

Onboard Density Modeling for Planetary Entry via Karhunen-Loève Expansion

Samuel W. Albert
Aerospace Engineering Sciences
University of Colorado Boulder
samuel.albert@colorado.edu

Alireza Doostan
Aerospace Engineering Sciences
University of Colorado Boulder

Hanspeter Schaub
Aerospace Engineering Sciences
University of Colorado Boulder

Abstract— Onboard density models are a key aspect of autonomous closed-loop guidance systems for hypersonic flight. Traditional approaches model density as a deterministic function of altitude, but a recent drive toward stochastic guidance approaches motivates onboard uncertainty propagation. Existing solutions for efficient uncertainty propagation generally treat density as an exponential function of altitude, but this approach is limited in its ability to capture relevant dispersions. This work models density as a Gaussian random field that is approximated by a Karhunen-Loève expansion, enabling a relatively high-fidelity, finite-dimensional parametric representation. Various normalization schemes are presented and compared by their efficiency in capturing density variability in a limited number of terms, and normalization by reference dynamic pressure is shown to be the most compact approach. The model alternatives are compared both by their approximations of density itself and their predictions of peak heat flux for dispersed direct-entry and aerocapture trajectories. An extension of this approach for modeling density as a function of multiple independent variables is also presented and demonstrated. Finally, it is shown that the density model can be sequentially updated according to noisy density observations by formulating the problem as a Kalman measurement function.

to expected density [2], [3]. Recent work contributes more sophisticated methods of incorporating in-flight observations, such as machine learning or an ensemble correlation filter [4], [5], [6]. However, these methods ultimately treat the density as known and update a nominal profile.

Recent and ongoing works propose stochastic approaches to closed-loop guidance with the aim of being robust to uncertainties without taking an overly-conservative approach [7], [8], and central to these methods is an onboard prediction of state and environmental uncertainty. Several non-Monte Carlo uncertainty quantification (UQ) techniques, including polynomial chaos expansion and linear covariance analysis [9], [10], [11], potentially enable onboard uncertainty propagation for hypersonic flight vehicles. However, these methods generally require a parametric, low-dimensional representation of uncertainty [12], [13]. Recent studies have explicitly incorporated a probabilistic atmosphere model into UQ approaches [14], [15], [16]; however, these approaches typically assume an exponential form for density and incorporate uncertainty by dispersing the atmospheric scale height and surface density, a method that always results in an exponential profile. The assumption of exponential density significantly limits the ability of the model to capture more complex perturbations due to its inability to capture short-period perturbations or other deviations of the density profile from the idealized exponential shape [17]. Semi-empirical models such as the Global Reference Atmospheric Models (GRAMs) from NASA provide much higher-fidelity representations of the atmosphere and its response to external factors, such as solar weather [18], but lack a convenient low-dimensional and parametric form. Estimating uncertainty using these models typically requires generating a large number of density profiles then computing statistics of the generated dataset, rather than estimating variability directly. Thus, GRAMs and similar models are not feasible for onboard uncertainty propagation purposes.

This motivates the development of a reduced-dimensionality model that retains the higher-fidelity properties of models like GRAM, and a method for in-flight updates to this model. Previous work by the authors treats density as a Gaussian random field with altitude the sole independent variable, and demonstrates a Karhunen-Loève expansion (KLE) for density [17]. Reference [10] shows that linear covariance analysis incorporating this model closely matches Monte Carlo results. This study expands on these results in the following ways. Practical implementation of the KLE is explored in greater detail, examining alternative methods of constructing the expansion. The efficiency of each approach in capturing density variability is compared both directly and through statistics of dispersed trajectories generated in Monte Carlo analyses using each model. New work outlining and demonstrating an expansion on the KLE model to treat density as a function of multiple variables (e.g. altitude, latitude, and longitude) is presented and its comparative utility is discussed. Finally,

TABLE OF CONTENTS

1. INTRODUCTION.....	1
2. REVIEW OF KARHUNEN-LOÈVE EXPANSION.....	2
3. SIMULATION METHODOLOGY	2
4. COLUMNAR ATMOSPHERE MODEL COMPARISON	3
5. MULTI-DIMENSIONAL KLE MODEL	6
6. KALMAN MEASUREMENT UPDATES	9
7. CONCLUSIONS.....	11
ACKNOWLEDGMENTS	11
REFERENCES	11
BIOGRAPHY	12

1. INTRODUCTION

Hypersonic flight mechanics are characterized by nonlinear dynamics and high sensitivity to variations in atmospheric density. Furthermore, the behavior of planetary atmospheres is complex and difficult to predict. Appropriate modeling of density is thus key to the analysis of hypersonic trajectories, including in the context of onboard modeling for closed-loop guidance schemes. Autonomous guidance algorithms typically treat density as a known function of altitude, either in analytical form as an exponential function of altitude or by interpolating from a table [1]. In-flight estimates of the current density are available by measuring sensed acceleration, and these observations are then incorporated by multiplying the nominal profile by the ratio of observed density

an approach to updating the KLE based on sequential noisy density measurements is presented and demonstrated, and the potential for onboard execution of this method is discussed.

2. REVIEW OF KARHUNEN–LOÈVE EXPANSION

A random field is a function that maps a random outcome to a continuous function across a (possibly multi-dimensional) domain in space. Somewhat more formally: for some measurable space (Ω, \mathcal{F}) of sample space Ω and σ -field \mathcal{F} of subsets of Ω , a random field $\{\Phi(z) : z \in \mathcal{Z} \subseteq \mathbb{R}\}$ is a collection of random variables $(X_z)_{z \in \mathcal{Z}}$ with values that map $\Omega \mapsto \mathbb{R}$ [19]. A Gaussian random field (GRF) $\Psi(z)$ is a random field for which any finite linear combination of the random variables X_z results in a Gaussian random variable; that is, at any point z_i in the domain \mathcal{Z} the probability density function of the value of the field $\Psi(z_i)$ is Gaussian [20]. A GRF is fully characterized by its mean function μ and covariance function Σ ,

$$\mu(z) = \langle \Psi(z) \rangle, \quad (1)$$

$$\Sigma(z_1, z_2) = \langle (\Psi(z_1) - \langle \Psi(z_1) \rangle)(\Psi(z_2) - \langle \Psi(z_2) \rangle) \rangle, \quad (2)$$

where $\langle \cdot \rangle$ is the expectation operator.

