AAS 24-048

TOWARDS PREDICTIVE CONTROL WITH ATMOSPHERIC ADAPTATION FOR MARTIAN ENTRY VEHICLES

Robert D. Halverson[*], Ryleigh McGiveron[†], Maziar S. Hemati[‡], and Ryan J. Caverly[§]

Atmospheric adaptation for Martian entry missions is a crucial aspect of guidance and control schemes that takes into account density variations from a nominal atmosphere model. This work considers the identification of a time-varying parameter that defines the ratio between true atmospheric density with that of an approximate model based on dispersed Mars-GRAM atmospheric data. A 3 degree-of-freedom model is used that takes into account the important hypersonic entry coefficients, and Monte Carlo simulations are performed to demonstrate the effect varying density has on the entry trajectory profile. Two techniques for predictive modeling using a least-squares estimation approach are introduced that provide static and time-varying methods for estimating the true atmospheric density. It is found that the predictive models provide an accurate estimate of the atmospheric density in the low portions of the Mars atmosphere. Results presented in this paper will inform the development of a predictive control scheme that incorporates an adaptive atmospheric model.

INTRODUCTION

This paper develops methods for real-time identification of the atmospheric model for a hypersonic entry vehicle, which may be used in conjunction with predictive control in the highly uncertain Martian atmospheric environment. Synthesizing a controller for a system operating in an uncertain environment is challenging. It is critical to develop an accurate dynamic model to be used in offline design and verification of guidance and control algorithms. Meanwhile, the environment in an EDL (entry, descent, and landing) application is often stochastic, which greatly complicates the design of guidance and control algorithms. The presence of uncertainty in the atmospheric qualities coupled with the extremely low air density of the Martian atmosphere causes significant degradation in landing precision.[1]

As of 2017, there were 21 total Martian landing attempts with several failing during launch or cruise. Of the 16 that reached the planet's sensible atmosphere, 8 were deemed successful landings whereas the other 8 failed.[2] These missions all required an autonomous guidance and control system that was heavily reliant on an understanding of Mars' atmospheric properties—primarily the air density. The Martian atmosphere is not as well studied as that of the Earth's,[3] where there still exists a heightened uncertainty regarding the qualities in a day-to-day and hour-to-hour timescale. Ensuring dynamic stability and achieving precise, reliable targeted reentry is still a challenging problem

[*]Ph.D. Student, Department of Aerospace Engineering & Mechanics. University of Minnesota, Minneapolis, MN 55455, USA.

[†]Undergraduate Student, Department of Aerospace Engineering & Mechanics. University of Minnesota, Minneapolis, MN 55455, USA.

[‡]Associate Professor, Department of Aerospace Engineering & Mechanics. University of Minnesota, Minneapolis, MN 55455, USA.

[§]Assistant Professor, Department of Aerospace Engineering & Mechanics. University of Minnesota, Minneapolis, MN 55455, USA.

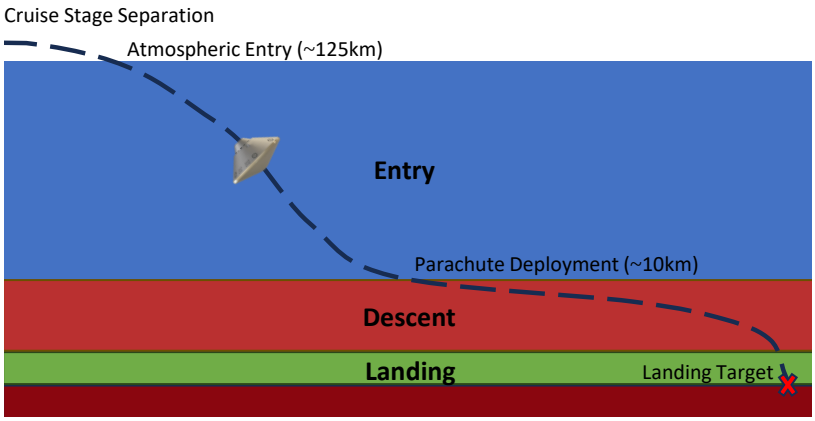

Figure 1: Entry, descent, and landing cartoon representation. Altitudes are representative of a nominal Martian EDL profile.

with significant room for improvement even close to home on Earth.[4,5] The low atmospheric density of the Martian atmosphere coupled with large environmental uncertainties (e.g., density, wind, dust storms, etc.) provides a massive challenge when synthesizing a guidance and control method.

Several approaches to autonomous guidance and control for atmospheric EDL exist that aim to deliver a precise landing. Legacy methods include either analytical or numerical predictor-corrector algorithms for entry guidance.[6] Recently, a convex predictor-corrector algorithm for entry guidance was introduced that allows for flexibility in vehicle design including control inputs and vehicle dynamics.[7] Other examples that utilize convex optimization methods include minimum fuel trajectory optimization via finite-dimensional second-order cone programming,[8] or model predictive control (MPC) techniques such as constrained,[9] robust,[10] or linear pseudospectral MPC.[11] There also exist many non-convex optimal control algorithms specific to a Mars pinpoint landing landing mission.[12,13] Other recent work is specifically related to entry vehicle guidance and control in uncertain environments.[14,15] Many of the EDL guidance & control solutions listed here focus on improving either performance, or robustness to model uncertainty, but not both. Very few existing approaches are capable of simultaneously achieving high performance and robustness in the presence of large amounts of environmental uncertainty.

A guidance trajectory is often developed offline using the nominal expected atmospheric density, which is then used as a reference trajectory that is tracked by the vehicle's control systems to lead it towards a final desired latitude and longitude before deployment of the parachute. Entry vehicles deploy parachutes as early as possible to provide a longer period of deceleration, however they must reach an airspeed slow enough for the parachute to survive.[16] In Mars EDL missions, this usually takes place at an altitude of 10km. Given the limited control authority under parachute, the deployment location should be reached as accurately and precisely as possible to set up the latter descent and landing portions for success. See Figure 1 for a representation of the normal EDL processes for a Martian landing.

This work aims to address one specific and important challenge in a Martian entry mission—imprecise knowledge of the true atmospheric density, and the capability to adapt to changes from a nominal atmospheric model. This is known as atmospheric adaptation. Previous work for Mars atmospheric adaptation has included using Kalman filtering online with a Convex Predictor-corrector

Entry Guidance algorithm,[17] or other Kalman filtering techniques for drag-modulated spacecraft for Earth reentry.[18] This paper is a stepping-stone towards using predictive modeling methods via parameter identification to approach the atmospheric adaptation problem. These methods can be expanded to be used in conjunction with a model predictive control architecture to inform an updated system model in real-time for use in the predictive control policy. Further, recent work has identified the efficacy of including atmospheric adaptation schemes online with a predictive control policy.[19] The results presented in this paper offer a step towards providing a systematic EDL control method that does not simply acknowledge uncertainty, but accounts and corrects for it while refining an environmental model.

This paper proceeds with a description of the simulation environment developed to support this work both in data collection and analysis. Atmospheric qualities including density, wind, dust storms, and dispersions on important values are generated via Mars-GRAM.[20] A simple system identification technique for density error estimation is introduced through a least-squares estimation approach, called sliding window batch estimation. This static estimate technique is expanded into time-varying predictive modeling through a fit to linear or nonlinear basis functions. The density error estimation results and subsequent discussion are presented, and future work to extend these results to be used in conjunction with a predictive control policy is described.

