possibly constraining fidelity of model used in optimization	No constraints on dynamics representation. Agent can learn in a high fidelity simulator (i.e., Navier-Stokes modeling of aerodynamics)
Open Loop (requires a controller to track the optimal trajectory)	Closed Loop (Integrated guidance and control)
Output feedback (co-optimization of state estimation and guidance law) an open problem for non-linear systems	Can learn from raw sensor outputs allowing fully integrated GNC (pixels to actuator commands). Can learn to compensate for sensor distortion.
Requires full state feedback	Does not require full state feedback
Elegantly handles state constraints	State constraints handled either via large negative rewards and episode termination or more recently, modification of policy gradient algorithm. Control constraints straightforward to implement
Deterministic	Stochastic, learning does not converge every time, may need to run multiple policy optimizations

helicopter flight [31]. More recently, in Ref. 32 the authors compared reinforcement learning to sliding mode control and linear control in a quadrotor control application. The authors used policy iteration [33] in a model-based formulation, where a model was learning from flight data using weighted least squares. Both the sliding mode and RL methods resulted in stable controllers, whereas the linear control methods failed.

To our knowledge, this is the first published work demonstrating an RL-derived integrated guidance and control system applied to planetary landing in a 6-DOF environment. In fact, we were unable to find any work using RL to learn even a 3-DOF guidance law. In previous work, we achieved good results in a 2-DOF environment, but over a limited range of initial conditions [34]. Open-AI gym has one environment for a 3-DOF planar lunar landing[1], but the problem is simplified by using a range of initial conditions that are unrealistically reduced from that required for an actual lunar landing (the crossrange dispersion is only three times the width of the landing site). In our work, we have found that solutions that work over a limited range of initial conditions often fail when the initial condition range is extended. Our ultimate goal is to use RL to develop a fully integrated guidance,

---

[1]https://gym.openai.com/envs/LunarLander-v2/

navigation, and control system, where the policy maps raw sensor observations directly to actuator commands, i.e., radar altimeter readings to commanded thrust for the lander's individual engines. The RL framework can be applied to solve many different types of guidance and control problems, including missile homing phase guidance, exoatmospheric intercepts, hypersonic reentry maneuvering for planetary landings, and booster recovery via powered landing.

This paper is organized as follows. In section II, the problem formulation together with the RL methodology is applied to the proposed work. In section III, we show the results from policy optimization and testing for 6-DOF planetary landing scenarios. In section IV, a comparison between RL-based 3-DOF/6-DOF closed-loop policy and open-loop optimal trajectory derived using direct methods is reported. In section V, conclusions and future work are reported.

II.   Problem Formulation

A.   Problem Statement

Consider the powered descent problem on a large planetary body such as Mars. In this section we describe the RL set-up that is employed to find a 6-DOF closed-loop policy that generates quasi-fuel optimal trajectories capable of a) driving the lander to the desired location to the planetary surface with pin-point accuracy and b) satisfying flight and systems constraints (e.g. glide slope, thrust and attitude constraints). To set the stage for the IGC RL , consider the standard formulation for a typical trajectory optimization problem (see for example [10]):

Minimum-Fuel Problem: Find the thrust program and flight time that minimize the following cost functions:

$$\min_{t_f, \mathbf{T}} \quad m_L(t_f) = \max_{t_f, \mathbf{T}} \int_0^{t_f} \|\mathbf{T}\| \mathbf{dt} \tag{1}$$

Subject to the following constraints (equations of motion):

$$\ddot{\mathbf{r}} = \mathbf{g} + \frac{\mathbf{T}}{m} \tag{2a}$$

$$\dot{m} = \frac{\|\mathbf{T}\|}{I_{\text{sp}} g_{\text{ref}}} \tag{2b}$$

and the following boundary conditions:

$$\mathbf{r}(0) = \mathbf{r}_0 \tag{3a}$$

$$\mathbf{v}(0) = \dot{\mathbf{r}}(0) = \mathbf{v}_0 \tag{3b}$$

$$\mathbf{r}(t_f) = \mathbf{r}_f \tag{3c}$$

$$\mathbf{v}(t_f) = \dot{\mathbf{r}}(t_f) = \mathbf{v}_f \tag{3d}$$

And additional flight (glide slope) and thrust constraints:

$$0 < T_{min} < \|\mathbf{T}\| < T_{max} \tag{4a}$$

$$\theta_{\text{alt}} = \arctan\left(\frac{\sqrt{r_y(t)^2 + r_z(t)^2}}{r_x(t)}\right) < \tilde{\theta}_{alt} \tag{4b}$$

Generally, one uses an inertial target-centered reference frame that neglects the rotation of Mars. Since the powered descent phase for a typical Mars mission is initiated at an altitude that is low w.r.t. the planet's radius, the gravity $\mathbf{g}$ is assumed to be constant (i.e., motion in a uniform gravitational field). The lander's position is given by the vector $\mathbf{r} \in \mathbb{R}^3$ and $\mathbf{r}^T = [r_x \quad r_y \quad r_z]^T$, where r_x, r_y, and r_z are the lander's downrange, crossrange, and altitude in the target centered reference frame. Likewise,the velocity is described by a vector $\mathbf{v} \in \mathbb{R}^3$ and $\mathbf{v}^T = [v_x \quad v_y \quad v_z]^T$, where v_x, v_y, and v_z are the lander's downrange, crossrange, and descent components, respectively. The vectors $\mathbf{r_0}, \mathbf{v_0}$, and $\mathbf{r_f}, \mathbf{v_f}$ are the initial and final position/velocity vetors, respectively. Importantly, the thrust magnitude $\|\mathbf{T}\|$ is limited between a maximum and minimum. The angle $\tilde{\theta}_{\text{alt}} < \frac{\pi}{2}$ represents that glide slope constraint. This is cast as a typical optimal control problem and generally solved using numerical methods (e.g. [10,12]). For a 6-DOF problem, one needs to provide additional information about the attitude of the lander. The body frame is defined with the z-axis passing vertically through the lander and orthogonal x and y axes completing the body reference frame. Although a more formal 6-DOF trajectory and attitude optimal control problem can be done (e.g. [11]), the RL framework is defined such that the agent (i.e., the lander) responds to a single cost (reward) signal which needs to be generally minimized (maximized) during the search process. Indeed, with the goal of the lander arriving at the origin of the target-centered reference frame at some specified velocity vector $\mathbf{v} = \dot{\mathbf{r}}$ and at a specified attitude, we can define the RL-based landing problem as follows:

RL 6-DOF Closed-Loop Landing Problem:

Minimize:

1. Terminal Position Error: $\|\mathbf{r}\|$ at $t = t_f$

2. Terminal Velocity Error: $\|\mathbf{v}\|$ at $t = t_f$

3. Terminal Attitude Error: Magnitude of Pitch ϕ and Roll θ at $t = t_f$

4. Terminal Rotational Velocity Error: $\|\boldsymbol{\omega}\|$ at $t = t_f$

5. Control Effort: $\sum \|\mathbf{T}\|$ where the sum is over a trajectory and T is the total thrust

Subject to:

1. Terminal Glideslope: $\|v_z\|/\|v_{x,y}\| > 5$ over final 2 m of descent (soft constraint)

2. Attitude Constraints: Magnitude of Pitch ϕ and Roll θ less than 80 degrees (hard constraint)

3. Equations of Motion (set by the environment)

Here, hard constraints are imposed by terminating the episode with a negative reward, whereas soft constraints are encouraged through rewards but without premature termination of the episode. Next, the lander's equations of motion, the proposed policy gradient approach and a more detailed RL-based formulation for the cost function are presented.

B. Equations of Motion

The force $\mathbf{F}^B$ and torque $\mathbf{L}^B$ in the lander's body frame for a given commanded thrust depends on the placement of the thrusters in the lander structure. We can describe the placement of each thruster through a body-frame direction vector $\mathbf{d}$ and a position vector $\mathbf{r}$, both in $\mathbb{R}^3$. The direction vector is a unit vector giving the direction of the body frame force that results when the thruster is fired. The position vector gives the body frame location where the force resulting from the thruster firing is applied for purposes of computing torque. For a lander with k thrusters, the body frame force and torque associated with one or more thrusters firing is given by,

$$\mathbf{F}^B = \sum_{i=1}^{k} \mathbf{d}^{(i)} T_{cmd}^{(i)} \tag{5a}$$

$$\mathbf{L}^B = \sum_{i=1}^{k} \mathbf{r}^{(i)} \times \mathbf{F}_B^{(i)} \tag{5b}$$

where $T_{cmd_i} \in [T_{min}, T_{max}]$ is the commanded thrust for thruster i, T_{min} and T_{max} are a thruster's minimum and maximum thrust, $\mathbf{d}^{(i)}$ the direction vector for thruster i, and $\mathbf{r}^{(i)}$ the position of thruster i. The total body frame force and torque are calculated by summing

the individual forces and torques. The dynamics model uses the lander's current attitude $\mathbf{q}$ to convert the body frame thrust vector to the inertial frame is given by

$$\mathbf{F}^N = A_B^N(\mathbf{q})^T \mathbf{F}^B \tag{6}$$

where $A_B^N(\mathbf{q})$ is the direction cosine matrix mapping the inertial frame to body frame obtained from the current attitude parameter $\mathbf{q}$. The attitude matrix is related to the quaternion by the following expression [35]:

$$A_B^N(\mathbf{q}) = \Xi^T(\mathbf{q})\Psi(\mathbf{q}) \tag{7}$$

where

$$\Xi(\mathbf{q}) \equiv \begin{bmatrix} q_4 I_{3\times3} + [\boldsymbol{\varrho}\times] \\ -\boldsymbol{\varrho}^T \end{bmatrix} \tag{8a}$$

$$\Psi(\mathbf{q}) \equiv \begin{bmatrix} q_4 I_{3\times3} - [\boldsymbol{\varrho}\times] \\ -\boldsymbol{\varrho}^T \end{bmatrix} \tag{8b}$$

where the quaternion is divided into the $\boldsymbol{\varrho}$ and q_4, the vector and scalar components, respectively, and $\mathbf{q}^T = [\boldsymbol{\varrho} \quad q_4]^T$. The body force component can be rotated into the inertial frame using the following expression

$$\mathbf{F}^N = A_B^N(\mathbf{q})^T \mathbf{F}^B \tag{9}$$

The rotational velocities $\boldsymbol{\omega}_{B/N}$ are then obtained by integrating the Euler rotational equations of motion, is as follows [36]:

$$\mathbf{J}\dot{\boldsymbol{\omega}}_{B/N} = -\left[\boldsymbol{\omega}_{B/N}\times\right]\mathbf{J}\boldsymbol{\omega}_{B/N} + \mathbf{L}^B + \mathbf{L}_{\text{env}}^{\ B} \tag{10}$$

with

$$[\mathbf{a}\times] \equiv \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix} \tag{11}$$