A Karhunen–Loève (also known as Kosambi–Karhunen–Loève) expansion represents a random field through an infinite linear combination of orthogonal basis functions (a Fourier expansion), in such a way that, when truncated to a fixed number of terms, the choice of the basis functions minimizes the mean-square error [21], [22]. This definition is shown by Eq. (3) where Φ is the random field, z is the independent variable, and λ_i and $\phi_i(z)$ are the eigenvalues and eigenfunctions of the covariance function of the random field $\Sigma(z_1, z_2)$, respectively, as shown in Eq. (4). Finally, each Y_i is a random variable described by Eq. (5).

$$\Phi(z) = \langle \Phi(z) \rangle + \sum_{i=1}^{\infty} \sqrt{\lambda_i} \phi_i(z) Y_i \quad (3)$$

$$\int_0^T \Sigma(z_1, z_2) \phi_i(z_2) dz_2 = \lambda_i \phi_i(z_1) \quad (4)$$

$$Y_i = \frac{1}{\sqrt{\lambda_i}} \int_0^T \Phi(z) \phi_i(z) dz \quad (5)$$

In practice, the eigenvalues and eigenfunctions are sorted by descending magnitude of the eigenvalues and then the sum in Eq. (3) is truncated after some d_K number of sufficient terms. Determining the required d_K is problem-dependent, but in general it is chosen such that the mean-square norm of the approximation is within some relative error of the exact mean-square norm. Eq. (6) gives one heuristic method, where k is some sufficiently large number and α is close to 1 based on the desired level of permissible error (for a relative mean-square norm error of $(1 - \alpha) \times 100\%$).

$$d_K = \min \left\{ j : \frac{\sum_{i=1}^j \lambda_i}{\sum_{i=1}^{j+k} \lambda_i} \geq \alpha \right\} \quad (6)$$

In the case where $\Phi(z)$ is a GRF $\Psi(z)$, the Y_i 's are all independent and identically distributed (i.i.d.) standard normal random variables:

$$Y_1, Y_2, \dots \sim \mathcal{N}(0, 1) \text{ i.i.d.} \quad (7)$$

Often a random field is not known exactly, but some sufficiently large dataset is available. In this case the sample covariance matrix is computed,

$$C_{ZZ} \approx Q_{ZZ} = \frac{1}{N-1} \Psi_c \Psi_c^T, \quad (8)$$

where Q_{ZZ} is the unbiased estimate of the sample covariance matrix C_{ZZ} , $\Psi_c \in \mathbb{R}^{n \times m}$ is a matrix such that each column is an observation vector less the sample mean, n is the number of datapoints per observation vector, and m is the number of observation vectors in the dataset. Having computed a covariance matrix, it is straightforward to find the eigenvalues and eigenvectors of that matrix and sort them according to descending order of the eigenvalues, and the results are the λ_i and ϕ in Eq. (3), respectively. The discrete KLE form of a GRF Ψ is thus summarized below:

$$\Psi \approx \langle \Psi \rangle + \sum_{i=1}^{d_K} \sqrt{\lambda_i} \phi Y_i, \quad (9)$$

$$Y_i, \dots, Y_d \sim \mathcal{N}(0, 1) \text{ i.i.d.}$$

3. SIMULATION METHODOLOGY

This section briefly describes the methodology for trajectory simulation used in this study and summarizes relevant vehicle parameters. Trajectories are simulated by numerically propagating the three degree-of-freedom equations of motion for atmospheric flight about a rotating ellipsoidal planet via explicit Runge-Kutta integration of order 4(5). Density is modeled using MarsGRAM 2010 [18], interpolating from a resulting table of density vs. altitude unless stated otherwise. Mars is assumed to have gravitational parameter $\mu = 4.305 \times 10^4 \text{ km}^3 \text{ s}^{-2}$, equatorial radius $R = 3397.2 \text{ km}$, oblateness spherical harmonic coefficient $J_2 = 0.001964$, and a planetary rotation period of $\omega_p = 1.02595675 \text{ days}$ [23]. Mach number is defined as the ratio of vehicle speed to the speed of sound $M = v/a$, where sound speed a for the Martian atmosphere is interpolated from a nominal tabular model [24]. Heat flux is modeled by computing convective heat flux \dot{q} at the stagnation point assuming a fully catalytic surface using the Sutton-Graves expression shown in Eq. (10), where ρ is density and a value of the heating coefficient $k = 1.904 \times 10^{-4} \text{ kg}^{0.5}/\text{m}$ is used based on nominal atmospheric composition at Mars [25]. Dynamic pressure q is defined by Eq. (11).

$$\dot{q} = k \sqrt{\frac{\rho}{R_n}} V^3 \quad (10)$$

$$q = \frac{1}{2} \rho v^2 \quad (11)$$

There are two types of trajectories used as representative examples in this study. The first is a steep direct entry trajectory at Mars for the Small High Impact Energy Landing Device, or SHIELD, a small, mostly-passive probe under development at NASA JPL intended for low-cost access to the Martian surface [26]. Once reaching subsonic conditions, SHIELD deploys a drag skirt, then jettisons the heatshield shortly thereafter. The drag coefficient C_D during each configuration varies with Mach number and is linearly interpolated from tabular data provided by the JPL SHIELD team. Ballistic coefficient $\beta = m/(C_D A)$ describes the ratio of inertial forces to aerodynamic forces, where m is vehicle mass and A is reference area; the ballistic coefficient for SHIELD ranges

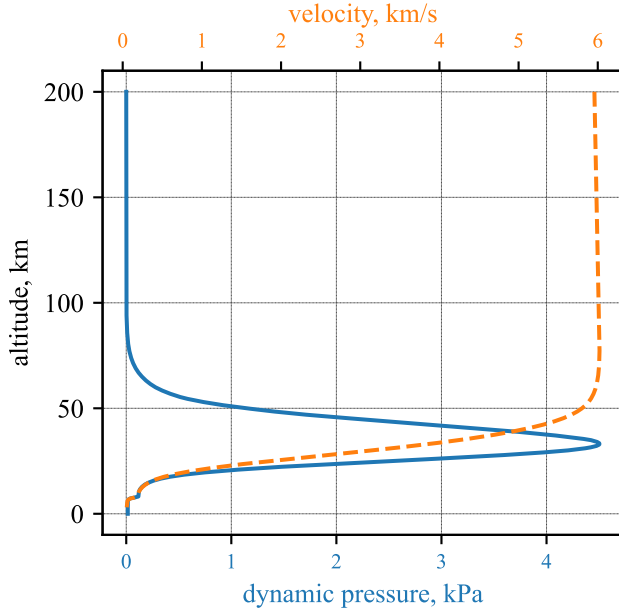


Figure 1. Dynamic pressure and velocity magnitude vs. altitude for reference SHIELD direct-entry trajectory

from about 20 kg m^{-2} shortly after entry to around 5 kg m^{-2} near the surface after drag skirt deployment and heatshield jettison. SHIELD has a lift-to-drag ratio of $L/D = 0$, and an assumed nose radius of $R_n = 0.85 \text{ m}$. The trajectory considered in this study is defined by an entry velocity of 6 km/s and an entry flight-path angle (EFPA) of -18° at the atmospheric interface altitude of 125 km , entering due-East at 0° latitude and 0° longitude, where flight-path angle is defined as the angle between the air-relative velocity of the vehicle and the local horizontal. The reference SHIELD direct-entry trajectory is shown in Fig. 1.