ENTRY VEHICLE DYNAMICS

This work considers a blunt-body entry vehicle in the Martian atmosphere, similar to that used in previous NASA missions.[1,21] Generally, these vehicles are capable of modulating their bank angle $\sigma \in \mathbb{R}$, however this work only considers the ballistic trajectory of an entry vehicle in the Martian atmosphere to test parameter identification techniques. For this reason, the bank angle is held to $\sigma = 0$ for all time t. The dynamics of the entry vehicle are described in a Mars-fixed Cartesian reference frame—a parameterization that is very common in powered-descent guidance literature given its well-posed numerical scaling.[8,22,23] The dynamics of the entry vehicle with this state representation are the following:[17]

$$\begin{aligned} \dot{\mathbf{r}} &= \mathbf{v}, \\ \dot{\mathbf{v}} &= \mathbf{a}_L + \mathbf{a}_D + \mathbf{a}_g - 2(\boldsymbol{\omega}^\times \mathbf{v}) - \boldsymbol{\omega}^\times(\boldsymbol{\omega}^\times \mathbf{r}), \end{aligned} \tag{1}$$

where $\mathbf{r}$ and $\mathbf{v}$ are the position and velocity vectors in a Mars-fixed frame in m and m/s, respectively, $\boldsymbol{\omega}$ is the angular velocity of Mars in rad/s, and $\mathbf{a}_L$, $\mathbf{a}_D$, and $\mathbf{a}_g$ are acceleration vectors due to lift, drag, and gravity, respectively. Regarding notation, boldface letters are representative of physical vectors resolved in the Mars-fixed Cartesian reference frame. The cross operator, $(\cdot)^\times : \mathbb{R}^3 \to \mathfrak{so}(3)$, is defined as

$$\mathbf{a}^\times = -\mathbf{a}^{\times^\mathsf{T}} = \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix},$$

where $\mathbf{a}^\mathsf{T} = [a_1 \;\; a_2 \;\; a_3]$ and $\mathfrak{so}(3) = \{\mathbf{S} \in \mathbb{R}^{3\times 3} \,|\, \mathbf{S} + \mathbf{S}^\mathsf{T} = \mathbf{0}\}$.

The magnitude of lift and drag accelerations are computed as

$$L = \frac{1}{2m}\rho A C_L ||\mathbf{v}_{\text{rel}}||^2, \tag{2}$$

$$D = \frac{1}{2m}\rho A C_D ||\mathbf{v}_{\text{rel}}||^2, \tag{3}$$

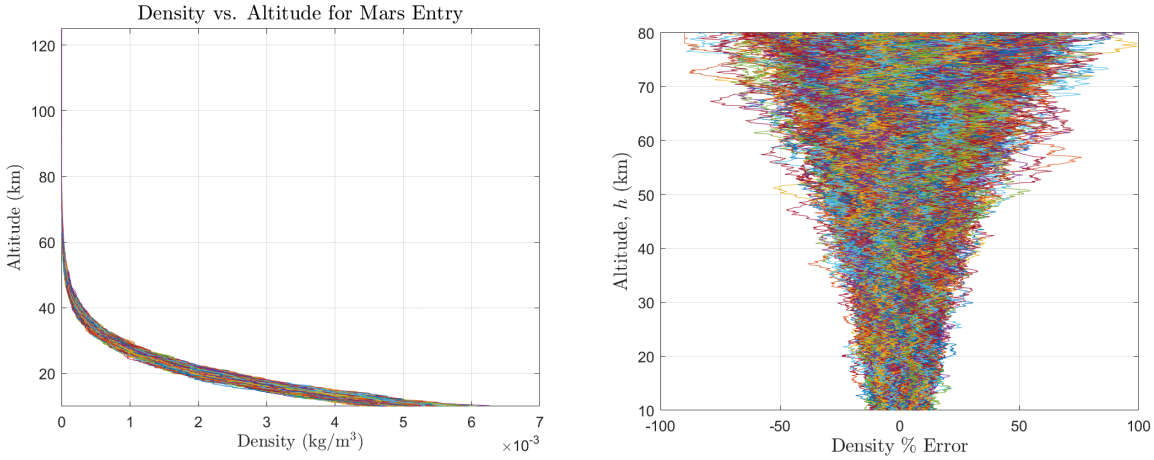

Figure 2: Relation of density and altitude for $N = 3000$ Monte Carlo runs from Mars-GRAM, both (a) density and (b) density error vs. altitude.

where m is the mass of the entry vehicle in kg, ρ is the atmospheric density in kg/m^3, A is the entry vehicle reference area in m^2, and $\mathbf{v}_{\text{rel}} = \mathbf{v} + \mathbf{w}$ is the relative velocity of the vehicle to the wind where $\mathbf{w}$ is the wind vector expressed in the Mars-fixed frame. The hypersonic coefficients C_L and C_D are evaluated through the Newtonian flow theory for a blunted body. Given a constant angle of attack, these coefficients will be constant. More details on these expressions can be found in Ref. 17.

The direction of the drag acceleration is opposite that of the relative velocity, expressed as:

$$\mathbf{a}_D = -D\frac{\mathbf{v}_{\text{rel}}}{||\mathbf{v}_{\text{rel}}||}.$$

Lift acceleration depends greatly on the bank angle σ, which becomes simplified via the assumption of zero bank angle. Aerodynamic basis vectors, $\hat{\mathbf{e}}_1 \in \mathbb{R}^3$ and $\hat{\mathbf{e}}_2 \in \mathbb{R}^3$, are constructed via the following:

$$\hat{\mathbf{e}}_1 = \frac{\mathbf{r}^{\times}\mathbf{v}_{\text{rel}}}{||\mathbf{r}^{\times}\mathbf{v}_{\text{rel}}||},$$
$$\hat{\mathbf{e}}_2 = \frac{\mathbf{v}_{\text{rel}}^{\times}\hat{\mathbf{e}}_1}{||\mathbf{v}_{\text{rel}}^{\times}\hat{\mathbf{e}}_1||},$$

which are used to solve for the direction of the lift acceleration as

$$\mathbf{a}_L = L\hat{\mathbf{e}}_2.$$

Acceleration due to gravity ($\mathbf{a}_g$) is calculated using the `gravitysphericalharmonic.m` function in MATLAB, which utilizes the Mars GMM2B model.[24] An order of 2 is used in calculating the gravitational potential function, which captures the nominal gravitational acceleration assuming a homogeneous central body plus one term in spherical harmonics (such as the J2 perturbation and other associated higher order terms).[25] A higher gravitational order may be used for high-fidelity simulations in order to capture a more accurate gravitational environment.

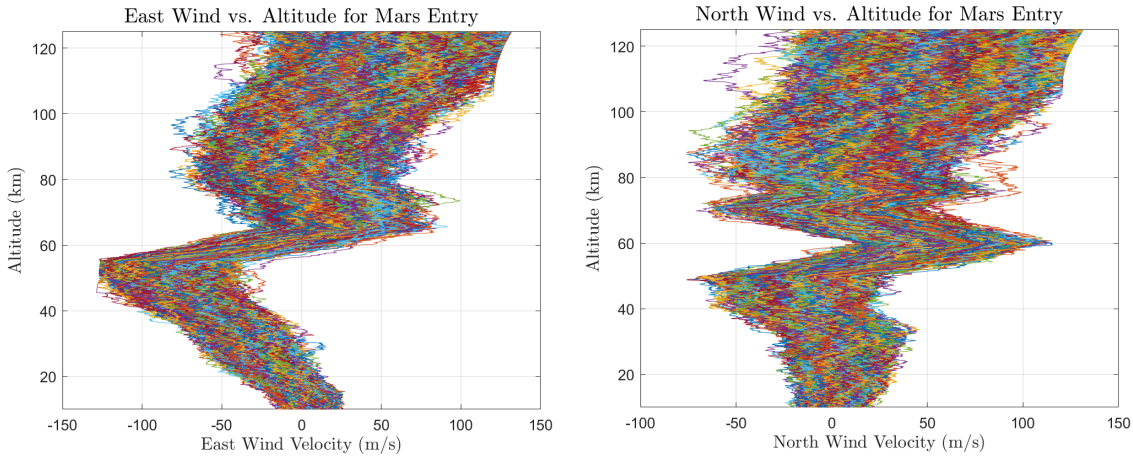

Figure 3: Relation of wind velocity and altitude between 10km and 125km for $N = 3000$ Monte Carlo runs via Mars-GRAM. Sub-figures show East and North wind velocity components.

Atmosphere Modeling and Dispersions

The Martian atmosphere is modeled via Mars-GRAM[20] to provide values of the atmospheric density ρ and the wind vector $\mathbf{w}$, which are known to be highly uncertain. To visualize this uncertainty, data from an $N = 3000$ Monte Carlo simulation are shown in Figure 2 and Figure 3. Mars-GRAM is sampled at altitudes ranging from 10 km to 125 km to provide a discrete density and wind lookup table with a spacing of 0.1 km. To recover a continuous data set for integration, a spline fit was used for density, East wind, and North wind. Table 1 shows the important nominal values and dispersions for the Mars-GRAM atmospheric simulations.