for any general 3×1 vector $\mathbf{a}$ defined such that $[\mathbf{a}\times]\mathbf{b} = \mathbf{a} \times \mathbf{b}$. The vector $\mathbf{L}^B$ is the body frame torque as given in Equation (5b), $\mathbf{L}_{\text{env}}$ is the body frame torque from external disturbances, and $\mathbf{J}$ is the lander's inertia tensor. The lander's attitude is then updated by integrating the differential kinematic given by

$$\dot{\mathbf{q}} = \frac{1}{2}\Xi(\mathbf{q})\boldsymbol{\omega}_{B/N} \tag{12}$$

where the lander's attitude is parameterized using the quaternion representation and $\boldsymbol{\omega}_i$ denotes the i^{th} component of the rotational velocity vector $\boldsymbol{\omega}_{B/N}$. The translational motion is modeled as follows

$$\dot{\mathbf{r}} = \mathbf{v} \tag{13a}$$

$$\dot{\mathbf{v}} = \frac{\mathbf{F}^N + {\mathbf{F}_{\text{env}}}^N}{m} + \mathbf{g} \tag{13b}$$

$$\dot{m} = -\frac{\sum_i^k \|\mathbf{F}^{B(i)}\|}{I_{\text{sp}} g_{\text{ref}}} \tag{13c}$$

where $\mathbf{F}^{N(i)}$ is the inertial frame force as given in Eq. (9), k is the number of thrusters, $g_{\text{ref}} = 9.8 \text{ m/s}^2$, $\mathbf{g} = \begin{bmatrix} 0 & 0 & -3.7114 \end{bmatrix} \text{ m/s}^2$ is used for Mars, $I_{\text{sp}} = 225$ s, and the spacecraft's mass is m. $\mathbf{F}_{\text{env}}$ is a vector of normally distributed random variables representing environmental disturbances such as wind and variations in atmospheric density. As a rough estimate of the force caused by wind gusts, a cross-sectional lander area of 100 square feet was assumed, and we used the Cornell University wind pressure calculator to calculate the resulting wind pressure in pounds per square foot, divided this by 168 to account for the difference in average surface air pressure between the two planets, multiplied by the assumed cross-section of the lander, and then converted the force to Newtons. Using this rough approximation, a wind of 100 m/s would result in a force of 162 N. Note that the Mars GRAM model shows a strong jet near 5-km altitude and 70 degrees N latitude with winds reaching 100 m/s, so it is reasonable to expect a Lander to be able to handle this case.

For purposes of modeling the lander's moments of inertia, we model the lander as a uniform density ellipsoid, with inertia matrix given by

$$\mathbf{J} = \frac{m}{5} \begin{bmatrix} b^2 + c^2 & 0 & 0 \\ 0 & a^2 + c^2 & 0 \\ 0 & 0 & a^2 + b^2 \end{bmatrix} \tag{14}$$

where a, b, and c correspond to the body frame x, y, and z axes. m is the lander's mass, which is updated as shown in Eq. (13c). We assume the lander has a wet mass of 2000 kg and four throttleable thrusters with a minimum and maximum thrust magnitude of 1000 N and 5000 N respectively. The four thrusters are located in the lander body frame as shown below in Table 2, where x, y, and z are the body frame axes. Roll is about the x-axis, yaw is about the z-axis, and pitch is about the y-axis. Note that this thruster configuration does not allow any direct control of the rotational velocity around the z-axis. However, the lander's yaw will change during the trajectory, but due to coupling with pitch via roll rather than due to torque caused by thrust. Direct yaw control could be implemented by

positioning the thrusters at a slight angle to the body-frame z-axis, which might result in faster convergence of policy optimization. The lander's translational motion is described using a target-centered inertial reference frame. The navigation system provides updates to the guidance system every 0.2 s, and we integrate the equations of motion using fourth order Runge-Kutta integration with a time step of 0.05 s.

Table 2. Body Frame Thruster Locations.

Thruster	x (m)	y (m)	z (m)
1	0	-2	-1
2	0	2	-1
3	-2	0	-1
4	2	0	-1

C.   Policy Gradient Method

In the RL framework, an agent learns through episodic interaction with an environment how to successfully complete a task by learning a *policy* that maps observations to actions. The environment initializes an episode by randomly generating an internal state, mapping this internal state to an observation, and passing the observation to the agent. These observations could be a corrupted version of the internal state (to model sensor noise) or could be raw sensor outputs such as Doppler radar altimeter readings, or a multi-channel pixel map from an electro-optical sensor. At each step of the episode, an observation is generated from the internal state and given to the agent. The agent uses this observation to generate an action that is sent to the environment; the environment then uses the action and the current state to generate the next state and a scalar reward signal. The reward and the observation corresponding to the next state are then passed to the agent. The environment can terminate an episode, with the termination signaled to the agent via a done signal. The termination could be due to the agent completing the task or violating a constraint. Initially, the agent's actions are random, which allows the agent to explore the state space and begin learning the value of experiencing a given observation, and which actions are to be preferred as a function of this observation. Here the value of an observation is the expected sum of discounted rewards received after experiencing that observation; this is similar to the cost-to-go in optimal control. As the agent gains experience, the amount of exploration is decreased, allowing the agent to exploit this experience. For most applications (unless a stochastic policy is required), when the policy is deployed in the field, exploration is turned off, as exploration gets quite expensive using an actual lander. The safe method of continuous learning in the field is to have the lander send back telemetry data, which can be used to improve the environment's dynamics model and update the policy via simulated experience.

In the following discussion, the vector $\mathbf{x}_k$ denotes the observation provided by the envi-

ronment to the agent. Note that in general $\mathbf{x}_k$ does not need to satisfy the Markov property. In those cases where it does not, several techniques have proven successful in practice. In one approach, observations spanning multiple time steps are concatenated, allowing the agent access to a short history of observations, which helps the agent infer the motion of objects in consecutive observations. This was the approach used in [24]. In another approach, a recurrent neural network is used for the policy and value function implementations. The recurrent network allows the agent to infer motion from observations, similar to the way a recursive Bayesian filter can infer velocity from a history of position measurements. The use of recurrent network layers has proven effective in supervised learning tasks where a video stream needs to be mapped to a label [37].

Each episode results in a trajectory defined by observation, actions, and rewards; a step in the trajectory at time t_k can be represented as $(\mathbf{x}_k, \mathbf{u}_k, r_k)$, where $\mathbf{x}_k$ is the observation provided by the environment, $\mathbf{u}_k$ the action taken by the agent using the observation, and r_k the reward returned by the environment to the agent. The reward can be a function of both the observation $\mathbf{x}_k$ and the action $\mathbf{u}_k$. The reward is typically discounted to allow for infinite horizons and to facilitate temporal credit assignment. Then the sum of discounted rewards for a trajectory can be defined as

$$r(\boldsymbol{\tau}) = \sum_{i=0}^{T} \gamma^i r_k(\mathbf{x}_k, \mathbf{u}_k) \tag{15}$$

where $\boldsymbol{\tau} = [\mathbf{x}_0, \mathbf{u}_0, ..., \mathbf{x}_T, \mathbf{u}_T]$ denotes the trajectory and $\gamma \in [0, 1)$ denotes the discount factor. The objective function the RL methods seek to optimize is given by

$$J(\boldsymbol{\theta}) = \mathbb{E}_{p(\boldsymbol{\tau})}\left[r(\boldsymbol{\tau})\right] = \int_{\mathbb{T}} r(\boldsymbol{\tau}) p_{\boldsymbol{\theta}}(\tau) d\boldsymbol{\tau} \tag{16}$$

where

$$p_{\boldsymbol{\theta}}(\boldsymbol{\tau}) = \left[\prod_{k=0}^{T} p(\mathbf{x}_{k+1}|\mathbf{x}_k, \mathbf{u}_{\boldsymbol{\theta}})\right] p(\mathbf{x}_0) \tag{17}$$

where $\mathbb{E}_{p(\boldsymbol{\tau})}[\cdot]$ denotes the expectation over trajectories and in general $\mathbf{u}_{\boldsymbol{\theta}}$ may be deterministic or stochastic function of the policy parameters, $\boldsymbol{\theta}$. However, it was noticed by Ref. 27 that if the policy is chosen to be stochastic, where $\mathbf{u}_k \sim \pi_{\boldsymbol{\theta}}(\mathbf{u}_k|\mathbf{x}_k)$ is a pdf for $\mathbf{u}_k$ conditioned on $\mathbf{x}_k$, then a simple policy gradient expression can be found.

$$\begin{aligned} \nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) =& \mathbb{E}_{p(\boldsymbol{\tau})}\left[\sum_{k=0}^{T} r_k(\mathbf{x}_k, \mathbf{u}_k) \nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}}(\mathbf{u}_k|\mathbf{x}_k)\right] \\ \approx& \sum_{i=0}^{M}\sum_{k=0}^{T} r_k(\mathbf{x}_k^i, \mathbf{u}_k^i) \nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}}(\mathbf{u}_k^i|\mathbf{x}_k^i) \end{aligned} \tag{18}$$

where the integral over $\boldsymbol{\tau}$ is approximated with M samples from $\boldsymbol{\tau}^i \sim p_{\boldsymbol{\theta}}(\boldsymbol{\tau})$ which are monte carlo roll-outs of the policy given the environment's transition pdf, $p(\mathbf{x}_{k+1}|\mathbf{x}_k)$. The expression in Eq. (18) is called the policy gradient and the form of this equation is referred to as the REINFORCE method [27]. Since the development of the REINFORCE method additional theoretical work improved on the performance of the REINFORCE method. In particular, it was shown that the reward $r_k(\mathbf{x}_k, \mathbf{u}_k)$ in Eq. (18) can be replaced with state-action value function $Q^\pi(\mathbf{x}_k, \mathbf{u}_k)$, this result is known as the Policy Gradient Theorem. Furthermore, the variance of the policy gradient estimate that is derived from the monte carlo roll-outs, $\boldsymbol{\tau}^i$, is reduced by subtracting a state-dependent basis from $Q^\pi(\mathbf{x}_k, \mathbf{u}_k)$. This basis is commonly chosen to be the state value function $V^\pi(\mathbf{x}_k)$, and we can define $A^\pi(\mathbf{x}_k, \mathbf{u}_k) = Q^\pi(\mathbf{x}_k, \mathbf{u}_k) - V^\pi(\mathbf{x}_k)$. This method is known as the Advantage-Actor-Critic (A2C) Method. The policy gradient for the A2C method is given by

$$\nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) \approx \sum_{i=0}^{M} \sum_{k=0}^{T} A_{\mathbf{w}}^\pi(\mathbf{x}_k^i, \mathbf{u}_k^i) \nabla_{\boldsymbol{\theta}} \log \pi_{\boldsymbol{\theta}}(\mathbf{u}_k^i|\mathbf{x}_k^i) \tag{19}$$

where the $\mathbf{w}$ subscript denotes a function parameterized by $\mathbf{w}$.