The other trajectory considered here is aerocapture at Mars by a vehicle similar to the Mars Science Laboratory (MSL) aeroshell. A ballistic coefficient of $\beta = 130 \text{ kg m}^{-2}$ and lift-to-drag ratio of $L/D = 0.24$ are assumed [27], and the vehicle flies full-lift-up for the duration of the trajectory. The entry is again due-East at 0° latitude and longitude, in this case with entry velocity of 5.8 km/s and EFPA of -11° . A nose radius of $R_n = 1 \text{ m}$ is assumed, which conveniently normalizes the value of \dot{q} for re-scaling to other vehicles. The reference aerocapture trajectory is shown in Fig. 2.

4. COLUMNAR ATMOSPHERE MODEL COMPARISON

In this section, atmospheric density is approximated as a Gaussian random field as a function of only altitude. In reality, atmospheres vary across 3D position and time, and are affected by external factors such as space weather. However, for applications like entry and aerocapture which traverse tens of vertical kilometers within the atmosphere, the dominant factor in density change is altitude. Thus, a columnar atmosphere model is assumed in this section, such that $\rho(h, \phi, \theta, t) \approx \rho(h)$ where ρ is density, h altitude, ϕ latitude, θ longitude, and t elapsed time. See Section 5 for a discussion of density variation with latitude and longitude.

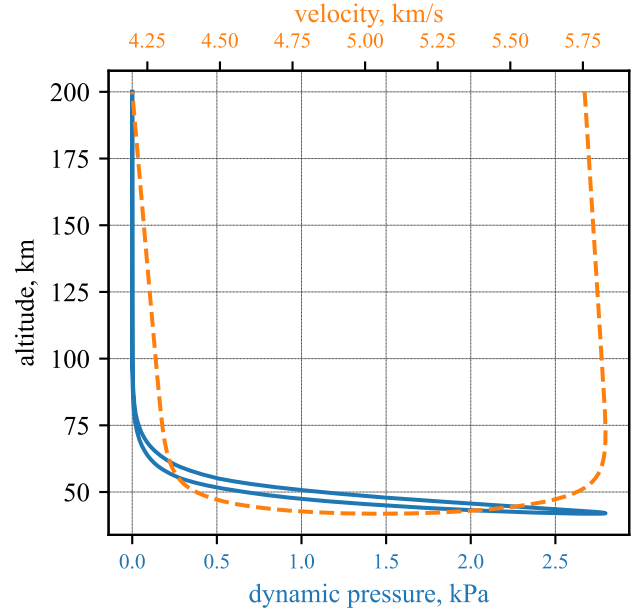


Figure 2. Dynamic pressure and velocity magnitude vs. altitude for reference aerocapture trajectory

The additional assumption of Gaussian probability is justified in Ref. [17] using Mars-GRAM 2010 data. It should also be noted that, to avoid a nonzero probability of producing a negative value, the density random field should in fact be treated as a truncated Gaussian.

Because density exhibits approximately Gaussian probability with correlation structure across a spatial domain, modeling density as a GRF is theoretically appropriate [17], [28]. However, a GRF is an infinite-dimensional object, and the non-Monte Carlo methods for onboard uncertainty propagation discussed earlier required a parametric, finite-dimensional representation of density variability [12], [13]. To obtain this form, a KLE is constructed and truncated after an appropriate number of terms. The sample covariance matrix is formed from any sufficiently large dataset of density values vs. altitude; typically, it is convenient to use simulated data from a relevant model such as a GRAM.

Figure 3 shows the result of constructing a KLE from a dataset of 5000 density profiles output by MarsGRAM, denoted KLE- ρ for shorthand. For the sake of later comparison, a fixed number of $d_K = 15$ terms is used for this and all subsequent expansions in this section unless noted otherwise. The horizontal axis of this plot shows normalized density perturbation $\delta\rho$, as defined in Eq. (12), rather than density itself because this better captures variability even as the value of density changes by orders of magnitude across this altitude range:

$$\delta\rho = \rho/\bar{\rho} - 1. \quad (12)$$

The thick dashed lines show the $\pm 3\sigma$ bounds, where σ is standard deviation. In the case of MarsGRAM these bounds are computed directly from the sample profiles; for the KLE, 5000 separate realizations are generated and evaluated, then standard deviation is computed from this generated dataset. In addition, three sample profiles from each model are shown in the thin solid lines.




Figure 3. KLE vs. Mars-GRAM for KLE constructed from density values; thick dashed lines are $\pm 3\sigma$ bounds, thin solid lines are sample profiles

Notably, the KLE 3σ bounds only align with the MarsGRAM bounds up to about 35 km, badly underestimating variability at higher altitudes. This occurs because the value of density is much greater at low altitudes: for Mars, about $1 \times 10^{-2} \text{ kg/m}^3$ at the surface, order of $1 \times 10^{-5} \text{ kg/m}^3$ at 50 km, and order of $1 \times 10^{-10} - 1 \times 10^{-9} \text{ kg/m}^3$ at the atmospheric interface of 125 km. The KLE is truncated based on eigenvalue magnitude, and the variability at low altitudes where density is high is prioritized as a result, even though as a percentage of nominal density varies more at high altitudes. For this reason, a KLE based on density values is an inefficient way to capture normalized density perturbations at high altitudes.

This shortcoming can be addressed by constructing the KLE differently. While columnar density remains the quantity of interest, the data can be pre-processed for KLE construction in a variety of ways, with a converse post-processing step recovering density values. For example, the KLE can be constructed from normalized density perturbation values in the following way. First, compute $\delta\rho$ values corresponding to each value in the dataset, forming a mean vector and covariance matrix for these $\delta\rho$ data. Second, construct a KLE using the $\delta\rho$ mean and covariance. Third, treat the output of this KLE as $\delta\rho$ values and re-arrange Eq. (12) to recover density values. The result of a KLE constructed in this way is shown in Fig. 4, denoted KLE- $\delta\rho$.

Figure 4 shows a clear improvement for the KLE- $\delta\rho$ model in terms of capturing overall density variability, and the sample profiles now look similar to the GRAM output. However, the 3σ bounds generally underestimate variability, and this becomes much more pronounced at low altitudes below about 55 km. This expansion in some ways suffers from the opposite problem as the KLE- ρ model: because normalized density perturbations are smaller near the surface, this region is poorly captured, whereas the model performs well at high




Figure 4. KLE vs. Mars-GRAM for KLE constructed from normalized density perturbations; thick dashed lines are $\pm 3\sigma$ bounds, thin solid lines are sample profiles

altitudes. That said, the KLE- $\delta\rho$ model is more compact, meaning that for a given number of terms it gives a better approximation of density variability with altitude.

However, it is important to keep the application of interest in mind. The goal of these KLE approximations is not to model the atmosphere as well as possible; the real goal is to provide an atmosphere model that results in accurate trajectory predictions when compared to trajectories predicted using MarsGRAM directly. Recall that aerodynamic force scales with dynamic pressure q . As seen in Fig. 1, for a planetary entry trajectory dynamic pressure peaks at mid to low altitudes, with the particular altitude depending on the vehicle and trajectory. Above this altitude density is too low for significant dynamic pressure, and below this altitude the vehicle has slowed down to the point that dynamic pressure greatly reduces. A similar phenomenon occurs in reverse for launch vehicles. Therefore, it would be of interest to form a KLE that prioritizes density variation where it matters most for a given trajectory of interest; that is, where dynamic pressure is highest.