Parameter	Value	Parameter	Value
Initial Altitude	125 km	NR1	1234
Final Altitude	10 km	INTENS	0.0
Latitude, Longitude	0°, 0°	zoffset	0 km
Date	6 August 2012	wlscale	10.0
Time	17:30:00	rpscale	2.0

Table 1: MarsGRAM parameter values

Figure 2 demonstrates the importance of taking into account density variance from an empirical standpoint. It is clear that the density greatly varies at the lower altitudes, however, at higher altitudes, the *percent error* is the greatest. The wind variation is also large, however this acts as a minor disturbance as the maximum expected wind velocity has a magnitude of $||\mathbf{w}||_2 < 150$ m/s, while the nominal entry velocity magnitude is $||\mathbf{v}||_2 > 5500$ m/s. Thus, wind is not as important of a driver in the perturbations off of a nominal entry profile as density.

Monte Carlo Entry Simulations

In order to show the effect that uncertainty on the air density and wind velocity vectors has on the nominal entry trajectory, several ($N = 1000$) Monte Carlo entry simulations are performed using the dynamics in Eq. (1). These results are displayed in Figure 4. Each simulation includes the same

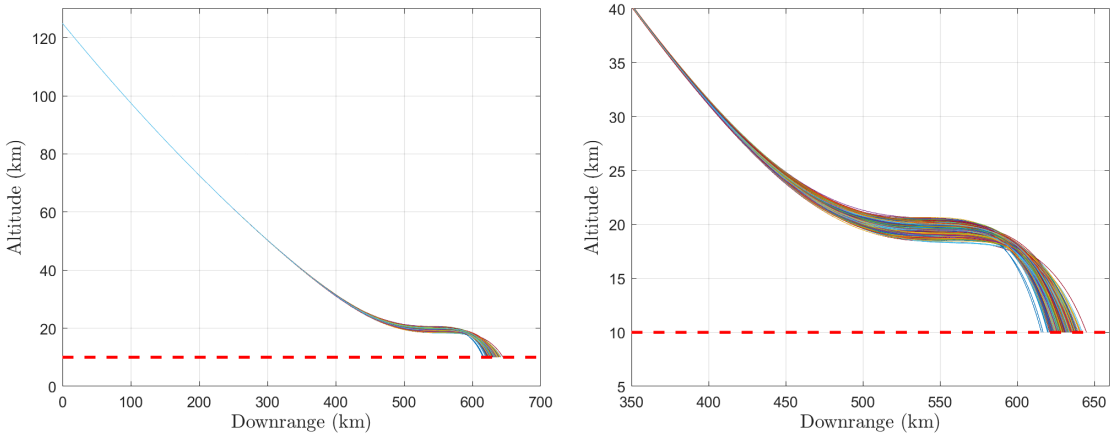

(a) Complete Trajectory

(b) Zoomed Trajectory

Figure 4: Entry trajectories from $N = 1000$ Monte Carlo simulations, where only atmospheric density and wind velocity vectors were dispersed according to Mars-GRAM.

exact initial conditions and entry vehicle parameters, only varying the atmospheric density and wind velocity vectors from their respective nominals along the entry profile.

Each simulation begins at 125 km in altitude and 0° Latitude and Longitude, facing due East along the equator. An initial, constant flight path angle was defined as $\gamma = -15.5^\circ$ with an initial entry velocity of 5.85 km/s. This is chosen from the Mars Science Laboratory (MSL) entry conditions.[21] A constant entry angle of attack of $\alpha = -9^\circ$ was also chosen to provide an entry profile that is similar to historic flight trajectories.[16] Normally the angle of attack is allowed (and sometimes required) to change, however this work only considers the ballistic trajectory under no control authority, and a constant angle of attack was deemed sufficient.

These Monte Carlo simulations display the importance of taking into account atmospheric deviations from the nominal. Given only dispersions on density and wind velocity vector per Mars-GRAM, the standard deviation of downrange final positions at an altitude of 10 km is 3.93 km off of a mean of 630.23 km for 1000 Monte Carlo runs of a ballistic trajectory simulation. This motivates the requirement for atmospheric adaptation in guidance and control algorithms to estimate the true atmospheric density in real time and account for it accordingly.

ATMOSPHERIC ADAPTATION

The proposed atmospheric adaptation algorithms aim to identify the time-varying parameter, k_ρ, which is the ratio between the observed and nominal atmospheric density. Specifically, this parameter is defined as

$$k_\rho = \frac{\rho_{\text{observed}}}{\rho_{\text{nominal}}},$$

where k_ρ is a dimensionless parameter that can be used to replace ρ in Equations (2) and (3) such that $\rho = k_\rho \rho_{\text{nominal}}$. In this case, ρ_{nominal} is the mean density at each altitude as defined by Mars-GRAM, which is commonly used as the nominal density profile for design of the initial guidance trajectory for a Mars entry mission.

The scaling parameter is used for atmospheric adaptation rather than directly identifying the true density, as the latter spans several orders of magnitude throughout the entry profile. Using k_ρ as a scaling on the nominal density allows for better numerical conditioning in the least-squares estimation approaches outlined in this section, as well as for future work that could include Kalman filtering techniques. Furthermore, much of the existing relevant atmospheric adaptation literature aims to identify this parameter rather than the true density for the same reasons as defined here.[17,18]

Two atmospheric adaptation algorithms are proposed. One is a 'static' estimate such that previous data is used to back out an estimate of the parameter k_ρ at the current time. This estimate could then be propagated forward and assumed constant for some time to provide a predictive model. The second algorithm makes use of basis functions that are a function of altitude, which could be propagated forward over some prediction horizon to provide a time-varying estimate. Ultimately, each algorithm is expected to support predictive control by giving an accurate estimate of the air density relative to the nominal.

Sliding Window Batch Estimation

This method for parameter estimation is rather intuitive—we use the known system dynamics to estimate k_ρ through algebraic manipulation of the equations of motion. This is especially useful if all other parameters are known to a high degree of accuracy, and if measurements of position, velocity, and acceleration are available. We return to the equations of motion in Equation (1), expanded to show the aerodynamic acceleration components:

$$\dot{\mathbf{v}} = \frac{1}{2m}\rho A C_L ||\mathbf{v}_{\text{rel}}||^2 \hat{\mathbf{e}}_2 - \frac{1}{2m}\rho A C_D ||\mathbf{v}_{\text{rel}}|| \mathbf{v}_{\text{rel}} + \mathbf{a}_g - 2(\boldsymbol{\omega} \times \mathbf{v}) - \boldsymbol{\omega} \times (\boldsymbol{\omega} \times \mathbf{r}),$$

where $\rho = k_\rho \rho_{\text{nominal}}$ is the actual atmospheric density. From here, we project the dynamic equations in the direction of the velocity vector and split the equation into parts that do and do not depend on k_ρ (i.e., aerodynamic accelerations multiplied by k_ρ, and all other accelerations). After some algebraic manipulation, the problem is written in the form $\mathbf{A}k_\rho = \mathbf{b}$, such that

$$\underbrace{\mathbf{v}_k^\mathsf{T} \left[C_L ||\mathbf{v}_{\text{rel},k}||^2 \hat{\mathbf{e}}_{2,k} - C_D ||\mathbf{v}_{\text{rel},k}|| \mathbf{v}_{\text{rel},k} \right] \frac{1}{2m} \rho_{\text{nom}} A}_{a_k} k_\rho$$
$$= \underbrace{\mathbf{v}_k^\mathsf{T} \left[\dot{\mathbf{v}}_k - \mathbf{a}_{g,k} - 2(\boldsymbol{\omega} \times \mathbf{v}_k) - \boldsymbol{\omega} \times (\boldsymbol{\omega} \times \mathbf{r}_k) \right]}_{b_k}. \quad (4)$$

This form is extremely useful for the simple parameter estimation problem, and allows for flexibility in choosing the number of data points (n) one would like to use. Here, n defines the 'bin size,' where a complete bin is defined as $\mathbf{A} = [a_1 \;\; a_2 \;\; \cdots \;\; a_n]^\mathsf{T}$ and similarly $\mathbf{b} = [b_1 \;\; b_2 \;\; \cdots \;\; b_n]^\mathsf{T}$. If $n = 1$, the solution is simply solved as $k_\rho = \frac{b}{A}$. If $n \geq 2$, the pseudo-inverse is used, such that

$$k_\rho = \mathbf{A}^+ \mathbf{b} = (\mathbf{A}^\mathsf{T}\mathbf{A})^{-1}\mathbf{A}^\mathsf{T}\mathbf{b},$$

where $\mathbf{A}$ and $\mathbf{b}$ contain consecutive data points as defined above. If there is no measurement noise and all parameters are known perfectly, this method will perfectly estimate the current k_ρ using a bin size $n = 1$. Should more time steps be used, k_ρ will be a sliding average estimate using several data points (hence 'Sliding Window'). Using too large of a bin size, however, will cause the estimate to lose some fidelity by being unable to capture smaller, faster variations in k_ρ.