D. Proximal Policy Optimization

The Proximal Policy Optimization (PPO) approach [29] is a type of policy gradient which has demonstrated state-of-the-art performance for many RL benchmark problem. The PPO approach is developed using the properties of the Trust Region Policy Optimization (TRPO) Method [28]. The TRPO method formulates the policy optimization problem using a constraint to restrict the size of the gradient step taken during each iteration [38]. The TRPO method policy update is calculated using the following problem statement:

$$\begin{aligned} \underset{\boldsymbol{\theta}}{\text{minimize}} \quad & \mathbb{E}_{p(\boldsymbol{\tau})}\left[\frac{\pi_{\boldsymbol{\theta}}(\mathbf{u}_k|\mathbf{x}_k)}{\pi_{\boldsymbol{\theta}_{\text{old}}}(\mathbf{u}_k|\mathbf{x}_k)} A_{\mathbf{w}}^\pi(\mathbf{x}_k, \mathbf{u}_k)\right] \\ \text{subject to} \quad & \mathbb{E}_{p(\boldsymbol{\tau})}\left[\text{KL}\left(\pi_{\boldsymbol{\theta}}(\mathbf{u}_k|\mathbf{x}_k), \pi_{\boldsymbol{\theta}_{\text{old}}}(\mathbf{u}_k|\mathbf{x}_k)\right)\right] \leqslant \delta \end{aligned} \tag{20}$$

where the function $\text{KL}(\cdot, \cdot)$ is the Kullback-Leibler (KL) divergence [39]. The parameter δ is a tuning parameter but the theory justifying the TRPO methods proves monotonic improvement in the policy performance if the policy change in each iteration is bounded a parameter C. The parameter C is computed using the Kullback-Leibler (KL) divergence [39]. Reference 28 computes a closed-form expression for C but this expression leads to prohibitively small steps, and therefore, Eq. (20) with a fix constraint is used. Additionally, Eq. (20) is approximately solved using the conjugate gradient algorithm, which approximates the constrained optimization problem given by Eq. (20) with a linearized objective function

and a quadratic approximation for the constraint. The PPO method approximates the TRPO optimization process by accounting for the policy adjustment constrain with a clipped objective function. The objective function used with PPO can be expressed in terms of the probability ratio $p_k(\boldsymbol{\theta})$ given by,

$$p_k(\boldsymbol{\theta}) = \frac{\pi_{\boldsymbol{\theta}}(\mathbf{u}_k|\mathbf{x}_k)}{\pi_{\boldsymbol{\theta}_{\mathrm{old}}}(\mathbf{u}_k|\mathbf{x}_k)} \tag{21}$$

where the PPO objective function is then as follows:

$$J(\boldsymbol{\theta}) = \mathbb{E}_{p(\boldsymbol{\tau})}\left[\min\left[p_k(\boldsymbol{\theta}), \mathrm{clip}(p_k(\boldsymbol{\theta}), 1-\epsilon, 1+\epsilon)\right] A^{\pi}_{\mathbf{w}}(\mathbf{x}_k, \mathbf{u}_k)\right] \tag{22}$$

This clipped objective function has been shown to maintain the KL divergence constraints, which aids convergence by insuring that the policy does not change drastically between updates. PPO uses an approximation to the advantage function that is the difference between the empirical return and a state value function baseline is given by the following:

$$A^{\pi}_{\mathbf{w}}(\mathbf{x}_k, \mathbf{u}_k) = \left[\sum_{\ell=k}^{T} \gamma^{\ell-k} r(\mathbf{x}_\ell, \mathbf{u}_\ell)\right] - V^{\pi}_{\mathbf{w}}(\mathbf{x}_k) \tag{23}$$

Here the value function $V^{\pi}_{\mathbf{w}}$ is learned using the cost function given by

$$L(\mathbf{w}) = \sum_{i=1}^{M}\left(V^{\pi}_{\mathbf{w}}(\mathbf{x}_k^i) - \left[\sum_{\ell=k}^{T} \gamma^{\ell-k} r(\mathbf{u}_\ell^i, \mathbf{x}_\ell^i)\right]\right)^2 \tag{24}$$

In practice, policy gradient algorithms update the policy using a batch of trajectories (rollouts) collected by interaction with the environment. Each trajectory is associated with a single episode, with a sample from a trajectory collected at step k consisting of observation $\mathbf{x}_k$, action $\mathbf{u}_k$, and reward $r_k(\mathbf{x}_k, \mathbf{u}_k)$. Finally, gradient accent is performed on $\boldsymbol{\theta}$ and gradient decent on $\mathbf{w}$ and update equations are given by

$$\mathbf{w}^+ = \mathbf{w}^- - \beta_{\mathbf{w}} \nabla_{\mathbf{w}} L(\mathbf{w})|_{\mathbf{w}=\mathbf{w}^-} \tag{25}$$

$$\boldsymbol{\theta}^+ = \boldsymbol{\theta}^- + \beta_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta})|_{\boldsymbol{\theta}=\boldsymbol{\theta}^-} \tag{26}$$

where $\beta_{\mathbf{w}}$ and $\beta_{\boldsymbol{\theta}}$ are the learning rates for the value function, $V^{\pi}_{\mathbf{w}}(\mathbf{x}_k)$, and policy, $\pi_{\boldsymbol{\theta}}(\mathbf{u}_k|\mathbf{x}_k)$, respectively. In our implementation, we adjust the clipping parameter ϵ to target a KL divergence between policy updates of 0.001. The policy and value function is learned concurrently, as the estimated value of a state is policy dependent. We use a Gaussian distribution with mean $\pi_{\boldsymbol{\theta}}(\mathbf{u}_k|\mathbf{x}_k)$ and a diagonal covariance matrix for the action distribution in the policy. Because the log probabilities are calculated using the exploration variance, the degree of

exploration automatically adapts during learning such that the objective function is maximized.

E. Implementation Details

In order to facilitate reproduction of our results, we include in this section several techniques we used in our implementation. We use the ADAM optimizer [40] to adjust the learning rate for both the policy and value function networks. We use an approximation to the KL divergence for a Gaussian that is given by the mean square difference between the pre and post log probabilities of policy actions. This approximation is close to the exact KL divergence, and is used in the open AI baseline PPO implementation PPO2.

$$
\epsilon = \begin{cases} \min(\epsilon_{max}, 1.5\epsilon) & \text{if kl} < \frac{1}{2}\text{kl}_{\text{targ}} \\ \max(\epsilon_{min}, \frac{1}{1.5}\epsilon) & \text{if kl} > 2\text{kl}_{\text{targ}} \end{cases} \tag{27}
$$

$$
\zeta = \begin{cases} 1.5\zeta & \text{if kl} < \frac{1}{2}\text{kl}_{\text{targ}} \text{ and } \epsilon > \frac{1}{2}\epsilon_{\max} \text{ and } \zeta < \zeta_{\max} \\ \frac{1}{1.5}\zeta & \text{if kl} > 2\text{kl}_{\text{targ}} \text{ and } \epsilon < 2\epsilon_{\min} \text{ and } \zeta > \zeta_{\min} \end{cases} \tag{28}
$$

We adjust the clipping parameter ϵ based off of the KL divergence between policy updates as shown in Equation (27), where kl is the KL divergence between policy updates, kl_{targ} is the target KL divergence, and $\epsilon_{\max}$ and $\epsilon_{\min}$ are set to 0.01 and 0.5. A similar implementation was suggested in Reference [29]. We also adjust the ADAM step size by multiplying it by a parameter ζ as shown in (28), where $\zeta_{\min}$ and $\zeta_{\max}$ are set to 0.1 and 10. A learning rate adjustment was used in [29] for the roboschool environments, but the exact implementation was not divulged.

When approximating functions using artificial neural networks, it is important to scale the inputs to avoid saturating the activation functions at each layer. Most implementations scale the network input using statistics calculated over a given rollout. Specifically, each element of an input vector is scaled by first subtracting the mean of the element and then dividing by three times the standard deviation, with the statistics calculated over a batch of rollouts. We use a variation on this approach that adjusts the statistics used for scaling incrementally over the entire optimization, which boosts performance by avoiding discontinuities in the scaling statistics. It is also important to insure that the magnitude of the neural network outputs are are reasonably close to unity, although full normalization is not required. For policy parameter learning, we scale the action such that maximum thrust for an engine corresponds to a value of one. For value function parameter learning, we use a heuristic that multiplies the rewards accumulated over an episode by a factor of $1 - \gamma$.

F. RL Problem Formulation

In order to apply the reinforcement learning framework developed in Section D to a particular problem, we need to define an environment and reward function and specify the policy and value function network architectures. To test the policy, the trained policy is substituted for the agent. The action taken by the agent based on the observation provided by the dynamics model is interpreted as a thrust command by the thruster model, which passes a body frame force and torque to the dynamics model, which computes the next state. An episode terminates when the lander's altitude falls below zero or the attitude constraint is violated. The initial condition generator generates random initial conditions for the lander with values uniformly distributed between the minimum and the maximum values given in Table 3. We used a reduced set of initial conditions in order to reduce simulation time to allow faster iteration over different approaches to solving the Mars landing problem using RL.

Table 3. Lander Initial Conditions for Optimization.

	Velocity		Position	
	min (m/s)	max (m/s)	min (m)	max (m)
Downrange	-70	-10	0	2000
Crossrange	-30	30	-1000	1000
Elevation	-90	-70	2300	2400
	min (rad/s)	max (rad/s)	min (rad)	max (rad)
Yaw	0.00	0.00	$-\pi/8$	$\pi/8$
pitch	-0.01	0.01	$\pi/4 - \pi/8$	$\pi/4 + \pi/16$
roll	-0.01	0.01	$-\pi/8$	$\pi/8$

To speed up the design process, we developed a 3-DOF RL environment where we could quickly assess the impact of different reward functions and hyperparameter settings. A policy optimized in the 3-DOF RL environment converges around 70 times faster than in the 6-DOF environment. The reward function for the 3-DOF environment is identical to that of the 6-DOF environment. We show optimization learning curves for the 3-DOF environment to allow comparison.

The policy and value functions are implemented using four-layer neural networks with tanh activations on each hidden layer. The network architectures are as shown in Table 4, where n_{hi} is the number of units in layer i, obs_dim is the observation dimension, and act_dim is the action dimension.

The most challenging part of solving the Mars landing problem using RL was the development of a reward function that works well in a sparse reward setting. If we only reward the agent for making a soft pinpoint landing at the correct attitude and with close to zero rotational velocity, the agent would never see the reward within a realistic number

Table 4. Policy and Value Function network architecture.

Layer	Policy Network # units	Policy Network activation	Value Network # units	Value Network activation
hidden 1	$10*\text{obs\_dim}$	tanh	$10*\text{obs\_dim}$	tanh
hidden 2	$\sqrt{n_{\text{h1}} * n_{\text{h3}}}$	tanh	$\sqrt{n_{\text{h1}} * n_{\text{h3}}}$	tanh
hidden 3	$10*\text{act\_dim}$	tanh	5	tanh
output	act_dim	linear	1	linear

of episodes, as the probability of achieving such a landing using random actions in a 6-DOF environment with realistic initial conditions is exceedingly low. The sparse reward problem is typically addressed using inverse reinforcement learning [41], where a per time step reward function is learned from expert demonstrations. With a reward given at each step of agent-environment interaction, the rewards are no longer sparse. In theory, demonstrations using optimal control could provide trajectories for inverse RL, but this would be very computationally expensive for 6-DOF trajectories, and many optimal control packages have trouble with complex non-convex and/or non-differentiable constraints.