To that end, Fig. 5 shows results for a KLE model built from normalized density perturbations that are then scaled by dynamic pressure along the reference SHIELD entry trajectory, denoted KLE- q for convenience. Specifically, the KLE is constructed by multiplying the $\delta\rho$ values by the corresponding dynamic pressure at that altitude along the SHIELD reference trajectory, and the output of the KLE is then divided by that same dynamic pressure before converting the normalized perturbation back to density. In effect, this informs the KLE which altitude range is most important when ordering and truncating the expansion. As seen in Fig. 5, the 3σ bounds computed by this KLE closely match GRAM from about 60 km down to about 20 km, corresponding closely to the dynamic pressure pulse shown in Fig. 1. Given the fixed number of terms in the expansion, this comes at the expense of accuracy outside of that altitude range, where this




Figure 5. KLE vs. Mars-GRAM for KLE constructed from normalized density perturbations scaled by SHIELD dynamic pressure profile; thick dashed lines are $\pm 3\sigma$ bounds, thin solid lines are sample profiles

expansion underestimates variability.

As previously mentioned, the quality test for these density models is how well they predict dispersed trajectories compared to GRAM. To that end, a 1000-trial Monte Carlo analysis is performed for each of the three KLE models and for GRAM, where the only dispersed parameter in each analysis is density. A box plot comparing the statistics of peak heat flux for each case is shown in Fig. 6. The KLE- q model is by far the best match with the statistics predicted by GRAM directly, whereas the KLE- ρ and KLE- $\delta\rho$ models significantly underestimate variability. These results demonstrate that scaling normalized density perturbations by reference dynamic pressure is the most compact of the models considered here.

A similar scaling approach can be applied based on the reference aerocapture trajectory. This process is slightly more involved because during aerocapture the vehicle passes through each relevant altitude twice, with differing dynamic pressures, and has a minimum altitude well above the surface, as seen in Fig. 2. Recall, however, that the reference dynamic pressure is simply useful for re-scaling, and does not need to be dynamically valid. Thus, the following approach is taken in this study to form the reference dynamic pressure. Above the minimum altitude of the reference trajectory, the dynamic pressure during the descending portion of the trajectory is used for scaling. For another 10 km below the minimum altitude a constant value equal to the dynamic pressure at the minimum altitude is used; this segment exists because some dispersed trajectories will fly below the minimum altitude of the reference. Finally, a small but nonzero value (0.01 in this case) is used for scaling at more than 10 km below the minimum altitude of the reference trajectory. The resulting KLE model is shown in Fig. 7, and the corresponding peak heat flux results for Monte Carlo analyses of the aerocapture trajectory are shown in Fig. 8.




Figure 6. Peak heat flux statistics for SHIELD trajectories

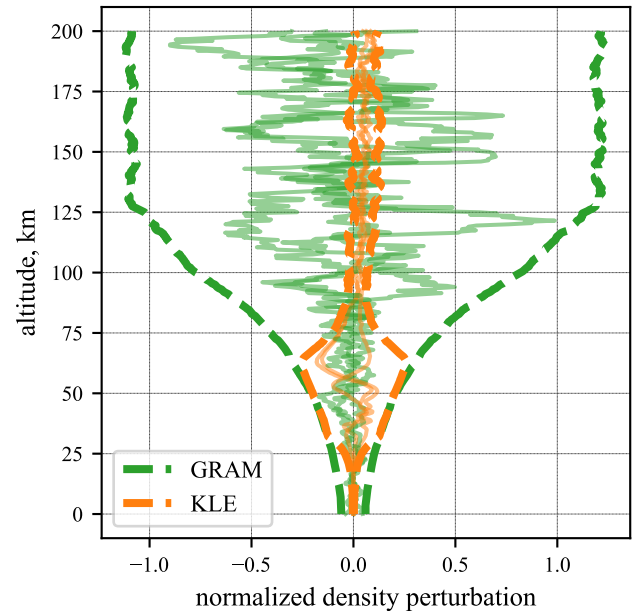


Figure 7. KLE vs. Mars-GRAM for KLE constructed from normalized density perturbations scaled by aerocapture dynamic pressure profile; thick dashed lines are $\pm 3\sigma$ bounds, thin solid lines are sample profiles




Figure 8. Peak heat flux statistics for aerocapture trajectories

Overall these results are similar to the corresponding results for SHIELD direct-entry, in that the KLE most efficiently captures density variation near the altitude of peak dynamic pressure and the KLE- q model performs best when predicting peak heat flux statistics. The altitude range where the aerocapture KLE- q accurately matches the GRAM 3σ bounds is shifted up by about 10 km compared to the SHIELD case due to peak dynamic pressure occurring at a higher altitude for the aerocapture trajectory.

These results demonstrate that the KLE- q model is the most compact representation of the three considered here as measured by the ability to predict statistics of peak heat flux. A relatively small number of terms ($d_K = 15$) is used for each KLE in order to highlight these differences. However, note that any of the KLE models should perform well if the number of included terms is sufficiently high, since the KLE representation of a GRF is exact for an infinite number of terms. For completeness, Fig. 9 shows the same results as Fig. 8 but with $d_K = 50$ terms included in each KLE model. The KLE- q model still performs the best, but the performance of the KLE- ρ and KLE- $\delta\rho$ models have improved considerably, such that there is less distinction between the three. The direct comparisons of density shown in Figs. 3–5 and 7 would also all improve with an increased d_K . This emphasizes an important point: any of these KLE flavors (and other possibilities not considered here) would form a valid density model if enough terms were included, but the comparisons in this paper illustrate that some approaches are more compact than others.

Furthermore, some differences in the density models are poorly captured by comparisons of peak heat flux statistics. Figure 10 directly shows density profiles for the altitude range near minimum altitude for the aerocapture Monte Carlo analyses using $d_K = 50$ for each model. This reveals another advantage to the KLE- q approach: because it captures density variability more efficiently, that model is able to capture more of the short-period oscillation behavior observed in GRAM




Figure 9. Peak heat flux statistics for aerocapture trajectories, with $d_K = 50$ terms

than the KLE- ρ and KLE- $\delta\rho$ models. Finally, note that the patterns that have been discussed here are somewhat tied to the choice of random variable; since peak heat flux occurs at mid-altitudes near peak dynamic pressure, the KLE- q model will be particularly efficient in capturing those statistics. For example, the landing locations for the SHIELD scenario and the apoapsis radius for the aerocapture scenario are both predicted reasonably well by the $d_K = 15$ KLE- ρ model, although the KLE- q model still performs better with the same number of terms.