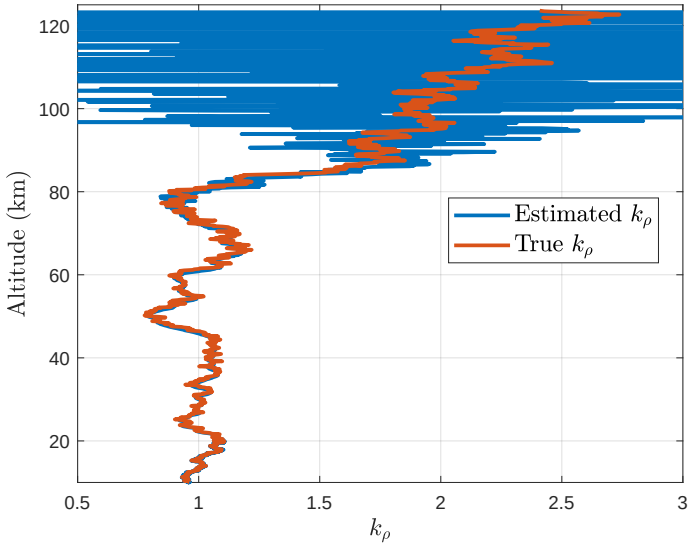

Figure 5: Sliding window batch estimation of density scaling parameter over the entire entry trajectory. The ratio between the estimated and nominal densities, k_ρ, is shown alongside the true values. Sliding window batch estimation is used with $n = 10$ time steps. Error due to simulated measurement noise greatly affects the estimation at higher altitudes where the density is very low.

If uncertainty is present in measurements or other vehicle parameters, a bin size of $n = 1$ would only capture the noisy measurements, and will be minimally useful. It is expected that as n increases, the closer the estimate of k_ρ will be to the true value, however with diminishing returns. This is same as the case with perfect measurements, where too large of a bin size will cause this parameter estimation technique to be unable to capture smaller changes in atmospheric density from the nominal. This indicates there is a trade-off with noise and losing fidelity/accuracy in estimating k_ρ, where the sensor qualities, guidance/control update frequency, and mission profile specifics would need to be considered in a selection of the bin size. Ultimately, using more data points will allow for a smoothing of the measurement noise.

In order to test the parameter identification algorithm outlined above, noise is added to the position, velocity, and acceleration vectors. Each state is appended with noise at each time step using a multivariate Gaussian distribution with a specified mean and covariance. The standard deviations include 100 m on position components, 10 m/s on velocity components, and 0.05 m/s^2 on the acceleration components. Further, all results herein use the same dispersed density profile from a single Monte Carlo run in Figure 2.

A bin size of $n = 10$ is used to demonstrate the sliding window batch estimation technique over the entire entry profile. Each timestep has length $\Delta t = 0.1$ seconds. Results are shown in Figure 5. It is clear that this estimation technique did not perform well at altitudes over roughly 80 km, although atmospheric adaptation for guidance and control algorithms usually does not begin until under 60 km, where the vehicle has greater control authority in the more dense portions of the atmosphere.[17]

A closer view on this parameter identification technique is shown in Figure 6, where three different bin sizes are used and compared between altitudes of 10 km and 60 km. It is seen that a bin size of $n = 1$ often gives a very noisy result, primarily driven by the measurement noise at the single

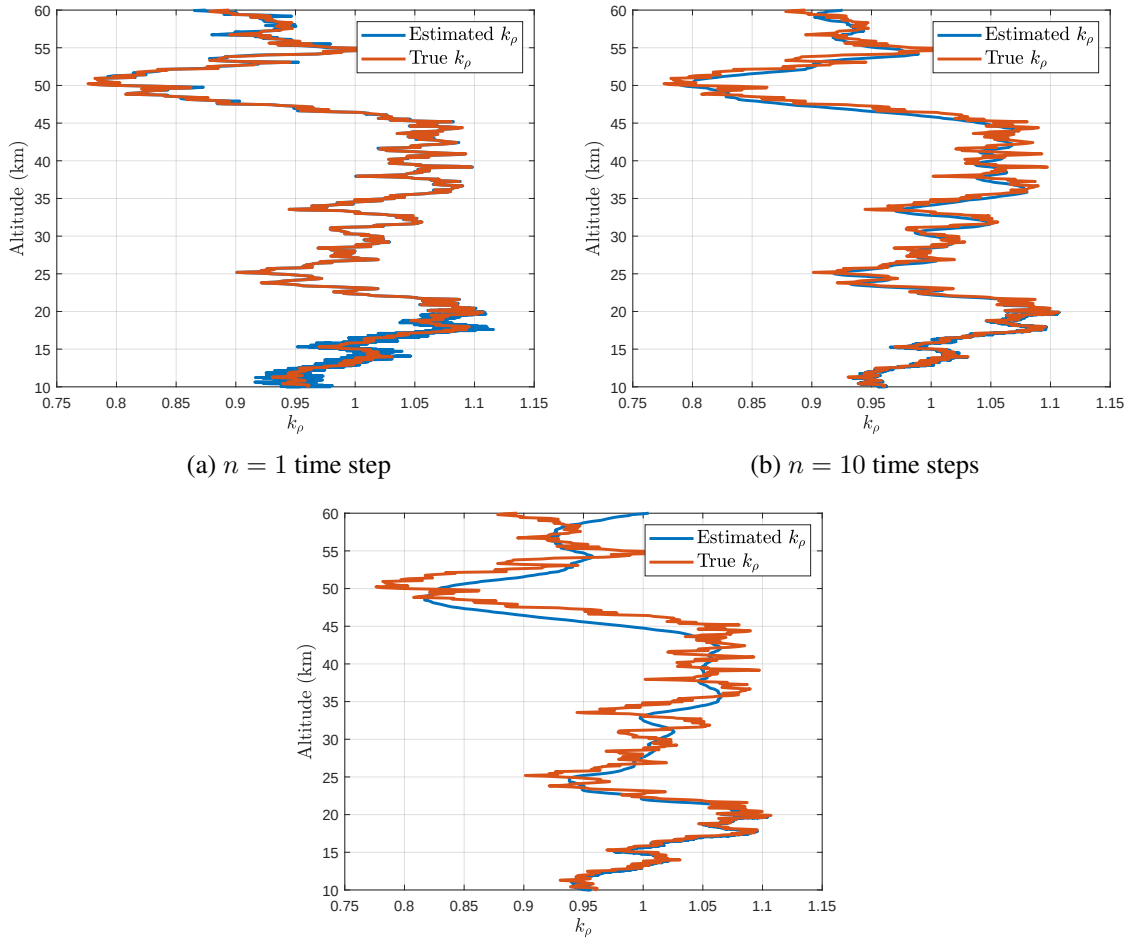

(a) $n = 1$ time step

(b) $n = 10$ time steps

(c) $n = 30$ time steps

Figure 6: Sliding window batch density estimation during entry where the ratio between the estimated and nominal densities, k_ρ, is shown alongside the true values with (a) 1, (b) 10, and (c) 30 time steps. Only the latter portion (< 60km) of the trajectory is shown.

time step. A bin size $n = 30$ is able to smooth out this noisy result, however it tends to lag in its estimation and cannot capture all of the fast changes in k_ρ. A bin size $n = 10$ seems to incorporate the best qualities for small and large bin sizes, where noise in the estimated k_ρ is minimized while also capturing the abrupt changes across altitude variations.