Instead, we chose a different approach, where we engineer a reward function that, at each time step, provides hints to the agent (referred to as shaping rewards) that drive it towards a soft pinpoint landing. The recommended approach for such a reward shaping function is to make the reward a difference of potentials, in which case theoretical results have shown that the additional reward does not change the optimal policy [31]. We experimented with several potential functions with no success. Instead, we drew inspiration from biological systems that use the gaze heuristic. The gaze heuristic is used by animals such as hawks and cheetahs to intercept prey (and baseball players to catch fly balls) and works by keeping the line of sight angle constant during the intercept. The gaze heuristic is also the basis of the well known PN guidance law used for homing phase missile guidance.

In our case, the landing site is not maneuvering, and we have the additional constraint that we want the terminal velocity to be small. Therefore we use a heuristic where the agent attempts to keep its velocity vector aligned with the line of sight vector. Since the target is not moving in the target-centered reference frame, the target's future position is its current position, and the optimal action is to head directly towards the target. Such a rule results in a pinpoint, but not necessarily soft, landing. To achieve the soft landing, the agent estimates time-to-go as the ratio of the range and the magnitude of the lander's velocity and reduces the targeted velocity as time-to-go decreases. It is also important that the lander's terminal velocity be directed predominantly downward, the lander's terminal attitude is upright, and there are negligible terminal rotational velocity components. To achieve these requirements, we use the piecewise reward shaping function given below in Eqs. (29a), (29b), (29c), (29d), and (29e), where τ_1 and τ_2 are hyperparameters and v_o is set to the magnitude of the lander's

velocity at the start of the powered descent phase. We see that the shaping rewards take the form of a velocity field that maps the lander's position to a target velocity. In words, we target a location 15 m above the desired landing site and target a z-component of landing velocity equal to -2 m/s. Below 15 m, the downrange and crossrange velocity components of the target velocity field are set to zero, which encourages a vertical descent. Targeting a vertical descent has the beneficial side effect of encouraging the agent to keep the attitude level with no rotational velocity. This results in a good landing attitude with small rotational velocity components and a velocity directed primarily downward.

$$\mathbf{v}_{targ} = -v_o \left(\frac{\hat{\mathbf{r}}}{\|\hat{\mathbf{r}}\|} \right) \left(1 - \exp\left(-\frac{t_{go}}{\tau} \right) \right) \tag{29a}$$

$$t_{go} = \frac{\|\hat{\mathbf{r}}\|}{\|\hat{\mathbf{v}}\|} \tag{29b}$$

$$\hat{\mathbf{r}} = \begin{cases} \mathbf{r} - \begin{bmatrix} 0 & 0 & 15 \end{bmatrix}, & \text{if } r_z > 15 \\ \begin{bmatrix} 0 & 0 & r_z \end{bmatrix}, & \text{otherwise} \end{cases} \tag{29c}$$

$$\hat{\mathbf{v}} = \begin{cases} \mathbf{v} - \begin{bmatrix} 0 & 0 & -2 \end{bmatrix}, & \text{if } r_z > 15 \\ \mathbf{v} - \begin{bmatrix} 0 & 0 & -1 \end{bmatrix}, & \text{otherwise} \end{cases} \tag{29d}$$

$$\tau = \begin{cases} \tau_1, & \text{if } r_z > 15 \\ \tau_2, & \text{otherwise} \end{cases} \tag{29e}$$

Finally, we provide a terminal reward bonus when the lander reaches an altitude of zero, and the terminal position, velocity, attitude, and rotational velocity are within specified limits. The reward function is then given by the following:

$$r = \alpha \|\mathbf{v} - \mathbf{v}_{targ}\| + \beta \|\mathbf{F}_B\| + \gamma \,\text{any}\,(\mathbf{q}(t) > \mathbf{q}_{lim}) + \delta \sum_{i=1}^{3} -\max(0, \mathbf{q}_i - \mathbf{q}_{mgn_i}) + \eta \tag{30}$$

$$+ \kappa (r_z < 0 \text{ and } \|\mathbf{r}\| < r_{lim} \text{ and } \|\mathbf{v}\| < v_{lim} \text{ and all}(\mathbf{q} < \mathbf{q}_{\lim}) \text{ and all}(\boldsymbol{\omega} < \omega_{\lim}))$$

where the various terms are described in the following:

1. α weights a term penalizing the error in tracking the target velocity.

2. β weights a term penalizing control effort.

3. γ weights a term penalizing exceeding yaw, pitch, and roll limits. The "any" function is set to Boolean True if any elements of its argument are Boolean True (just like the

Python np.any() function). If any attitude component exceeds its limit, the episode is terminated.

4. δ weights a term that increases as the lander's attitude passes a threshold q_{mgn}; this gives the agent a hint that it is approaching the attitude limit q_{lim}

5. η is a constant positive term that encourages the agent to keep making progress along the trajectory. Since all other rewards are negative, without this term, an agent would be incentivized to violate the attitude constraint and prematurely terminate the episode to maximize the total discounted rewards received starting from the initial state.

6. κ is a bonus given for a successful landing, where terminal position, velocity, attitude, and rotational velocity are all within specified limits. The "all" function is set to Boolean True if all elements of its argument are Boolean True (just like the Python np.all() function). The limits are $r_{lim} = 5$ m, $v_{lim} = 2$ m, $q_{lim} = 0.2$ rad (except for yaw, which is not limited), and $w_{lim} = 0.2$ rad/s.

This reward function allows the agent to trade off between tracking the target velocity given in Eq. (29a), conserving fuel, satisfying the attitude constraints, and maximizing the reward bonus given for a good landing. Note that the constraints are not hard constraints such as might be imposed in an optimal control problem solved using collocation methods. However, the consequences of violating the constraints (a large negative reward and termination of the episode) are sufficient to ensure they are not violated once learning has converged. Hyperparameter settings and coefficients used in this work are given below in Table 5, note that due to lander symmetry, we do not impose any limits on the lander's yaw. The observation

Table 5. Hyperparameter Settings.

v_o (m/s)	τ_1 (s)	τ_2 (s)	α	β	γ	δ	η	κ	$\mathbf{q}_{lim}$ (rad)			$\mathbf{q}_{mgn}$ (rad)		
$\|\mathbf{v}_o\|$	20	100	-0.01	-0.05	-100	-20	0.01	10	$\begin{bmatrix} 2\pi \end{bmatrix}$	$\frac{7\pi}{16}$	$\frac{7\pi}{16}$	$\begin{bmatrix} 0 \end{bmatrix}$	$\frac{5\pi}{16}$	$\frac{5\pi}{16}$

given to the agent during learning and testing is given by

$$\text{obs} = \begin{bmatrix} \mathbf{v}_{\text{error}} & \mathbf{q} & \boldsymbol{\omega} & r_z & t_{\text{go}} \end{bmatrix} \tag{31}$$

where $\mathbf{v}_{\text{error}} = \mathbf{v} - \mathbf{v}_{\text{targ}}$, with $\mathbf{v}_{\text{targ}}$ given in Eq. (29a), the lander's estimated altitude, the time-to-go, as well as an estimate of the lander's attitude ($\mathbf{q}$) and rotational velocity ($\boldsymbol{\omega}$). Note that aside from the altitude, the lander translational coordinates do not appear in the observation. This results in a policy with good generalization in that the policy's behavior can extend to areas of the full state space that were not experienced during learning. To provide a robust final policy, we optimize with parameter uncertainty. Specifically, at the beginning of each episode both the lander's initial mass and the acceleration due to gravity

are perturbed to give a random value within 5% and 2% of nominal, respectively. In addition, we apply random uniform noise to the inertial tensor at the beginning of each episode. There is no physical significance to the inertia tensor noise; it is just a method of introducing parameter uncertainty in order to learn a robust policy.

Table 6. Parameter Uncertainty for Policy Optimization.

	min	max
Initial Mass (kg)	1900	2100
Gravitational Acceleration (m/s^2)	$\begin{bmatrix} 0.07 & 0.07 & 3.64 \end{bmatrix}$	$\begin{bmatrix} -0.07 & -0.07 & 3.79 \end{bmatrix}$
Inertia Tensor Diagonal Noise (kg-m^2)	-100	100
Inertia Tensor Off-Diagonal Noise (kg-m^2)	-10	10

It turns out that the when a terminal reward is used (as we do), it is advantageous to use a relatively large discount rate. However, it is also advantageous to use a lower discount rate for the shaping rewards. To our knowledge, a method of resolving this conflict has not been reported in the literature. In this work, we resolve the conflict by introducing a framework for accommodating multiple discount rates. Let γ_1 be the discount rate used to discount $r_1(k)$, the reward function term (as given in 30) associated with the κ coefficient). Moreover, let γ_2 be the discount rate used to discount $r_2(k)$, the sum of all other terms in the reward function. We can then rewrite the Eq. (24) and (23) in terms of these rewards and discount rates, as given by the following:

$$J(\mathbf{w}) = \sum_{i=1}^{M} \left(V_{\mathbf{w}}^{\pi}(\mathbf{x}_k^i) - \left[\sum_{\tau=k}^{T} \gamma_1^{\ell-k} r_1(\mathbf{u}_\ell^i, \mathbf{x}_\ell^i) + \gamma_2^{\ell-k} r_2(\mathbf{u}_\ell^i, \mathbf{x}_\ell^i) \right] \right)^2 \tag{32a}$$

$$A_{\mathbf{w}}^{\pi}(\mathbf{x}_k, \mathbf{u}_k) = \left[\sum_{\ell=k}^{T} \gamma_1^{\ell-k} r_1(\mathbf{u}_\ell, \mathbf{x}_\ell) + \gamma_2^{\ell-k} r_2(\mathbf{u}_\ell, \mathbf{x}_\ell) \right] - V_{\mathbf{w}}^{\pi}(\mathbf{x}_k) \tag{32b}$$

Although the approach is simple, the performance improvement is significant, although we had to give up the use of generalized advantage estimation [42] as it is not compatible with multiple discount rates. Without the use of multiple discount rates, the performance was actually worsened by including the terminal reward term.