5. MULTI-DIMENSIONAL KLE MODEL

Although the columnar assumption is typical for onboard models of density as previously discussed, in some cases it may be of interest to represent density as a random function of multiple independent variables. The KLE approximation demonstrated in Section 4 can be straightforwardly extended to model longitudinal and latitudinal variations in density as well as in altitude. Thus, in this section the necessary steps for constructing a multi-dimensional KLE are presented, models are compared following the approach taken in Section 4, and finally there is a brief discussion of the potential utility of these models for onboard use.

Recall that the first step in forming a KLE approximation from some discrete dataset is computing the sample covariance matrix as shown in Eq. 8. The data matrix Ψ_c is formed such that each column is one observation vector with the sample mean subtracted. In the columnar KLE model, the observation vectors are ordered such that they correspond with a reference altitude vector. For the more general case, however, the indexing of the data matrix Ψ_c need not refer to a single independent variable. Rather, the index corresponds to a specific variable being observed, whether that be defined as density at 100 km or as density at 100 km, 20° E, and 40° N. Any arbitrary set of points in a multi-dimensional domain can be uniquely identified via sequential indexing,




Figure 10. Density profiles near minimum altitude on aerocapture trajectories, with $d_K = 50$ terms




Figure 11. Illustration of reshaping between an arbitrary set of points in a multi-dimensional domain and an observation vector

and then observations at these points can be reshaped into a column vector following that ordering; this process is conceptually illustrated in Fig. 11. The process of computing the covariance matrix and constructing and evaluating the KLE is unchanged. The original reshaping is then reversed to reshape the column vectors produced by realizations of the KLE to a set of values for each point in the multi-dimensional domain.

As an example, MarsGRAM is used to generate 1000 density values at each point in an evenly-spaced 2D grid going from 0 to 200 km in altitude, from 0 to 10° in longitude, and at 0° latitude. Figure 12 visualizes the resulting data as a heatmap of the $+3\sigma$ value of $\delta\rho$; in other words, the heatmap values correspond to the right dashed line in figures like Fig. 3.

Figure 13 shows the equivalent statistics for a KLE approximation of the 2D MarsGRAM data; in contrast to Section 4, in this case a value of $\alpha = 0.99$ is used to truncate the KLE to $d_K = 884$ terms. From visual inspection, these results are




Figure 12. 3σ values of normalized density perturbation from MarsGRAM dataset

virtually indistinguishable from the GRAM data in Fig. 12.

As before, the real test of the KLE approximation is its ability to accurately predict trajectory dispersions. To this end, Figs. 14 and 15 show the peak heat flux statistics and a portion of the density profiles, respectively, resulting from 1000-trial Monte Carlo analyses of the same SHIELD direct-entry trajectory previously considered. In each case except GRAM 1D, bivariate spline approximation is used to compute density at the altitude and longitude of the vehicle based on a grid of density values. Recall that the reference SHIELD trajectory is ballistic and enters due-East, so the trajectory remains in the equatorial plane and thus, for this scenario, this approach is equivalent to computing density based on the 3D position of the vehicle. The GRAM 2D case interpolates from a set of density samples output by MarsGRAM directly, whereas the KLE $\alpha = 0.99$ case interpolates from values produced by a realization of an 884-term KLE approximation. The KLE $d_K = 50$ case also uses a KLE approximation, but in this case the expansion is limited to 50 terms. Finally, the GRAM 1D case interpolates from the same MarsGRAM data but always assumes a longitude of 0° , corresponding to a columnar assumption. This case should be exactly equivalent to the GRAM results shown for SHIELD in Fig. 6, but is slightly different. This occurs due to a quirk in how MarsGRAM density perturbations are computed. Thus, in this section the full 2D dataset is used but assuming a constant longitude of 0° in order to create an apples-to-apples comparison.

From Fig. 14, it is clear that the peak heat flux statistics predicted by the 2D GRAM and 2D KLE ($\alpha = 0.99$) models are very similar, and Fig. 15 shows a characteristic similarity between the density profiles predicted by these two models. These results and the direct comparison of density values in Figs. 12 and 13 demonstrate the successful use of a multi-dimensional KLE to approximate density as a function of both altitude and longitude. In contrast, the 50-term KLE approximation performs very poorly, significantly under-predicting both the mean and uncertainty of peak heat




Figure 13. 3σ values of normalized density perturbation from 2D KLE approximation

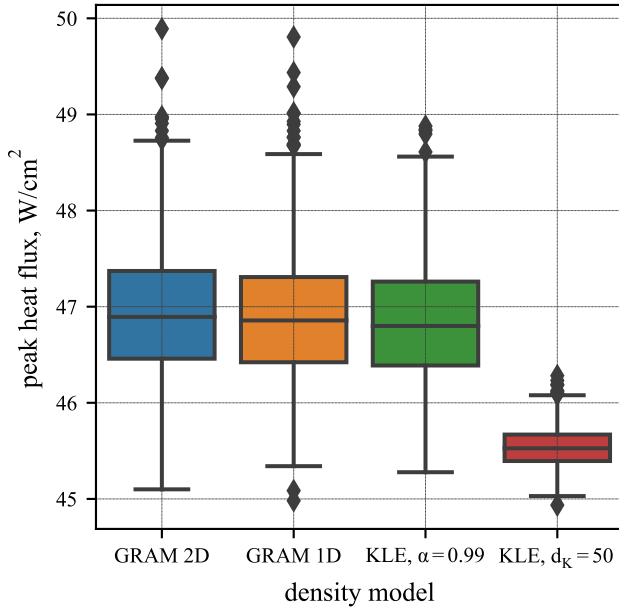


Figure 14. Peak heat flux statistics for SHIELD trajectories in 2D atmosphere models




Figure 15. Density profiles on SHIELD trajectories for 2D density models

flux. The expansion fails to capture much of the variability in density, as is clear from Fig. 15. The KLE $d_K = 50$ case performs worse than the KLE $\alpha = 0.99$ case because it has a much lower number of terms (50 vs. 884), and the expansion is truncated before sufficiently capturing the modes of variability present in the data.

These comparisons merit a broader discussion of the columnar atmosphere approximation for onboard density modeling. Figure 14 shows that the *GRAM 1D* case, which is equivalent to a columnar atmosphere assumption, almost exactly matches the 2D *GRAM* case in predictions of peak heat flux, and from Fig. 15 the sample density profiles themselves also appear to be very similar. This is not surprising when considering Fig. 12, which shows no significant horizontal gradient to indicate changes in density variability with longitude. Note that, despite this uniformity in longitude, the KLE requires roughly 10x as many terms to accurately predict dispersed trajectories when constructed from the 2D density data as opposed to the columnar atmosphere case. This would require an increase in both memory and computational expense for onboard use. Furthermore, to sample across the entire 2D grid in altitude and longitude requires 8505 datapoints for the discretization used here, as compared to 405 datapoints for a columnar profile, further exacerbating the onboard computational burden. These results suggest that, based on the dataset used here, a columnar atmosphere model is likely a good enough approximation for onboard use, and is significantly less demanding of both memory and computational effort than a multi-dimensional model.