This technique is strongly capable of estimating accurate values of k_ρ at a specific point in time. This estimate could be propagated forwards as a constant value across some prediction horizon and still be beneficial at some points (e.g., at altitudes of 35–45 km), but otherwise the trend can widely vary. The sliding window batch estimate could severely limit capabilities to predict future trends given the simplicity of the static least-squares estimation approach.

Predictive Dual-Horizon Adaptation

A static estimate of the current scaling of true to nominal density is useful, however as one aims to predict future trends of k_ρ, more sophisticated methods could be used to fit available data.

Ultimately, a predictive modeling approach should be used through fitting previous data to certain basis functions, which are to be propagated forwards in time. This leads to the proposed method of predictive dual-horizon adaptation (PDHA). PDHA employs a simple and intuitive predictive modeling approach for atmospheric adaptation. It is capable of collecting data from n timesteps in the past, and propagate the expected k_ρ forwards m timesteps in the future (hence 'dual-horizon').

PDHA uses the same datasets as the sliding window batch estimation technique as defined in Equation (4). However, the purpose is now to solve for the 'best' coefficients in a least-squares sense for nonlinear basis functions that fit the trend of k_ρ in the given dataset. These basis functions then can be given a set of altitudes expected in the future to be used to effectively update the model available to a predictive control policy.

This method takes the form

$$\mathbf{A}_{\text{diag}}\boldsymbol{\Psi}(h)\boldsymbol{\Xi} = \mathbf{b},$$

where $\mathbf{A}_{\text{diag}} \in \mathbb{R}^{n\times n}$ is the data from the left-hand side of Equation (4), diagonalized, and $\mathbf{b} \in \mathbb{R}^n$ is defined by the right-hand side of Equation (4). The matrix $\boldsymbol{\Xi} = [\xi_1 \;\; \xi_2 \;\; \dots \;\; \xi_p]^\mathsf{T}$ contains the coefficients to be fit to the basis functions defined in $\boldsymbol{\Psi}(h) \in \mathbb{R}^{n\times p}$. Specifically, $\boldsymbol{\Psi}(h)$ contains the desired basis functions evaluated at the height corresponding to the data in $\mathbf{A}$ and $\mathbf{b}$, where $\boldsymbol{\Psi}^\mathsf{T} = \left[\hat{\psi}^\mathsf{T}(h_1) \;\; \hat{\psi}^\mathsf{T}(h_2) \;\; \dots \;\; \hat{\psi}^\mathsf{T}(h_n)\right]$. This can be used to solve for the coefficients of the basis functions in a least-squares sense such that

$$\boldsymbol{\Xi} = \left(\mathbf{A}_{\text{diag}}\boldsymbol{\Psi}(h)\right)^+ \mathbf{b}, \tag{5}$$

where again the pseudo-inverse is used to solve the non-square inverse operation. The solution of Equation (5) provides the coefficients for basis functions—thus, a predictive model of the expected air density deviation from the nominal is supplied through the solution of $k_\rho(h) = \hat{\psi}(h)\boldsymbol{\Xi}$.

This technique is applied to a few different snapshots of the same entry trajectory and density profile as the sliding window batch estimator. In these examples, PDHA is tasked to fit to the previous $n = 50$ timesteps and propagate forwards $m = 50$ timesteps into the future. Each case herein includes the same noise on the kinematic vectors as used in the sliding window batch estimation. First, a linear fit is used such that $\hat{\psi}_{\text{lin}} = [1 \;\; h_i]$ where i denotes the height at each timestep in the calculation. Figure 7 shows two examples of linear-fit snapshots. Figure 7a shows a case where the propagated linear fit matches well with the actual k_ρ, and figure 7b shows a case where the propagated linear fit diverges significantly after the measurement point. This result provides interesting insight into the simple linear fit case, where there may be certain altitude bands within the atmosphere where the trend in k_ρ is roughly linear, and other areas where such a linear fit is a poor estimate.

Next, a quadratic fit is used such that $\hat{\psi}_{\text{quad}} = \left[1 \;\; h_i \;\; h_i^2\right]$. Figure 8 shows two examples of quadratic-fit snapshots, where again Figure 8a shows a good match of k_ρ with the propagated quadratic fit and Figure 8b shows a case where the propagated quadratic fit diverges significantly over the prediction horizon. A cubic fit could be used, however this omitted in this work so as to avoid issues that arise with over-fitting. Regardless, the basis function matrix $\boldsymbol{\Psi}(h)$ would take the same form, i.e., a cubic fit would be defined as $\hat{\psi}_{\text{cubic}} = \left[1 \;\; h_i \;\; h_i^2 \;\; h_i^3\right]$.

The results from snapshots of the proposed PDHA approach to predictive modeling are promising. Ultimately, it is expected that an estimate of k_ρ will be most useful in atmospheric adaptation when it is close to the true value. However, when the prediction is far away from the true k_ρ, the

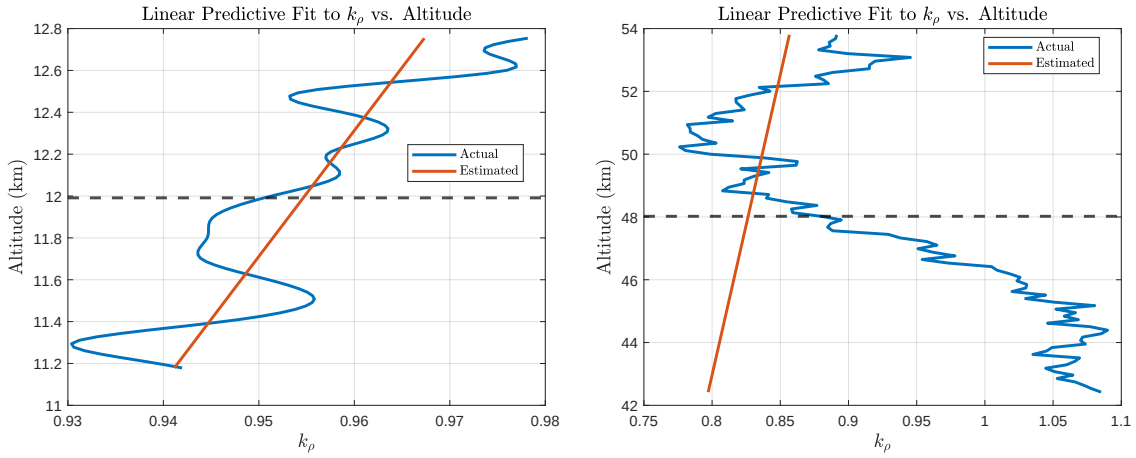

(a) Good linear fit (b) Poor linear fit

Figure 7: Two examples of a PDHA approach using linear basis functions, with 50 timesteps of data propagated 50 timesteps forwards. The horizontal dashed line defines the current point in time (i.e., above the line is data that is fitted and below is propagated forwards).