III. Results

A. Policy Optimization

Rollouts are generated by the agent interacting with the environment for 120 episodes, with the resulting trajectories used to compute the advantages and update the value and policy function approximators. Optimization was terminated after 300,000 episodes. Learning curves are shown below. In Figure 1 (Figure 4 for 3-DOF optimization), we plot statistics

for the undiscounted rewards over 120 episodes, which is the number of episodes the agent accumulates before updating the policy and value function. "Steps" refers to the number of interactions between the agent and environment for the episode; this can be converted to the trajectory duration by multiplying by the navigation period of 0.2s. Figure 2 (Figure 5 for 3-DOF optimization) gives the policy entropy and the KL divergence between policy updates. We also plot the explained variance as a measure of how well the value function explains the observed returns; if the explained variance is 1, the value function perfectly explains the observed returns, if it is less than zero, you would have been better off predicting a constant value for the value of being in any state. Figure 3(a) and Figure 3(b) (Figures 6(a) and 6(b) for 3-DOF) plot the lander's end of episode position and velocity magnitudes as learning progresses, where again the statistics are accumulated over the 120 episodes used for the policy and value function updates. Note in Figures 3(c) through 3(d) that the standard deviation of the attitude and rotational velocity seems inconsistent with a good landing. This is partly due to the parameter uncertainty introduced during learning, and partly due to policy exploration. Although we terminated optimization at 300,000 episodes, the performance was still improving, and the standard deviation of the action was still around 4% for exploration. Since this is 4% of the maximum thrust, exploration noise still had a significant impact on the policy. During policy testing, we see that the magnitude of the attitude and rotational velocity is bounded more tightly.

B.   Policy Testing

To test the guidance policy, we simulate the policy for 10000 episodes using the same initial conditions as used for policy optimization, as given in Table 3. During testing, the dynamics model adds Gaussian noise to the force acting on the lander. The noise has a mean that is computed at the start of each episode uniformly distributed between -100 and 100 N and held constant during the episode. The noise standard deviation is 100 N. At the start of each episode, the nominal wet mass of 2000 kg is set to a uniformly distributed value between +/- 5% of nominal. Table 7 tabulates the touchdown statistics accumulated over testing. The glideslope statistic is given by the average $v_z/\|[v_x, v_y]\|$ over the final 2 m of descent; here a high value is desirable, as it indicates the lander's velocity is directed primarily in the downward direction. Note that due to symmetry in the lander design, we do not care about the terminal yaw value.

We compared the fuel efficiency of the 6-DOF RL agent to that of a 3-DOF RL agent using the same reward shaping function and to a 3-DOF DR/DV controller; the fuel statistics are at the bottom of Table 7. We use the DR/DV results as a proxy for optimal performance, as the algorithm is energy-optimal for the case of unlimited thrust (although we limit the thrust for the comparison). Because DR/DV has an unacceptable terminal glideslope (less than 1),
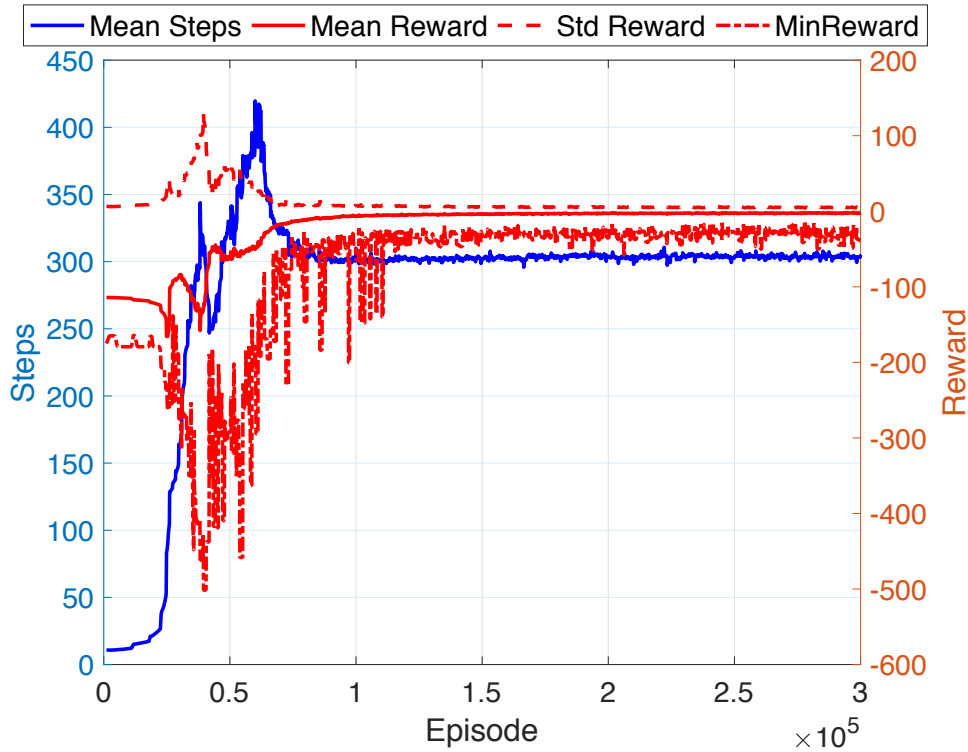
Figure 1. 6-DOF Optimization: Learning Curves.

Table 7. Touchdown Statistics (same initial condition range as for optimization).

	Mean	Std. Dev.	Min	Max
Downrange Position (m)	0.4	0.7	-2.9	4.5
Crossrange Position (m)	-0.1	0.7	-5.2	2.5
Downrange Velocity (m/s)	0.06	0.03	-0.04	0.14
Crossrange Velocity (m/s)	-0.01	0.04	-0.20	0.14
Elevation Velocity (m/s)	-0.93	0.08	-1.32	-0.49
Pitch (rad)	-0.016	0.010	-0.060	0.028
Roll (rad)	-0.003	0.011	-0.043	0.050
Rot. Velocity - Roll (rad/s)	-0.000	0.021	-0.105	0.084
Rot. Velocity - Pitch (rad/s)	-0.005	0.013	-0.083	0.051
Rot. Velocity - Yaw (rad/s)	0.000	0.000	0.000	0.000
Glideslope	22.02	17.77	7.50	851.54
Fuel Consumed - 6-DOF RL Policy (kg)	291	15	257	352
Fuel Consumed - 3-DOF RL Policy (kg)	291	14	259	358
Fuel Consumed - 3-DOF DR/DV (kg)	279	14	251	335

we used a piecewise trajectory with a single waypoint 15 m above the landing site, similar to the approach we used for the RL policy, which achieved a minimum glideslope of close to 8. This added about 20 kg to the DR/DV fuel consumption. We see a 4% increase in fuel consumption for the 6-DOF RL agent as compared to 3-DOF DR/DV. To put this increase in perspective, note that tracking the DR/DV trajectory in a 6-DOF environment would
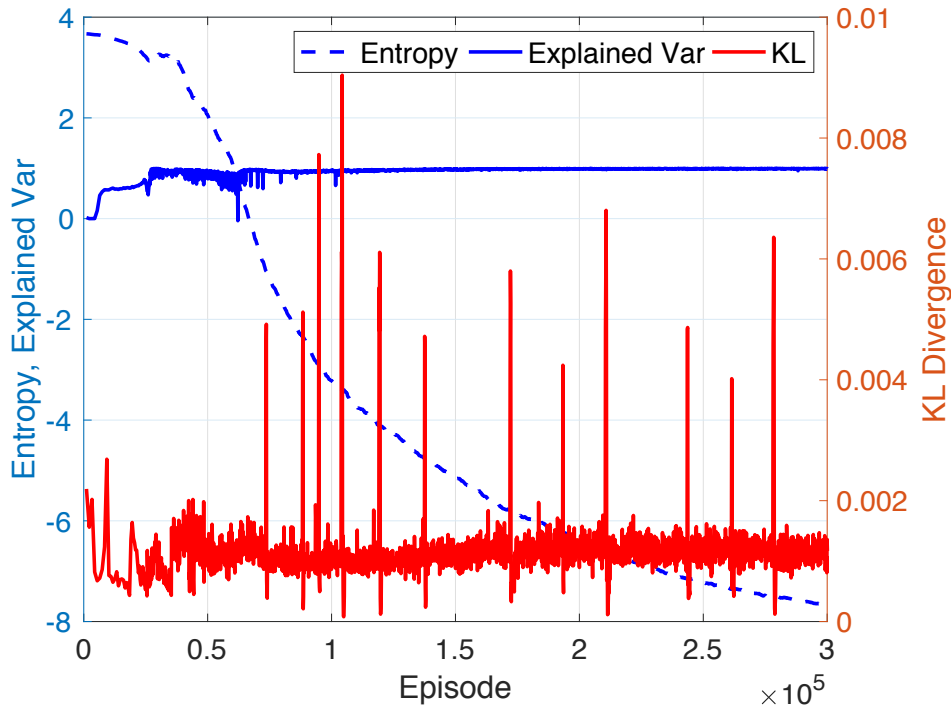
Figure 2. 6-DOF Optimization: Learning Statistics.

certainly increase full consumption. Finally, note that the 6-DOF RL agent achieves fuel efficiency close to that of the 3-DOF RL agent; this tells us that the selected reward shaping function has a critical impact on fuel efficiency. There is certainly room for improvement here, as the exponential decrease in the magnitude of the target velocity is probably suboptimal. Indeed, an optimal trajectory with an initial position far from the target landing site might accelerate towards the target to quickly close the downrange and crossrange distance in order to reduce trajectory time and use less fuel to maintain altitude.

Divert functionality was tested by running 5000 Monte Carlo simulations over the same initial conditions, but triggering a divert of 800 m downrange and 800 m crossrange when the lander reached an altitude of 1500 m. On average, the divert maneuver resulted in a 30 kg average increase in fuel consumption but otherwise did not impact performance.