This is decidedly *not* to say that regional variations in density can be neglected. Density gradients occur due to a range of factors including gravity waves, time of day, and winds, and are relevant for both vehicle performance prediction and trajectory reconstruction [29], [30], [31]. MarsGRAM data is used in this study as an example only, and is not necessarily well-suited to capturing these types of regional

density variation. Any hypersonic vehicle using closed-loop guidance would need to be simulated in a wide range of possible atmospheric conditions, regardless of the assumptions used for the onboard density model. The resulting vehicle performance, taken together with the relevant computational limitations, is ultimately what determines whether or not the onboard density model meets requirements.

Note also that, for a scenario where density is expected to change significantly along the groundtrack of an entry trajectory, a columnar model could be constructed using data generated along the reference trajectory. In other words, the raw data is generated along a 3D trajectory, but is then treated as a function of only altitude in the KLE approximation. This approach begins to fail if altitude is not monotonically decreasing, such as in the case of aerocapture. However, the procedure for onboard measurement updates presented in the next section would potentially result in different density predictions for the descending and ascending portions of the trajectory, and this could partly mitigate the limitations of a columnar model.

6. KALMAN MEASUREMENT UPDATES

During atmospheric flight, observations of estimated density $\rho^*(h_k)$ are typically available by taking estimated sensed accelerations measurements from an accelerometer or inertial measurement unit (IMU) and rearranging the equation for aerodynamic acceleration,

$$a(h_k) = \frac{v^2(h_k)}{2\beta} \rho(h_k) \rightarrow \rho^*(h_k) = \frac{2\beta a^*(h_k)}{v^2(h_k)}, \quad (13)$$

where estimates of the ballistic coefficient and current velocity magnitude are known. Thus, for any onboard density model to be useful in practice, it should accommodate some method of updating the model in real-time with noisy measurements. It is well-demonstrated in literature and in practice that appropriate onboard density estimation can significantly improve targeting performance [2].

The novel benefit of a KLE density model is the representation of both a nominal density profile and associated uncertainty. Therefore, it is desirable to formulate an approach that updates both the mean and covariance represented by the KLE. Furthermore, this should be done in a way that respects the correlation structure assumed in the pre-update model, as opposed to replacing a single diagonal element of the covariance matrix. For clarity, this section returns to the columnar atmosphere assumption.

In this work a Bayesian approach for sequential estimation is applied, such that the mean and covariance of density from the previous update (or the initial model) form the prior, and these are updated with the noisy density measurement to form the posterior mean and covariance of density. The density estimates are assumed to be corrupted by additive white Gaussian noise, based on the assumption that some pre-processing removes artifacts such as IMU drift; note that this also implies accurate estimates for ballistic coefficient and velocity magnitude. The state uncertainty is also Gaussian based on the earlier assumption treating density as a Gaussian random process. Finally, density estimates are assumed to arrive at altitude points included in the original a priori density model, either by judiciously timing measurement updates or by interpolating multiple measurements.

Based on the above assumptions, density can be optimally

estimated by the Kalman measurement update via the following formulation [32]. Take the series of density values at each altitude to be the state vector. The dynamic equation is trivial, since the density profile is assumed not to vary in time, so the state propagation step from the Kalman filter is unnecessary. The measurement equation is simply a direct observation of a single state component and is thus linear. Therefore, the optimal estimate of the vector of atmospheric density at each altitude $\hat{\rho}^+ \in \mathbb{R}^n$ and its covariance $\mathbf{P}_k^+ \in \mathbb{R}^{n \times n}$ can be computed according to a scalar noisy density measurement $\rho_k^* \in \mathbb{R}$ according to the following equations:

$$\hat{\rho}^+ = \hat{\rho}^- + \mathbf{K}(\rho_k^* - \mathbf{H}_k \hat{\rho}^-), \quad (14)$$

$$\mathbf{P}^+ = \mathbf{P}^- - \mathbf{K} \mathbf{H}_k \mathbf{P}^-, \quad (15)$$

$$\mathbf{K} = \mathbf{P}^- \mathbf{H}_k^\top (\mathbf{H}_k \mathbf{P}^- \mathbf{H}_k^\top + \mathbf{R})^{-1}, \quad (16)$$

$$\mathbf{H}_k = [\delta_{1k}, \delta_{2k}, \dots, \delta_{nk}], \quad (17)$$

where $\mathbf{K} \in \mathbb{R}^{n \times 1}$ is the Kalman gain matrix, $\mathbf{H}_k \in \mathbb{R}^{1 \times n}$ is the measurement matrix, $\mathbf{R} \in \mathbb{R}$ is the measurement noise variance, δ_{ij} is the Kronecker delta, n is the number of discrete altitudes considered, and k is the index of the altitude at which density is currently being observed. Notably, because only one density is measured at a time the bracketed term in Eq. (16) is a scalar, so taking its inverse is computationally inexpensive.

For notational clarity, consider an example where the discretization of density values is from 100 to 0 km in altitude steps of 0.5 km, in descending order, resulting in $n = 201$. Then $\hat{\rho}^-$ and $\hat{\rho}^+$ are the prior and posterior 201-vectors, respectively, containing density values at each altitude. Assume the scalar density measurement ρ_k^* is at an altitude of 80 km, such that $k = 41$ (indexing from 1 in this notation). Then, \mathbf{H}_k becomes a row matrix with all elements equal to zero except in the 41st column, which is equal to one.

Equations (14) – (17) can be applied to sequentially ingest noisy density measurements and update the onboard model of the density profile and its covariance. By re-solving for the eigenvalues and eigenvectors of the \mathbf{P}^+ , the KLE representation can be updated accordingly. This process is demonstrated in Figs. 16 and 17; here, the prior mean and covariance are formed from a dataset of 3000 density profiles from MarsGRAM, where density perturbations are normalized by the sample mean and thus the normalized prior mean falls exactly along 0. The true profile to be estimated is also computed by MarsGRAM, but is not included in the prior dataset. Five density values are observed, corrupted by measurement noise with a standard deviation of 1×10^{-9} kg/m³, a value selected purely for illustrative purposes. In this example the assumed measurement noise \mathbf{R} is equal to the true noise value, but note that this can instead be treated as a tuning parameter in practice and need not be the same value at each altitude.