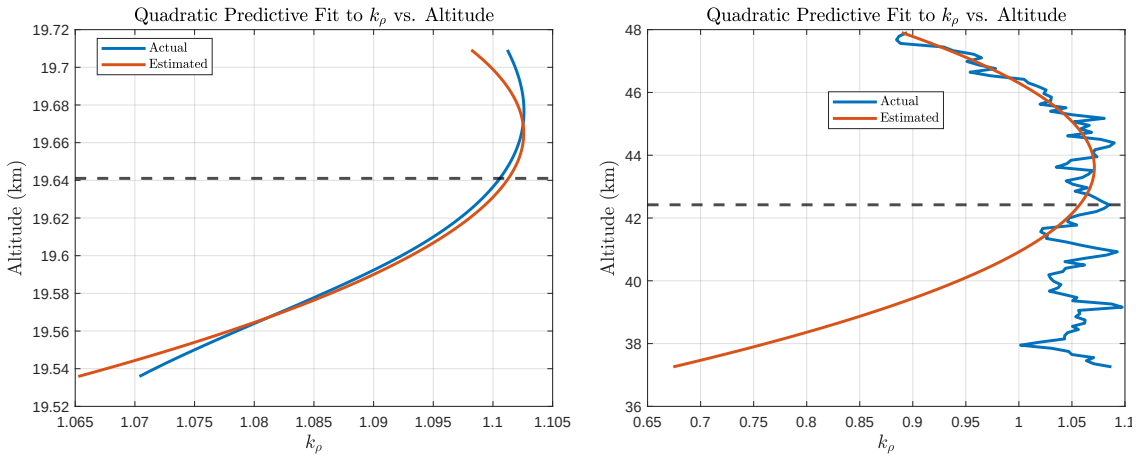

(a) Good quadratic fit (b) Poor quadratic fit

Figure 8: Two examples of a PDHA approach using quadratic basis functions, with 50 timesteps of data propagated 50 timesteps forwards. The horizontal dashed line defines the current point in time (i.e., above the line is data that is fitted and below is propagated forwards).

use of this prediction could actually hinder the performance of a predictive control approach. This effect is exacerbated in the worst-case when the predicted k_ρ is greater than 1, but the true k_ρ is less than 1.

PREDICTIVE MODELING ERROR ANALYSIS

This section takes a deeper dive into the prediction error uncovered when using sliding window batch estimation or PDHA with different basis functions throughout the latter portion of the entry trajectory. Figure 9 shows the error in k_ρ relative to the true value for a constant propagation, and linear and quadratic fits for PDHA over the portion of the entry trajectory after 60 km. The

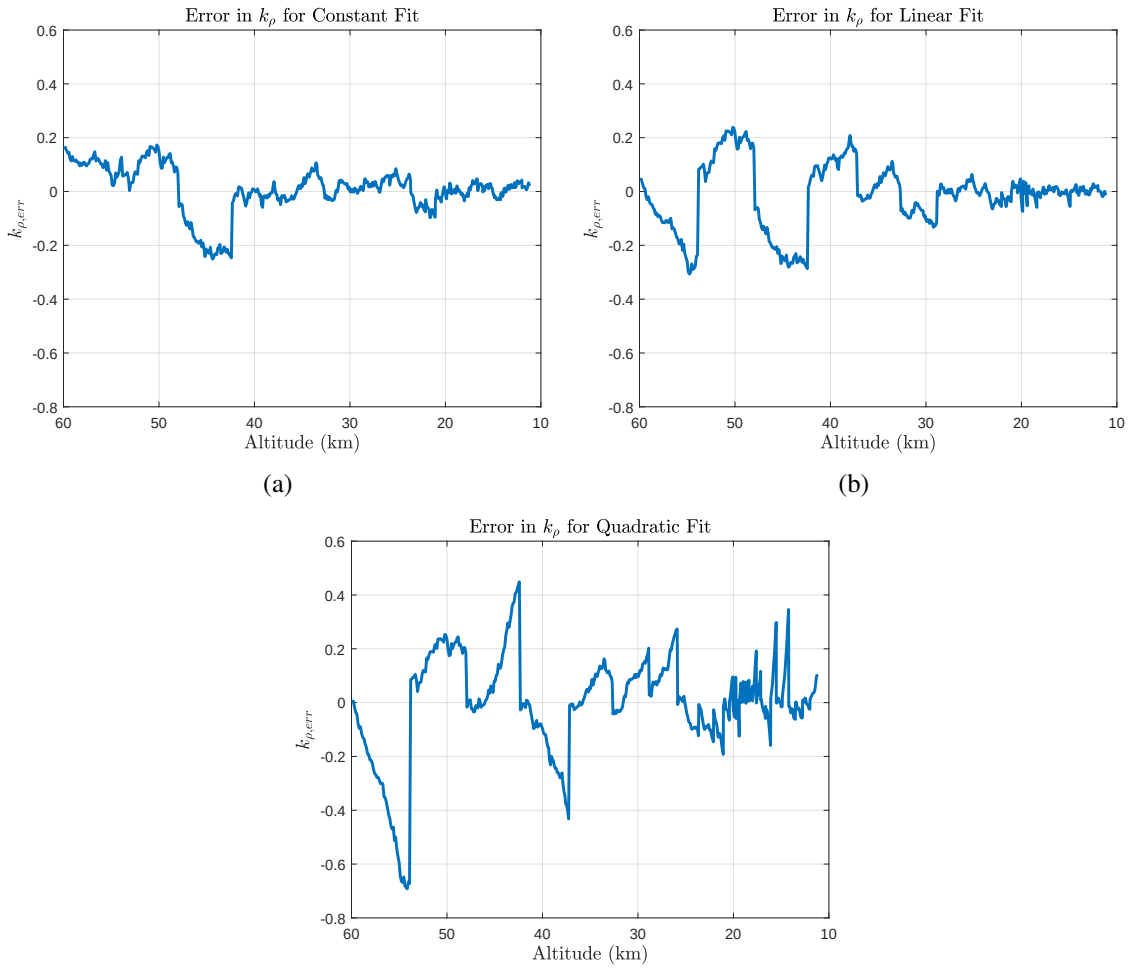

Figure 9: Error in k_ρ estimate relative to the true k_ρ for (a) constant propagation from the sliding window batch estimation, (b) a linear fit using PDHA, and (c) a quadratic fit using PDHA.

values plotted are defined as $k_{\rho,\text{err}} = k_{\rho,\text{estimated}} - k_{\rho,\text{actual}}$. A timestep of $\Delta t = 0.1$ seconds is used. Each fit uses 50 timesteps from the past to predict the value of k_ρ 50 timesteps in the future (i.e., $n = m = 50$, or 5 seconds). In practice, this computation would likely be repeated each timestep, however for the sake of clarity in the results presented, the propagated error is shown each 50 timesteps as repeated snapshots of the prediction.

Results from Figure 9 offer some insight into the quality of fit for each proposed method. However, a pure error presented in this way does not offer a full explanation of the predictive quality of this estimate. It is important to take into account the direction in which the k_ρ scaling takes the new density estimate. Should the estimated k_ρ scale the nominal density in the wrong direction (e.g., the estimate gives $k_\rho = 0.8$ while the true value is $k_\rho = 1.1$), the atmospheric adaptation would hinder the quality of the prediction compared to simply using the nominal air density.

A quantity to be used for investigating this sort of sign error is defined as $Q_\rho = k_\rho - 1$, such that the magnitude and direction of error due to each predictive modeling technique can be better understood. Results showing Q_ρ versus altitude from the constant fit propagation via the sliding