A sample trajectory is plotted in Figure 7, where x, y, and z are the downrange, crossrange, and altitude trajectory components in the target-centered reference frame. Thrust is shown in the inertial frame. The left side plot that is second from the top gives the altitude as a function of the norm of the cross range and downrange position. This plot is from initial conditions of 1500 m to -70 m/s downrange, -500 m to -30 m/s crossrange, and 2400m to -90 m/s elevation. For comparison, we also plot (Figure 8) a trajectory from the 3-DOF RL policy that starts from the same initial conditions. Note that there are a couple of places
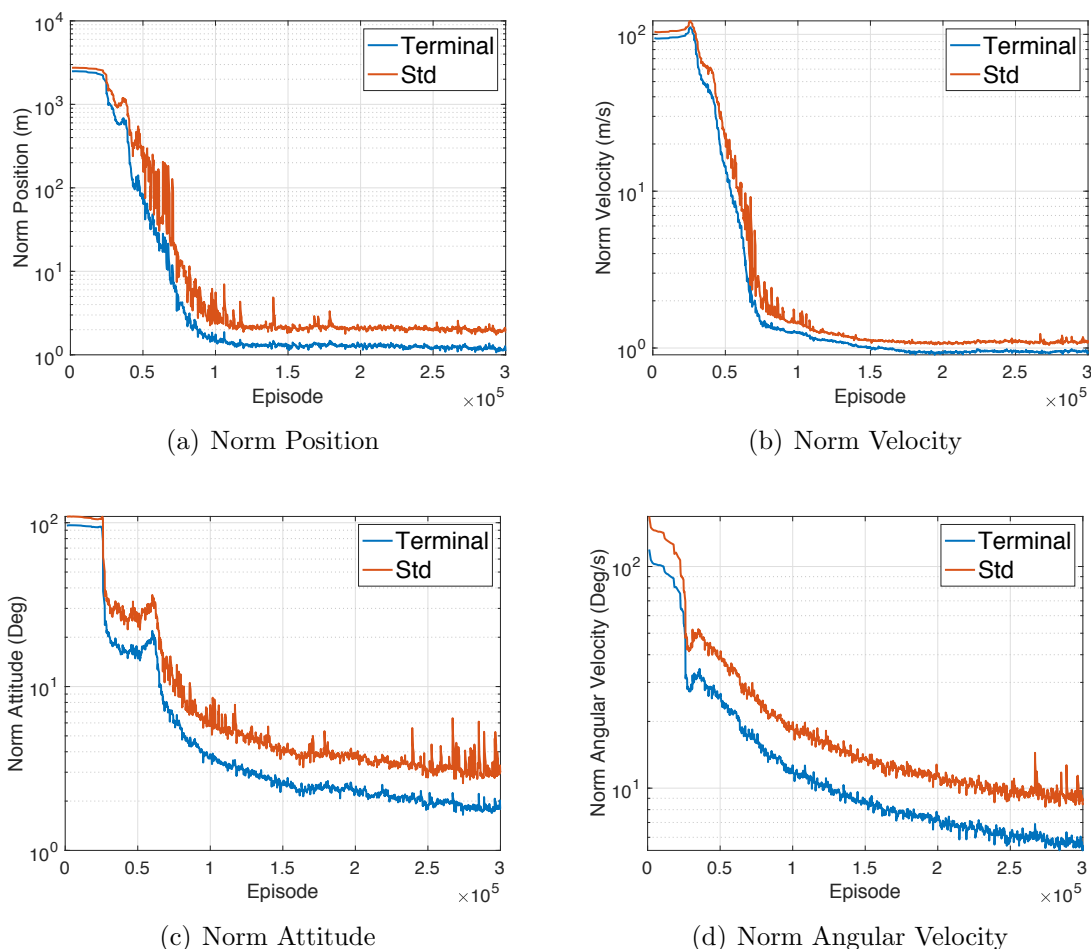
(a) Norm Position

(b) Norm Velocity

(c) Norm Attitude

(d) Norm Angular Velocity

Figure 3. Optimization: Lander Position, Velocity, Attitude, and Angular Velocity at end of Episode.

where the system exhibits small oscillations in the lander's thrust vector. Training the agent longer possibly have a smoothing effect on the resulting solution.

From Figures 7 and 8, it is apparent that the 3-DOF and 6-DOF position and velocity trajectories are almost identical. The thrust differs toward the start and end of the landing due to the need to adjust attitude in the 6-DOF case in order to change the thrust direction, although overall the thrust trajectories are similar. Specifically, since the simulations begin with an average pitch of 45 degrees, the policy rotates the lander to a smaller pitch angle in order to allow a burn with the thrust vector pointing downwards to reduce the magnitude of the vertical velocity component. In addition, the policy must rotate the lander to allow a burn that reduces the crossrange velocity until the lander is on a trajectory pointing towards the target. Towards the end of the trajectory, where the lander must transition to a vertical descent, there is another area where the thrust trajectories differ.

We also re-tested the policy over two extended deployment ellipses, as shown in Table
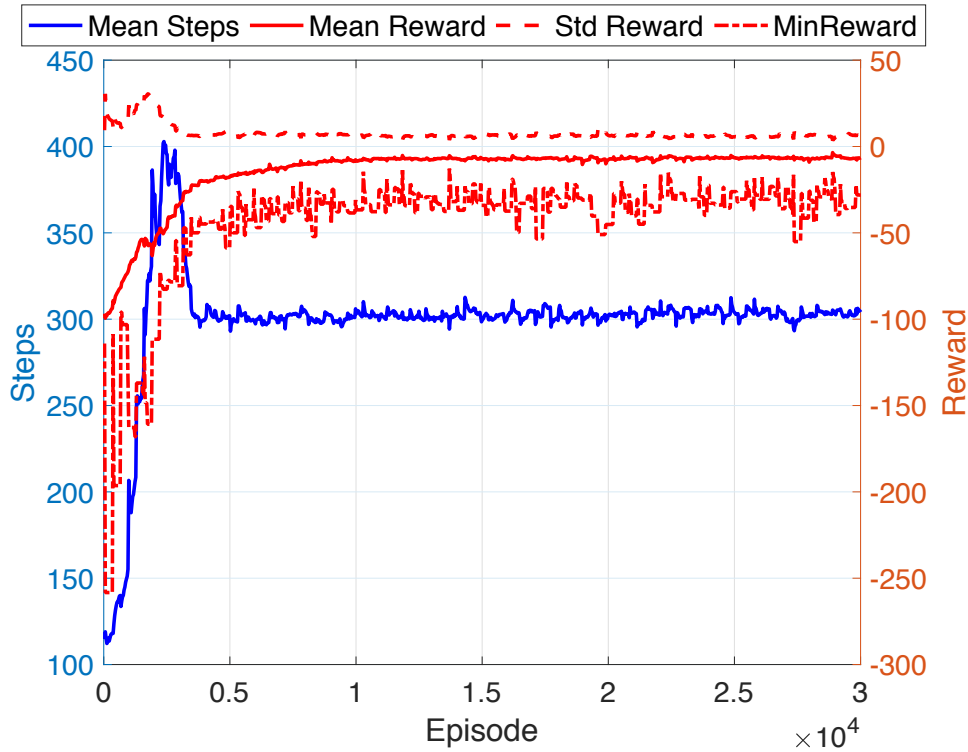
Figure 4. 3-DOF Optimization: Learning Curves.

8, where we extend the deployment ellipse to 9 km^2, and again to 12 km^2. Here we use the same environmental noise and mass uncertainty as in the testing of the policy over the 4 km^2 deployment ellipse. Note that in order to obtain satisfactory performance over the 12 km^2 deployment ellipse, we had to raise the altitude of the deployment ellipse to 3000m. Importantly, this is the same policy that was optimized over the initial conditions given in Table 3, i.e., a 4 km^2 deployment ellipse. Landing statistics over 20,000 episodes were similar to that given in Table 7, but average and maximum fuel consumption increased. Since the reward shaping function does not require the lander's full translational state, the policy generalizes quite well to regions of state space not experienced during optimization. Figures 10(a) and 10(b) illustrate 100 randomly sampled trajectories using the 9 km^2 deployment ellipse for the 3-DOF and 6-DOF policies, respectively.

IV.   Comparison with GPOPS Solution

Figure 9 shows experimental results comparing the 3-DOF and 6-DOF agents to solution provided by GPOPS [43], an optimal control solver, using a 3-DOF problem formulation with the same initial conditions used to generate Figures 7 and 8. The GPOPS solution, which also required a vertical descent over the final 15 m of altitude loss, had a fuel consumption
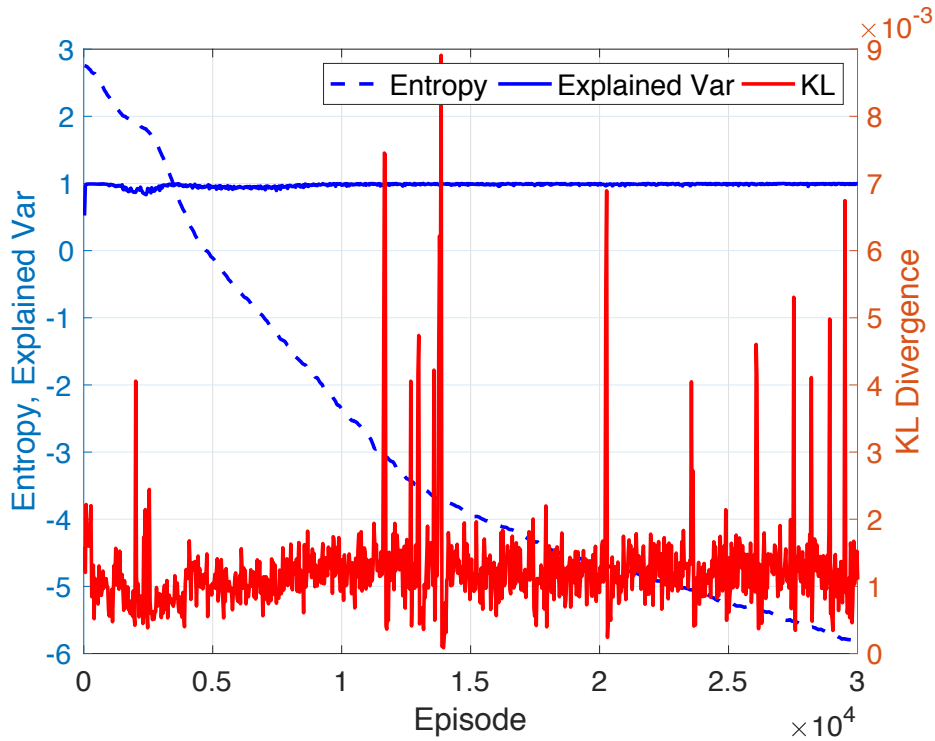
Figure 5.   3-DOF Optimization: Learning Statistics.

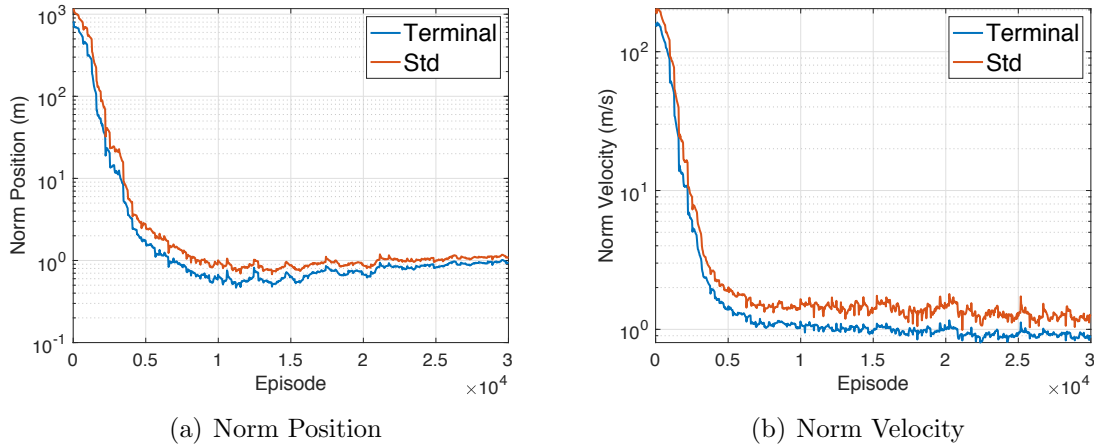(a) Norm Position

(b) Norm Velocity

Figure 6.   3-DOF Optimization: Lander Position, Velocity, Attitude, and Angular Velocity at end of Episode.