Note that the posterior mean passes nearly through each observation (with one exception), but reverts to the mean for altitudes above and below the observation altitudes. The posterior uncertainty bounds are also only weakly affected at these higher and lower altitudes. This occurs because the correlation structure in the prior covariance dictates the degree to which new information at one altitude affects the estimated density at another altitude; since in this model density perturbation at 80 km is only weakly correlated with density perturbation at 50 km, the posterior mean has reverted to nominal by that lower altitude. This can also be achieved




Figure 16. Mean and 3σ bounds for prior and posterior density profiles, given five sequential noisy observations

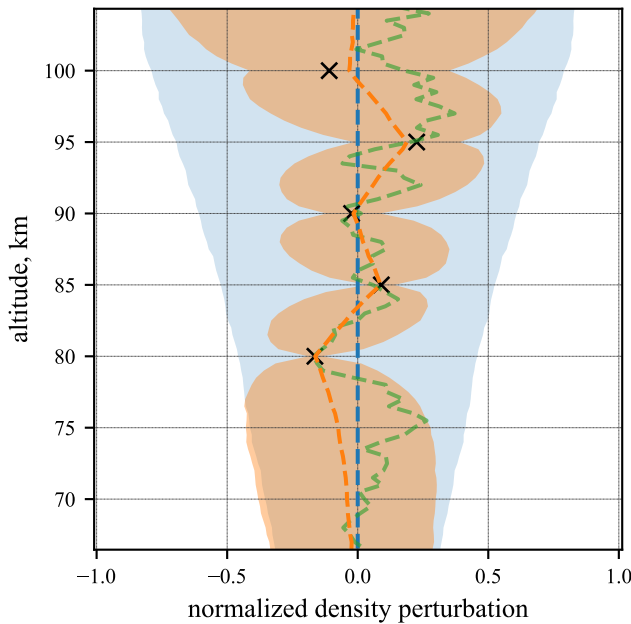


Figure 17. Zoomed-in view of Fig. 16

by exponentially correlating a corrective scale factor that is estimated onboard, but the approach presented here has two advantages. First, the correlation length is inferred from the prior model (MarsGRAM in this case) rather than defined by the user, and second, the correlation length is not necessarily constant with altitude.

The reason that the posterior mean passes more closely through the lower three measurements than the first two is related to how measurement noise was defined. Measurement noise is applied to the density values directly and is constant across all altitudes, but the data is then converted to normalized density perturbations for estimation and visualization. Thus, at higher altitudes where nominal density is significantly lower, the measurement noise has a more significant effect, and the filter tends to trust the prior. This is also observable by the much wider posterior uncertainty bounds for the higher-altitude measurements. At lower altitudes the same measurement noise has relatively less effect and the situation is reversed; since the actual measurement noise and the value used by the filter are the same, this also means the lower-altitude measurements fall closer to the truth values. It should be reiterated that the assumption of a measurement noise constant with altitude is made here for demonstration purposes and is not required.

The application of a Kalman measurement update demonstrated here provides a way of updating the mean and covari-

ance for atmospheric density based on noisy measurements, which could inform onboard predictions of state uncertainty for the purpose of closed-loop guidance. A significant drawback of this approach, however, is the requirement to re-solve the eigenvalues and eigenvectors after each measurement update in order to obtain the updated KLE representation. This adds significant computational expense to the update process, potentially to the point of infeasibility for onboard computation, depending on the resolution of the density profile and the choice of flight hardware. This motivates an approach that updates the eigenvectors and eigenvalues directly in a way that approximates the result of the Kalman measurement update at a lower computational expense. Such a method could take advantage of the fact that the covariance matrix is low-rank, meaning there is approximately zero covariance between altitudes more than a certain distance apart. Approaches such as low-rank partial Hessian approximations or sequential updates to singular value decompositions of a matrix provide potential pathways to significant computational efficiency improvement [33], [34]; this remains an area for future work.

7. CONCLUSIONS

This work presents the mathematical foundation and practical implementation for modeling density as a KLE. For the direct-entry and aerocapture scenarios considered here, a KLE constructed by scaling normalized density perturbations by the reference dynamic pressure is shown to be the best predictor of peak heat flux. Directly forming the KLE from density or normalized density perturbations is less compact but also gives accurate predictions, and could be the more straightforward approach if the necessary number of terms is allowable based on computational limitations. A KLE formed over a multi-dimensional domain is demonstrated, but for the MarsGRAM data considered here the gains compared to a columnar model are unlikely to outweigh the additional computational expense. Finally, a Kalman measurement update is used to update the density covariance matrix based on new density measurements, and the example results show promising behavior. However, further work is necessary to improve the computational efficiency of this approach for onboard implementation.

ACKNOWLEDGMENTS

This work was supported by a NASA Space Technology Research Fellowship.

REFERENCES

- [1] S. N. D'Souza and N. Sarigul-Klijn, "Survey of planetary entry guidance algorithms," *Progress in Aerospace Sciences*, vol. 68, pp. 64–74, Jul. 2014.
- [2] E. Perot and S. Rousseau, "Importance of an on-board estimation of the density scale height for various aerocapture guidance algorithms," in *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, 2002.
- [3] Z. R. Putnam and R. D. Braun, "Drag-modulation flight-control system options for planetary aerocapture," *Journal of Spacecraft and Rockets*, vol. 51, no. 1, pp. 139–150, 2014.
- [4] J. Wagner, A. Wilhite, D. Stanley, and R. Powell, "An adaptive real time atmospheric prediction algorithm for entry vehicles," in *3rd AIAA Atmospheric Space Environments Conference*, 2011.
- [5] D. Amato, S. Hume, B. Grace, and J. McMahon, "Robustifying mars descent guidance through neural networks," in *AAS Guidance, Navigation, and Control Conference*, no. AAS 20-073, 2020.
- [6] E. Roelke, P. Hattis, and R. Braun, "Improved atmospheric estimation for aerocapture guidance," in *AAS/AIAA Astrodynamics Specialist Conference*, 08 2019, pp. 1–16.
- [7] J. Ridderhof and P. Tsiotras, "Chance-Constrained Covariance Steering in a Gaussian Random Field via Successive Convex Programming," *Journal of Guidance, Control, and Dynamics*, vol. 45, no. 4, pp. 599–610, Apr. 2022. [Online]. Available: <https://arc.aiaa.org/doi/10.2514/1.G005941>
- [8] J. W. McMahon, D. Amato, D. Kuettel, and M. J. Grace, *Stochastic Predictor-Corrector Guidance*.
- [9] X. Jiang and S. Li, "Mars entry trajectory planning using robust optimization and uncertainty quantification," *Acta Astronautica*, vol. 161, pp. 249–261, Aug. 2019.
- [10] J. Ridderhof, S. W. Albert, P. Tsiotras, and H. Schaub, "Linear covariance analysis of entry and aerocapture trajectories in an uncertain atmosphere," in *AIAA SCITECH 2022 Forum*, San Diego, CA, 2022.
- [11] D. Woffinden, S. Robinson, J. Williams, and Z. R. Putnam, "Linear covariance analysis techniques to generate navigation and sensor requirements for the safe and precise landing integrated capabilities evolution (SPLICE) project," in *AIAA Scitech Forum*, San Diego, CA, Jan. 7 – 11 2019, p. 0662.
- [12] A. Doostan and G. Iaccarino, "A least-squares approximation of partial differential equations with high-dimensional random inputs," *Journal of Computational Physics*, vol. 228, no. 12, pp. 4332–4345, 2009.
- [13] A. Doostan and H. Owhadi, "A non-adapted sparse approximation of pdes with stochastic inputs," *Journal of Computational Physics*, vol. 230, no. 8, pp. 3015–3034, 2011.
- [14] X. Jiang, "Uncertainty quantification for mars atmospheric entry using polynomial chaos and spectral decomposition," in *2018 AIAA Guidance, Navigation, and Control Conference*, 2018.
- [15] A. Halder and R. Bhattacharya, "Dispersion analysis in hypersonic flight during planetary entry using stochastic liouville equation," *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 2, pp. 459–474, 2011.
- [16] C. R. Heidrich and R. D. Braun, "Aerocapture trajectory design in uncertain entry environments," in *AIAA Scitech 2020 Forum*, 2020.
- [17] S. W. Albert, A. Doostan, and H. Schaub, "Finite-dimensional density representation for aerocapture uncertainty quantification," in *AIAA Scitech 2021 Forum*, no. AIAA 2021-0932, Virtual event, Jan. 2021.
- [18] H. Justh, "Mars global reference atmospheric model 2010 version: Users guide," NASA, Tech. Rep. NASA/TM—2014–217499, 2014.
- [19] J.-F. Le Gall, *Brownian Motion, Martingales, and Stochastic Calculus*, ser. Graduate Texts in Mathematics. Springer International Publishing, 2016.
- [20] C. E. Rasmussen and C. K. I. Williams, *Gaussian*