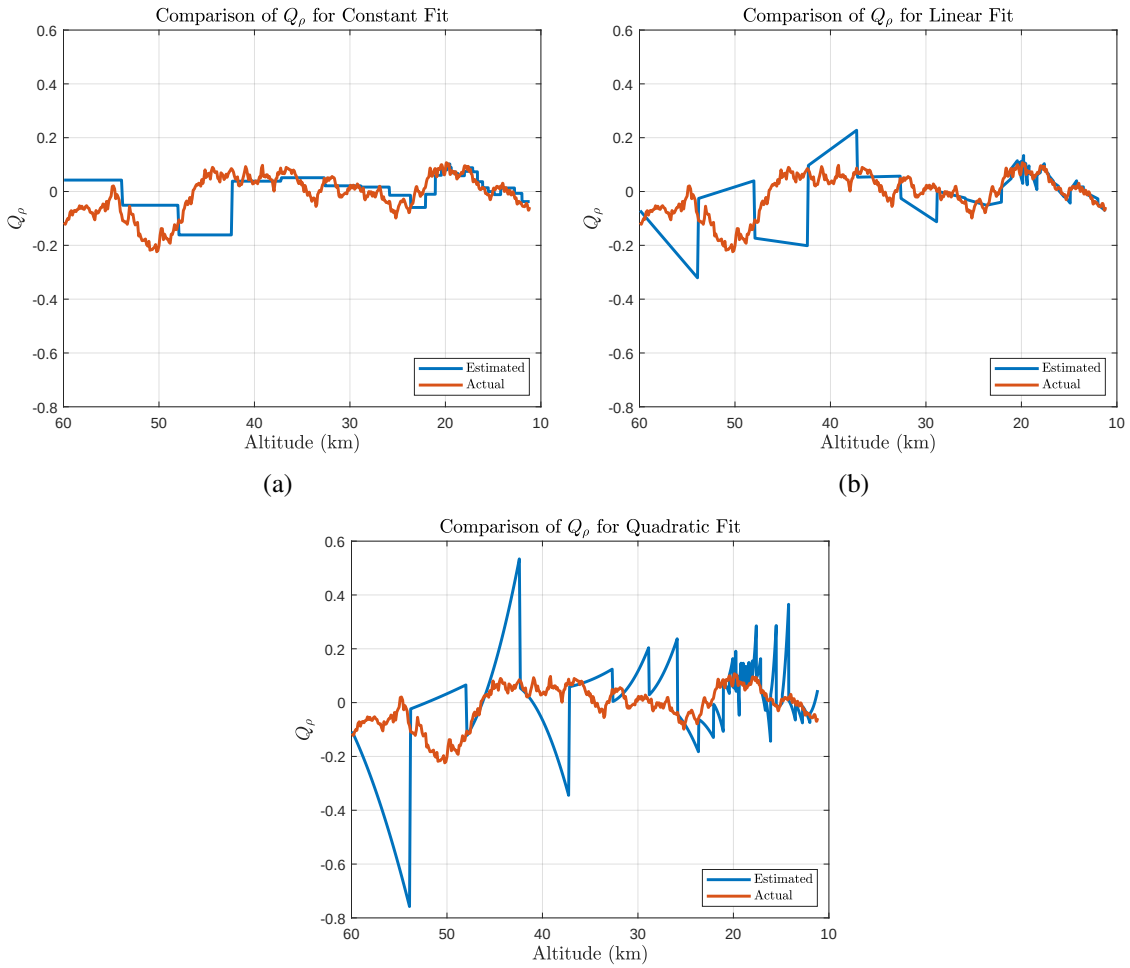

Figure 10: Q_ρ vs altitude of a 50-timestep fit for (a) constant propagation from the sliding window batch estimation, (b) a linear fit using PDHA, and (c) a quadratic fit using PDHA

window batch estimation and PDHA using linear propagation and quadratic propagation can be found in Figure 10. This figure uses the same propagation techniques as used to generate data in Figure 9. Finally, the results for Q_ρ using each of these techniques with a batch of 100 timesteps used to predict k_ρ 50 timesteps in the future (i.e., $n = 100$ and $m = 50$) are shown in Figure 11.

Discussion

The implementation of the sliding window batch estimator performed extremely well with regards to a static estimate through carefully-selected bin lengths, as seen in Figure 6. The primary limitation of using this technique on a time-varying parameter with noisy data can clearly be mitigated through the use of the proper number of time steps. When this static estimate is propagated forwards in the future, however, the fit may often suffer great errors. This error is attempted to be mitigated through the use of time-varying basis functions in the PDHA methodology. The linear fit and subsequent propagation seems to perform extremely well in the lower regime (< 30 km) of

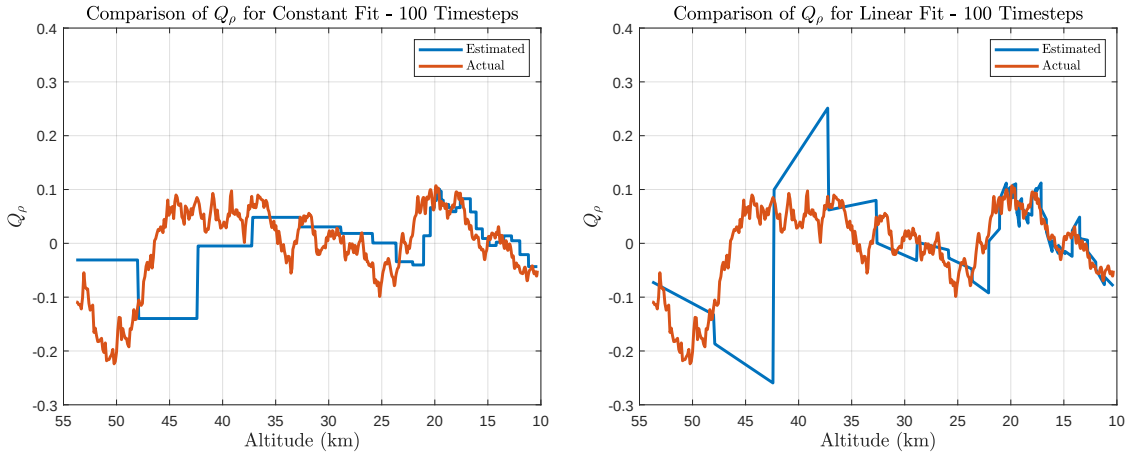

(a) Q_ρ for constant fit

(b) Q_ρ for linear fit

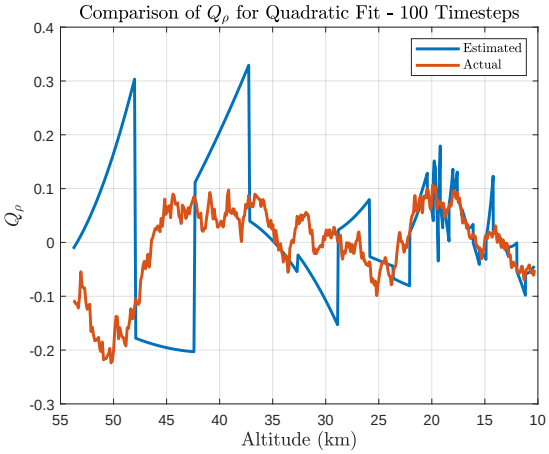

(c) Q_ρ for quadratic fit

Figure 11: Q_ρ vs altitude of a 100-timestep fit for (a) constant propagation from the sliding window batch estimation, (b) a linear fit using PDHA, and (c) a quadratic fit using PDHA

the atmosphere, however the fit in upper portions of the atmosphere tends to diverge quickly from the true values of k_ρ. Finally, the quadratic fit sees similar divergence across most altitudes. This is likely due to over-fitting of a limited amount of data and too-far of a forward propagation, however this effect is significantly mitigated with a bin length of 100 timesteps used to fit.

The results in Figures 10 and 11 seem to imply that atmospheric adaptation using the predictive modeling techniques proposed in this work only perform well for a small fraction of the entry trajectory (altitudes $<$ 30 km). While these approaches may only seem reasonable for the lower portions of the atmosphere, this only tells part of the story. It is important to take into account the time spent at different altitudes. As seen in Figure 1, the entry vehicle spends a significant portion of the downrange trajectory within 25-15 km. This is also seen clearly in Figures 10 and 11, where the updates to the propagated estimate are farther apart in upper portions of the entry trajectory, even though the timestep length remains constant.