of 250 kg, as compared to 290 kg for the 6-DOF and 3-DOF policies, making the RL policy fuel consumption 18 percent higher than GPOPS optimal. First, note that the GPOPS trajectories are open loop, and when combined with a trajectory tracking controller the fuel consumption will be higher. Second, since the 6-DOF and 3-DOF policies have almost identical fuel consumption, we can attribute the difference in fuel efficiency between the
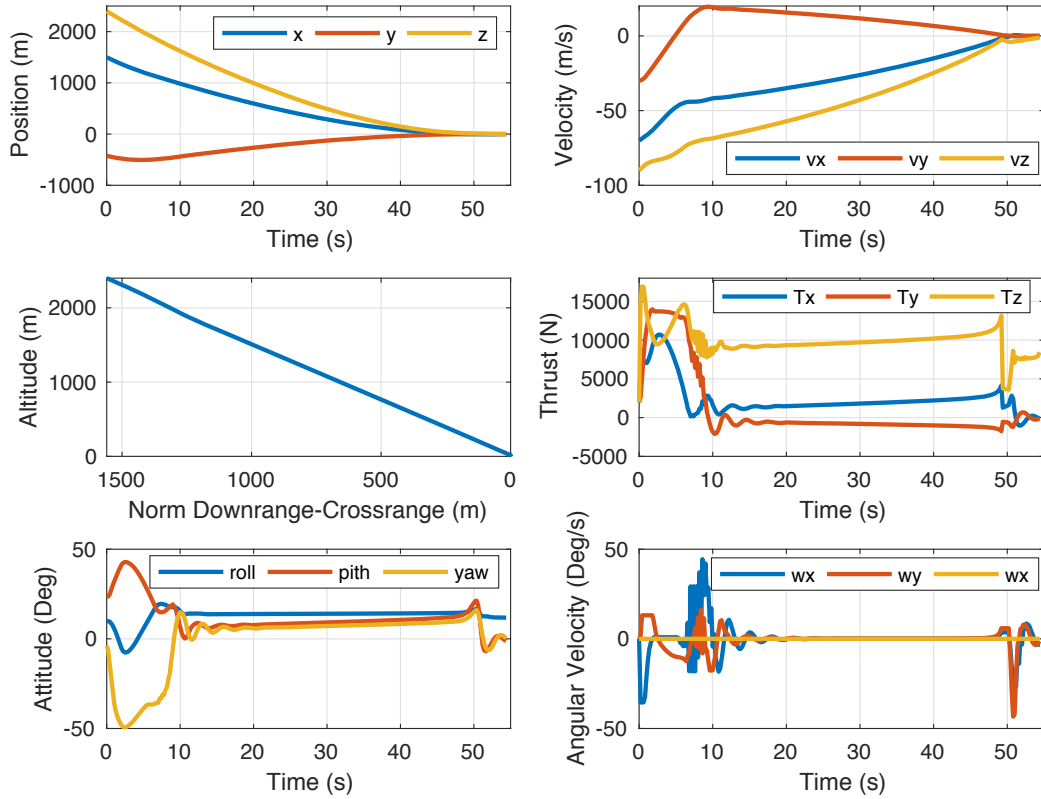
Figure 7. 6-DOF Sample Trajectory.

Table 8. Lander Extended Initial Conditions for Optimization.

	Velocity		Position	
	min (m/s)	max (m/s)	min (m)	max (m)
9 km Deployment Ellipse				
Downrange	-70	-10	0	3000
Crossrange	-30	30	-1500	1500
Elevation	-90	-70	2400	2500
12 km Deployment Ellipse				
Downrange	-70	-10	0	4000
Crossrange	-30	30	-1500	1500
Elevation	-90	-70	2900	3100

RL policies and optimal to the reward shaping function used during optimization, which although effective, is not optimal. Consequently, future work to improve the fuel efficiency of the RL derived integrated guidance and control system should focus on the reward shaping function. Since this function takes the form of a velocity field, it may prove productive to learn a reward shaping function parameterized as a neural network from optimal trajectories. Figures 11(a), 11(b), 11(c), and 11(d) show the fuel performance for the 3-DOF and 6-DOF
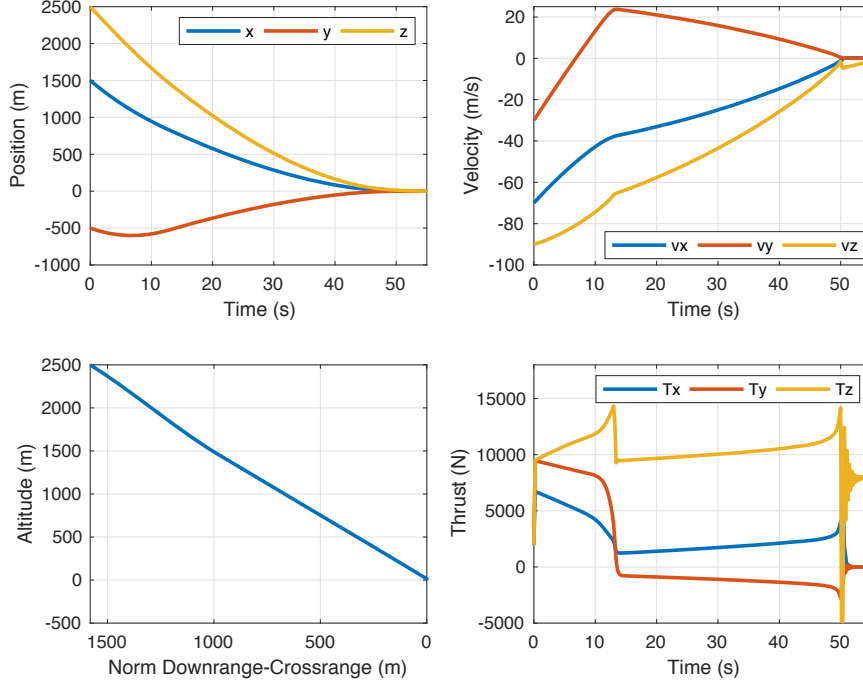
**Figure 8. Same Initial Conditions using 3-DOF policy.**

Table 9. Extended Range Fuel Consumption.

|	Mean	Std. Dev.	Min	Max
9 km Deployment Ellipse				
Fuel Consumed - 6-DOF RL Policy (kg)	308	25	257	412
Fuel Consumed - 3-DOF RL (kg)	309	25	260	400
Fuel Consumed - 3-DOF DR/DV (kg)	297	23	252	381
12 km Deployment Ellipse				
Fuel Consumed - 6-DOF RL Policy (kg)	340	31	282	468
Fuel Consumed - 3-DOF RL (kg)	341	30	285	437
Fuel Consumed - 3-DOF DR/DV (kg)	319	27	268	414

policies over the 9 square kilometer deployment ellipse. Interestingly, the zero divert fuel consumption in Figures 11(a) and 11(b) is around 240kg, close to that of the MSL if we subtract fuel used in the sky-crane maneuver.

V.   Conclusion

We have applied a novel approach to the Mars landing powered descent phase problem using an integrated guidance and control system developed using reinforcement learning, and validated the approach through Monte Carlo simulation. Coupled with a navigation system such as that described in [3], this guidance and control system can achieve pinpoint accuracy
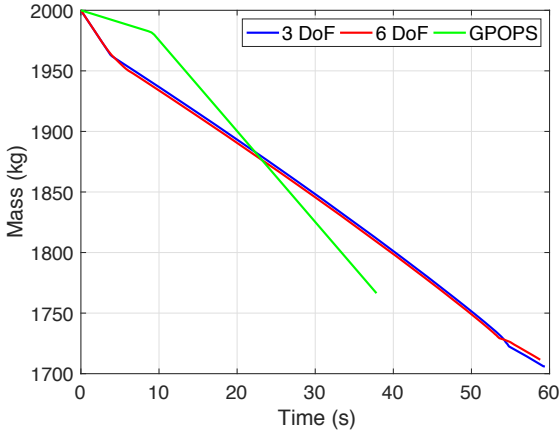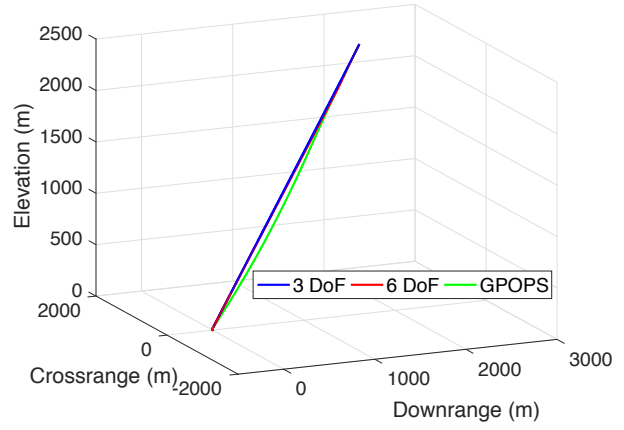
(a) 3-DOF Trajectories   (b) 6-DOF Trajectories

Figure 9. Comparison of 3-DOF and 6-DOF policy with GPOPS solution.

(a) 3-DOF Trajectories   (b) 6-DOF Trajectories

Figure 10. 100 Sample Trajectories over 9 square km Deployment Ellipse.