Processes for Machine Learning, ser. Adaptive Computation and Machine Learning. The MIT Press, 2005.

- [21] M. Loeve, "Elementary probability theory," in *Probability theory i*. Springer, 1977, pp. 1–52.
- [22] D. D. Kosambi, "Statistics in function space," *Journal of the Indian Mathematical Society*, vol. 7, pp. 76–88, 1943.
- [23] D. A. Vallado, *Fundamentals of Astrodynamics and Applications*, 4th ed. Microcosm Press, 2013, ch. Appendix D, pp. 1041–1042.
- [24] C. G. Justus and R. D. Braun, "Atmospheric environments for entry, descent and landing (edl)," in *5th International Planetary Probes Workshop and Short Course*, 2007, pp. 1–37, <https://ntrs.nasa.gov/citations/20070032693>.
- [25] M. G. Trainer, M. H. Wong, T. H. McConnochie, H. B. Franz, S. K. Atreya, P. G. Conrad, F. Lefèvre, P. R. Mahaffy, C. A. Malespin, H. L. Manning, J. Martín-Torres, G. M. Martínez, C. P. McKay, R. Navarro-González, A. Vicente-Retortillo, C. R. Webster, and M.-P. Zorzano, "Seasonal variations in atmospheric composition as measured in gale crater, mars," *Journal of Geophysical Research: Planets*, vol. 124, no. 11, pp. 3000–3024, 2019.
- [26] N. Barba, T. Komarek, R. Woolley, L. Giersch, V. Stamenković, M. Gallagher, and C. D. Edwards, "Mars small spacecraft studies: Overview," in *2019 IEEE Aerospace Conference*, 2019, pp. 1–10.
- [27] D. Way, R. Powell, A. Chen, A. Steltzner, M. San Martín, P. Burkhart, and G. Mendeck, "Mars science laboratory: Entry, descent, and landing system performance," in *IEEE Aerospace Conference Proceedings*, March 2006, pp. 1–19.
- [28] J. Ridderhof and P. Tsiotras, "Chance-constrained covariance steering in a Gaussian random field via successive convex programming," 2021. [Online]. Available: <https://arxiv.org/abs/2101.09634>
- [29] M. A. Mischna, G. Villar, D. M. Kass, S. Dutta, S. Rafkin, D. Tyler, J. Barnes, B. Cantor, S. R. Lewis, D. Hinson, J. Pla-García, A. Kleinböhl, and C. Karlgaard, "Pre- and post-entry, descent and landing assessment of the martian atmosphere for the mars 2020 rover," *The Planetary Science Journal*, vol. 3, no. 6, p. 147, jun 2022. [Online]. Available: <https://doi.org/10.3847/psj/ac7148>
- [30] P. Withers, "Mars global surveyor and mars odyssey accelerometer observations of the martian upper atmosphere during aerobraking," *Geophysical Research Letters*, vol. 33, no. 2, 2006.
- [31] S. Bougher, G. Keating, R. Zurek, J. Murphy, R. Haberle, J. Hollingsworth, and R. T. Clancy, "Mars Global Surveyor aerobraking: Atmospheric trends and model interpretation," *Advances in Space Research*, vol. 23, no. 11, pp. 1887–1897, 1999.
- [32] D. Simon, *Optimal State Estimation*. John Wiley & Sons, Ltd, 2006.
- [33] H. P. Flath, L. C. Wilcox, V. Akçelik, J. Hill, B. van Bloemen Waanders, and O. Ghattas, "Fast algorithms for bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial hessian approximations," *SIAM Journal on Scientific Computing*, vol. 33, no. 1, pp. 407–432, 2011.
- [34] M. Brand, "Fast low-rank modifications of the thin

singular value decomposition," *Linear Algebra and its Applications*, vol. 415, no. 1, pp. 20–30, 2006, special Issue on Large Scale Linear and Nonlinear Eigenvalue Problems.

BIOGRAPHY




by Dr. Hanspeter Schaub.

Samuel W. Albert received his B.S. in Aeronautical and Astronautical Engineering from Purdue University in 2018 and an M.S. in Aerospace Engineering Sciences from the University of Colorado Boulder in 2020. He is currently a PhD Candidate and NASA Space Technology Research Fellow in the Autonomous Vehicle Systems laboratory at CU Boulder, where he is advised



Dr. Alireza Doostan is a Professor of Aerospace Engineering Sciences Department at the University of Colorado Boulder and an affiliated faculty of the Applied Mathematics Department. Prior to his appointment at CU Boulder in 2010, he was an Engineering Research Associate in the Center for Turbulence Research at Stanford University. Alireza received his PhD in Structural Engineering and M.A. in Applied Mathematics and Statistics from the Johns Hopkins University both in 2007. He is a recipient of a DOE (ASCR) and an NSF (Engineering Design) Early Career awards, as well as multiple teaching awards from CU Boulder and AIAA. His research interests include: Uncertainty quantification, reduced order and data-driven modeling, optimization under uncertainty, machine learning, and computational stochastic mechanics.



Dr. Hanspeter Schaub is a professor and chair of the University of Colorado aerospace engineering sciences department. He holds the Schaden leadership chair. He has over 28 years of research experience, of which 4 years are at Sandia National Laboratories. His research interests are in astrodynamics, relative motion dynamics, charged spacecraft motion as well as spacecraft autonomy. This has led to about 207 journal and 327 conference publications, as well as a 4th edition textbook on analytical mechanics of space systems. Dr. Schaub has been the ADCS lead in the CICERO mission, the ADCS algorithm lead on a Mars mission and supporting ADCS for a new asteroid mission. He has been awarded the H. Joseph Smead Faculty Fellowship, the Provost's Faculty Achievement Award, the faculty assembly award for excellence in teaching, as well as the Outstanding Faculty Advisor Award. He is a fellow of AIAA and AAS, and has won the AIAA/ASEE Atwood Educator award, AIAA Mechanics and Control of Flight award, as well as the Collegiate Educator of the Year for the AIAA Rocky Mountain section.