A supplemental result showing Q_ρ versus time since entry interface is shown in Figure 12. For brevity, an example of only the linear propagated fit is shown, which is fit to $n = 50$ timesteps and

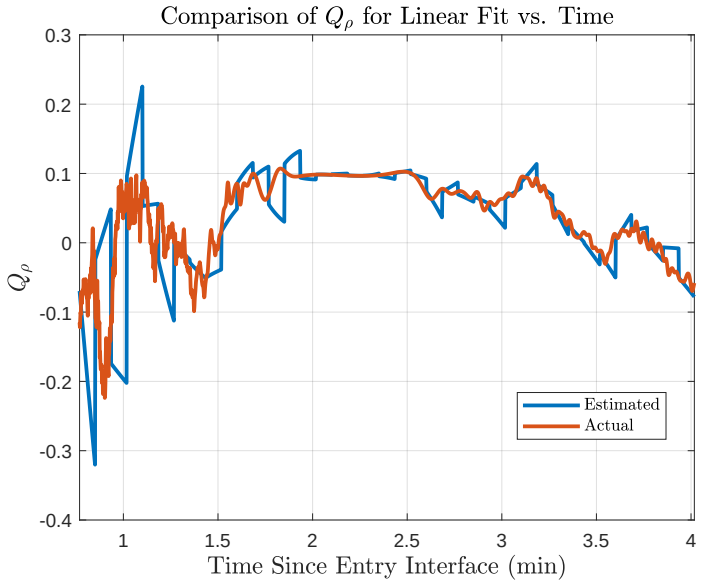

Figure 12: Q_ρ vs time since entry interface of a 50-timestep fit for a linear fit using PDHA

a forward propagation of the same $m = 50$ timesteps. In this figure, it is clear that the PDHA technique offers a great estimate of k_ρ over a large portion of the atmospheric entry. Specifically, the linear-fit predictive modeling gives accurate, time-varying estimates of k_ρ over the last $\approx$ 2 & 1/2 minutes of the entire 4-minute entry trajectory, in the portion of the atmosphere where air density is the greatest throughout the entry trajectory where bank-angle control is most effective.

These results support the use of sliding window batch estimation and time-varying predictive modeling via the PDHA methods for atmospheric adaptation in online guidance generation and predictive control algorithms. The predicted models could update the model used in a predictive control policy in real time according to the experienced environment. An understanding of current density compared to the nominal ρ (used to develop the initial guidance trajectory) will likely lead to improved performance in precision entry guidance control.

CONCLUSIONS

This work presents two intuitive methods for estimating the time-varying parameter that defines the ratio of observed atmospheric density to a nominal model. A Martian entry simulation was developed to gather perturbed density data and provide a time-history of entry trajectories. Monte Carlo results for different density vs. altitude and entry trajectory profiles were provided. A sliding window batch estimation technique was introduced that uses the entry equations of motion to identify a static estimate of k_ρ, and a method for predictive dual-horizon adaptation was introduced to provide a predictive estimate of air density. Results implied great potential for using forward propagation and predictive modeling for the use in a predictive control approach in entry guidance and control using atmospheric adaptation.

The techniques described in this work will be expanded through use of more sophisticated predictive modeling and system identification techniques. Further, the use of Kalman filtering will be implemented in conjunction with the predictive modeling methods introduced here, to more explicitly take into account measurement noise. These atmospheric adaptation methods will be applied to

entry guidance and control problems using model predictive control in order to update the environment model available to the control policy.

ACKNOWLEDGMENTS

Robert D. Halverson acknowledges partial support by the Science, Mathematics, and Research for Transformation (SMART) Scholarship-for-Service Program within the Department of Defense (DoD). Ryleigh McGiveron acknowledges support by the University of Minnesota Undergraduate Research Opportunities Program (UROP).

REFERENCES

[1] R. D. Braun and R. M. Manning, "Mars exploration entry, descent and landing challenges," *IEEE Aerospace Conference*, 2006, pp. 18–pp.

[2] Z. Yu, P. Cui, and J. L. Crassidis, "Design and optimization of navigation and guidance techniques for Mars pinpoint landing: Review and prospect," *Progress in Aerospace Sciences*, Vol. 94, 2017, pp. 82–94.

[3] B. M. Jakosky, R. P. Lin, J. M. Grebowsky, J. G. Luhmann, D. Mitchell, G. Beutelschies, T. Priser, M. Acuna, L. Andersson, D. Baird, *et al.*, "The Mars atmosphere and volatile evolution (MAVEN) mission," *Space Science Reviews*, Vol. 195, 2015, pp. 3–48.

[4] A. D. Hayes, I. Nompelis, R. Caverly, J. Mueller, and D. Gebre-Egziabher, "Dynamic stability analysis of a hypersonic entry vehicle with a non-linear aerodynamic model," *AIAA Aviation Forum*, 2020, p. 3201.

[5] S. Omar and R. Bevilacqua, "Guidance, navigation, and control solutions for spacecraft re-entry point targeting using aerodynamic drag," *Acta Astronautica*, Vol. 155, 2019, pp. 389–405.

[6] J. Davis, A. Dwyer Cianciolo, R. Powell, J. Shidner, and E. Garcia-Llama, "Guidance and control algorithms for the Mars entry, descent and landing systems analysis," *AIAA/AAS Astrodynamics Specialist Conference*, 2010, p. 7972.

[7] K. Tracy and Z. Manchester, "CPEG: A convex predictor-corrector entry guidance algorithm," *IEEE Aerospace Conference*, 2022.

[8] 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.

[9] W. Van Soest, Q. Chu, and J. Mulder, "Combined feedback linearization and constrained model predictive control for entry flight," *Journal of Guidance, Control, and Dynamics*, Vol. 29, No. 2, 2006, pp. 427–434.

[10] J. Recasens, Q. Chu, and J. Mulder, "Robust model predictive control of a feedback linearized system for a lifting-body re-entry vehicle," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, 2005, p. 6147.

[11] L. Yang, X. Liu, W. Chen, and H. Zhou, "Autonomous entry guidance using linear pseudospectral model predictive control," *Aerospace Science and Technology*, Vol. 80, 2018, pp. 38–55.

[12] B. A. Steinfeldt, M. J. Grant, D. A. Matz, R. D. Braun, and G. H. Barton, "Guidance, navigation, and control system performance trades for Mars pinpoint landing," *Journal of Spacecraft and Rockets*, Vol. 47, No. 1, 2010, pp. 188–198.

[13] S. Li and X. Jiang, "Review and prospect of guidance and control for Mars atmospheric entry," *Progress in Aerospace Sciences*, Vol. 69, 2014, pp. 40–57.

[14] R. Derollez and Z. Manchester, "Sample-based robust uncertainty propagation for entry vehicles," *AAS/AIAA Astrodynamics Specialist Conference*, 2020.

[15] J. Williams, D. Woffinden, and Z. R. Putnam, "Mars entry guidance and navigation analysis using linear covariance techniques for the Safe and Precise Landing–Integrated Capabilities Evolution (SPLICE) project," *AIAA Scitech Forum*, 2020, p. 0597.

[16] G. L. Jacob, G. Neeler, and R. Ramanan, "Mars entry mission bank profile optimization," *Journal of Guidance, Control, and Dynamics*, Vol. 37, No. 4, 2014, pp. 1305–1316.

[17] K. S. Tracy, G. Falcone, and Z. Manchester, "Robust entry guidance with atmospheric adaptation," *AIAA SciTech Forum*, 2023, p. 0301.

[18] A. Hayes and R. J. Caverly, "Atmospheric density estimation in low-Earth orbit for drag-modulated spacecraft," *AAS Guidance, Navigation and Control Conference*, 2022.

[19] A. D. Hayes and R. J. Caverly, "Model predictive tracking of spacecraft deorbit trajectories using drag modulation," *Acta Astronautica*, Vol. 202, 2023, pp. 670–685.

[20] H. Justh, "Mars Global Reference Atmospheric Model 2010 Version: Users Guide," tech. rep., 2014.

[21] A. Steltzner, D. Kipp, A. Chen, D. Burkhart, C. Guernsey, G. Mendeck, R. Mitcheltree, R. Powell, T. Rivellini, M. San Martin, *et al.*, "Mars Science Laboratory entry, descent, and landing system," *IEEE Aerospace Conference*, 2006, pp. 15–pp.

[22] L. Blackmore, B. Açıkmeşe, and J. M. Carson III, "Lossless convexification of control constraints for a class of nonlinear optimal control problems," *Systems & Control Letters*, Vol. 61, No. 8, 2012, pp. 863–870.

[23] 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.

[24] F. G. Lemoine, D. E. Smith, D. D. Rowlands, M. Zuber, G. Neumann, D. Chinn, and D. Pavlis, "An improved solution of the gravity field of Mars (GMM-2B) from Mars Global Surveyor," *Journal of Geophysical Research: Planets*, Vol. 106, No. E10, 2001, pp. 23359–23376.

[25] D. A. Vallado, *Fundamentals of Astrodynamics and Applications*. El Segundo, CA: Microcosm Press, 2nd ed., 2004.