and a soft landing with minimal deviation from an ideal landing attitude and rotational velocity, with large divert distance capability. As compared to current practice, this system provides a significant increase in landing precision. The policy was shown to be robust to noise and parameter uncertainty. Compared to proposed systems such as that described in [9], this system has the advantage of not requiring a cone-shaped glideslope constraint, allowing the targeting of locations such as the bottom of a deep crater. The computational requirements of running the policy are modest, with only four matrix multiplications required to map the estimated state from the navigation system to a body frame thrust command,
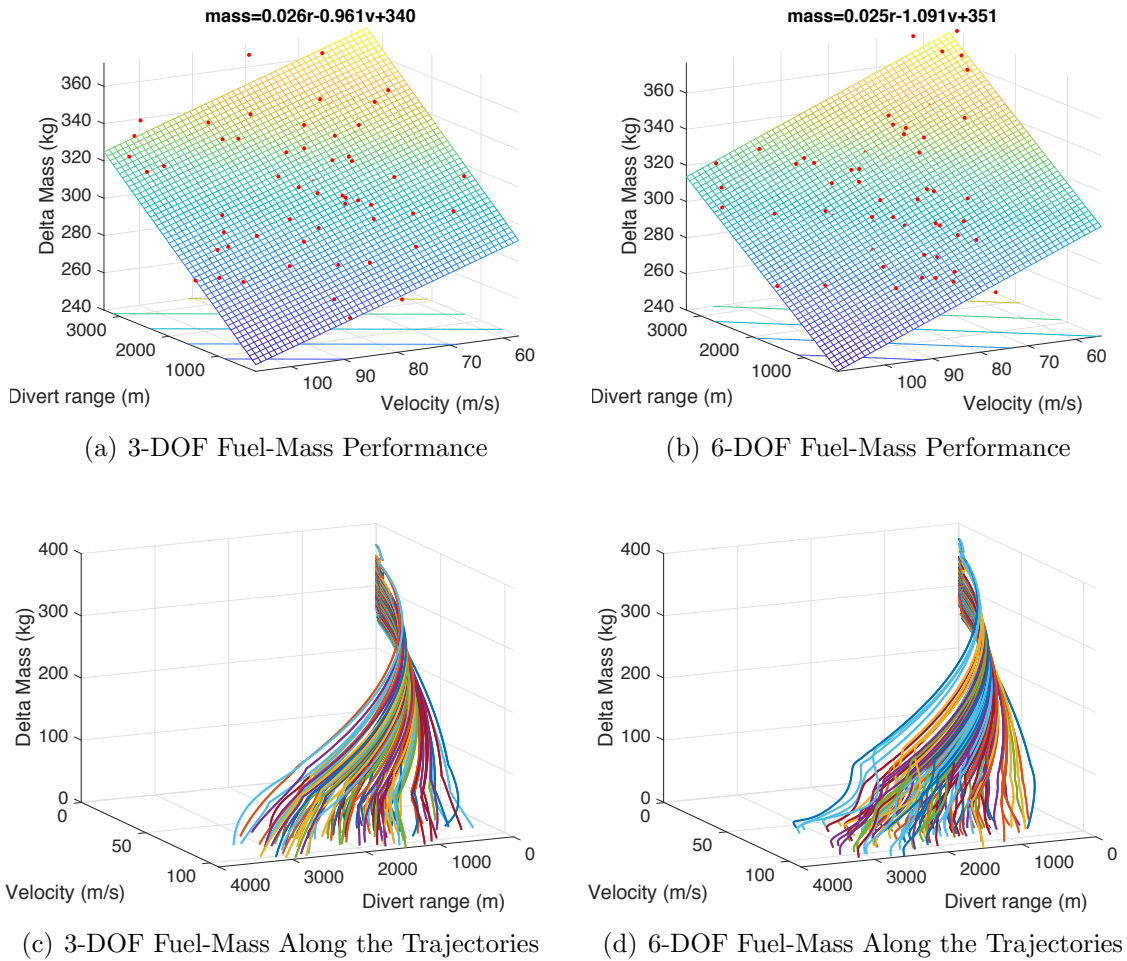
mass=0.026r-0.961v+340

(a) 3-DOF Fuel-Mass Performance

mass=0.025r-1.091v+351

(b) 6-DOF Fuel-Mass Performance

(c) 3-DOF Fuel-Mass Along the Trajectories

(d) 6-DOF Fuel-Mass Along the Trajectories

Figure 11. 6-DOF and 3-DOF Fuel-Mass Performance.

and then a trained policy should have no problem running on the current generation of flight computers.

VI.   Acknowledgment

References

[1]R. Shotwell, "Phoenixthe first Mars Scout mission," *Acta Astronautica*, Vol. 57, No. 2-8, 2005, pp. 121–134.

[2]R. D. Braun and R. M. Manning, "Mars exploration entry, descent, and landing challenges," *Journal of spacecraft and rockets*, Vol. 44, No. 2, 2007, pp. 310–323.

[3]B. Gaudet and R. Furfaro, "A navigation scheme for pinpoint mars landing using radar altimetry, a digital terrain model, and a particle filter," *2013 AAS/AIAA Astrodynamics Specialist Conference, Astrodynamics 2013*, Univelt Inc., 2014.

[4]R. H. Battin, *An Introduction to the Mathematics and Methods of Astrodynamics, revised edition.* American Institute of Aeronautics and Astronautics, 1999.

[5]C. D'Souza and C. D'Souza, "An optimal guidance law for planetary landing," *Guidance, Navigation, and Control Conference*, 1997, p. 3709.

[6]A. D. Steltzner, P. D. Burkhart, A. Chen, K. A. Comeaux, C. S. Guernsey, D. M. Kipp, L. V. Lorenzoni, G. F. Mendeck, R. W. Powell, T. P. Rivellini, *et al.*, "Mars science laboratory entry, descent, and landing system overview," 2010.

[7]B. Acikmese, J. Casoliva, J. Carson, and L. Blackmore, "G-fold: A real-time implementable fuel optimal large divert guidance algorithm for planetary pinpoint landing," *Concepts and Approaches for Mars Exploration*, Vol. 1679, 2012.

[8]G. Singh, A. M. SanMartin, and E. C. Wong, "Guidance and control design for powered descent and landing on Mars," *Aerospace Conference, 2007 IEEE*, IEEE, 2007, pp. 1–8.

[9]B. Açıkmeşe, J. M. Carson, and L. Blackmore, "Lossless convexification of nonconvex control bound and pointing constraints of the soft landing optimal control problem," *IEEE Transactions on Control Systems Technology*, Vol. 21, No. 6, 2013, pp. 2104–2113.

[10]B. Acikmese and S. R. Ploen, "Convex programming approach to powered descent guidance for mars landing," *Journal of Guidance, Control, and Dynamics*, Vol. 30, No. 5, 2007, pp. 1353–1366.

[11]M. Szmuk and B. Acikmese, "Successive Convexification for 6-DoF Mars Rocket Powered Landing with Free-Final-Time," *2018 AIAA Guidance, Navigation, and Control Conference*, 2018, p. 0617.

[12]P. Lu, "Propellant-Optimal Powered Descent Guidance," *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 4, 2017, pp. 813–826.

[13]D. Dueri, B. Açıkmeşe, D. P. Scharf, and M. W. Harris, "Customized real-time interior-point methods for onboard powered-descent guidance," *Journal of Guidance, Control, and Dynamics*, Vol. 40, No. 2, 2016, pp. 197–212.

[14]L. Blackmore, B. Acikmese, and D. P. Scharf, "Minimum-landing-error powered-descent guidance for Mars landing using convex optimization," *Journal of guidance, control, and dynamics*, Vol. 33, No. 4, 2010, pp. 1161–1171.

[15]P. Lu, R. R. Sostaric, and G. F. Mendeck, "Adaptive Powered Descent Initiation and Fuel-Optimal Guidance for Mars Applications," *2018 AIAA Guidance, Navigation, and Control Conference*, 2018, p. 0616.

[16]R. Yanushevsky, *Modern missile guidance.* CRC Press, 2007.

[17]M. Xin, S. Balakrishnan, and E. J. Ohlmeyer, "Integrated Guidance and Control of Missiles With $theta hbox - D$ Method," *IEEE Transactions on Control Systems Technology*, Vol. 14, No. 6, 2006, pp. 981–992.

[18]H. Yan and H. Ji, "Integrated guidance and control for dual-control missiles based on small-gain theorem," *Automatica*, Vol. 48, No. 10, 2012, pp. 2686–2692.

[19]N. F. Palumbo, B. E. Reardon, and R. A. Blauwkamp, "Integrated guidance and control for homing missiles," *Johns Hopkins APL Technical Digest*, Vol. 25, No. 2, 2004, pp. 121–139.

[20]Y. B. Shtessel and I. A. Shkolnikov, "Integrated guidance and control of advanced interceptors using

second order sliding modes," *Decision and Control, 2003. Proceedings. 42nd IEEE Conference on*, Vol. 5, IEEE, 2003, pp. 4587–4592.

[21]Y. Zhang and T. Yang, "An Integrated Trajectory Guidance and Attitude Control Law with Enhanced Anti-Disturbance Capability for Mars Pinpoint Landing," *2017 4th International Conference on Information Science and Control Engineering (ICISCE)*, IEEE, 2017, pp. 840–846.

[22]D. R. Wibben and R. Furfaro, "Integrated guidance and attitude control for pinpoint lunar guidance using Higher Order Sliding Modes," *22nd AAS/AIAA Space Flight Mechanics Meeting*, 2012.

[23]C. J. Watkins and P. Dayan, "Q-learning," *Machine learning*, Vol. 8, No. 3-4, 1992, pp. 279–292.

[24]V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, *et al.*, "Human-level control through deep reinforcement learning," *Nature*, Vol. 518, No. 7540, 2015, p. 529.

[25]D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller, "Deterministic policy gradient algorithms," *ICML*, 2014.

[26]T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.

[27]R. J. Williams, "Simple statistical gradient-following algorithms for connectionist reinforcement learning," *Machine learning*, Vol. 8, No. 3-4, 1992, pp. 229–256.

[28]J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," *International Conference on Machine Learning*, 2015, pp. 1889–1897.

[29]J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[30]J. Achiam, D. Held, A. Tamar, and P. Abbeel, "Constrained policy optimization," *arXiv preprint arXiv:1705.10528*, 2017.

[31]A. Y. Ng, *Shaping and policy search in reinforcement learning*. PhD thesis, University of California, Berkeley, 2003.

[32]S. L. Waslander, G. M. Hoffmann, J. S. Jang, and C. J. Tomlin, "Multi-agent quadrotor testbed control design: Integral sliding mode vs. reinforcement learning," *Intelligent Robots and Systems, 2005.(IROS 2005). 2005 IEEE/RSJ International Conference on*, IEEE, 2005, pp. 3712–3717.

[33]R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*, Vol. 1. MIT press Cambridge, 1998.

[34]B. Gaudet and R. Furfaro, "Adaptive pinpoint and fuel efficient mars landing using reinforcement learning," *IEEE/CAA Journal of Automatica Sinica*, Vol. 1, No. 4, 2014, pp. 397–411.

[35]M. D. Shuster, "A Survey of Attitude Representations," *Journal of the Astronautical Sciences*, Vol. 41, Oct.-Dec. 1993, pp. 439–517.

[36]J. L. Junkins and H. Schaub, *Analytical mechanics of space systems*. American Institute of Aeronautics and Astronautics, 2009.

[37]M. Baccouche, F. Mamalet, C. Wolf, C. Garcia, and A. Baskurt, "Sequential deep learning for human action recognition," *International Workshop on Human Behavior Understanding*, Springer, 2011, pp. 29–39.

[38]D. C. Sorensen, "Newtons method with a model trust region modification," *SIAM Journal on Numerical Analysis*, Vol. 19, No. 2, 1982, pp. 409–426.

[39]S. Kullback and R. A. Leibler, "On information and sufficiency," *The annals of mathematical statistics*, Vol. 22, No. 1, 1951, pp. 79–86.

[40]D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[41]A. Y. Ng, S. J. Russell, *et al.*, "Algorithms for inverse reinforcement learning.," *Icml*, 2000, pp. 663–670.

[42]J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, "High-dimensional continuous control using generalized advantage estimation," *arXiv preprint arXiv:1506.02438*, 2015.

[43]A. V. Rao, D. A. Benson, C. Darby, M. A. Patterson, C. Francolin, I. Sanders, and G. T. Huntington, "Algorithm 902: Gpops, a matlab software for solving multiple-phase optimal control problems using the gauss pseudospectral method," *ACM Transactions on Mathematical Software (TOMS)*, Vol. 37, No. 2, 2010, p. 22.