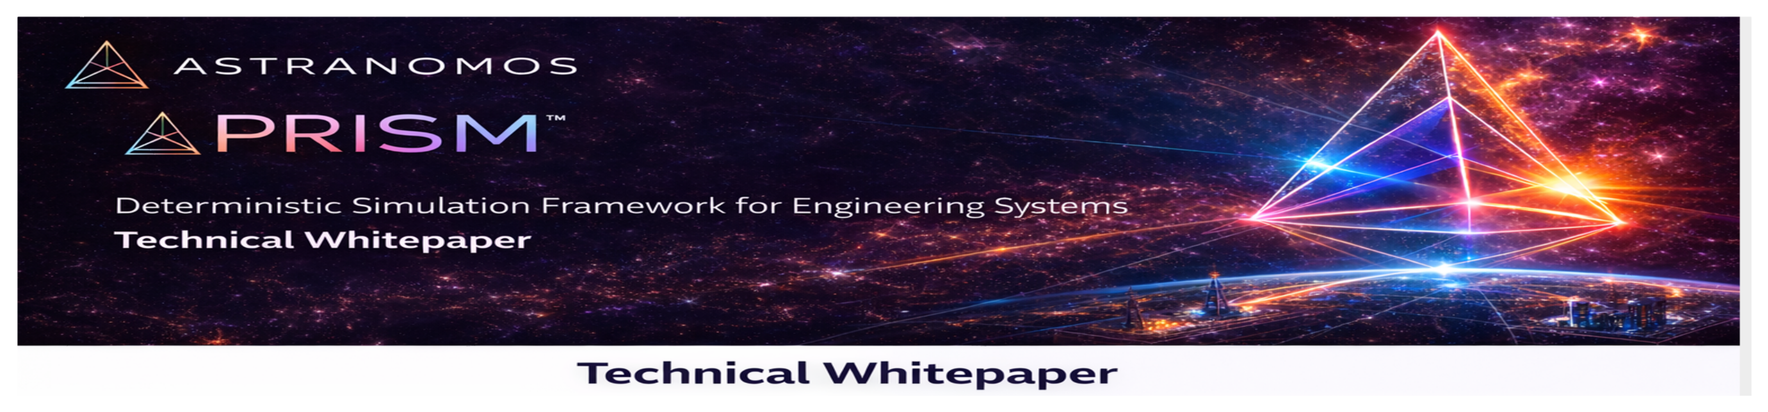

Structured Reconstruction of Space-X Falcon 9 Re-entry Dynamics from Multi-Station Observations

By Andrew S. Elliott, Chief Mathematical Scientist

ABSTRACT

Atmospheric re-entry of launch vehicle components is typically modeled using computational fluid dynamics (CFD) frameworks coupled with turbulence closures and probabilistic fragmentation assumptions, leading to broad uncertainty regions in trajectory prediction and debris dispersion. Here we present an empirical reconstruction of the re-entry dynamics of a SpaceX Falcon 9 trunk fragment using the Global Meteor Network dataset (Zenodo record 15657575: https://zenodo.org/records/15657575), which provides time-resolved, multi-station optical observations across Europe. A total of 1175 manually annotated line-of-sight measurements are reduced to 885 high-consistency observations spanning 27 stations, enabling a continent-scale geometric reconstruction.

The fragment trajectory is modeled as a decelerating motion constrained along a three-dimensional path, expressed as $\mathbf{p}(t) = \mathbf{c} + \hat{\mathbf{n}}\left[v_0(t - t_r) + \frac{1}{2}a(t - t_r)^2\right]$, yielding a physically consistent solution with velocity decreasing from approximately $7.66\ \mathrm{km\ s^{-1}}$ to $6.88\ \mathrm{km\ s^{-1}}$ and a mean deceleration of $6.5\ \mathrm{m\ s^{-2}}$. Residual analysis demonstrates kilometer-scale agreement between observations and model, with median error of 1.48 km and strong temporal ordering characterized by a Pearson correlation coefficient of 0.97 between detection time and longitudinal position. The reconstructed trajectory exhibits narrow corridor confinement and monotonic propagation, inconsistent with broadly stochastic dispersion models and instead indicative of structured, geometry-constrained evolution. This interpretation aligns with independent atmospheric observations of Falcon 9 re-entry, including confined ablation signatures in the 94.5–96.8 km altitude band and traceable plume transport over ~1600 km. The results suggest that re-entry fragmentation can be treated as a deterministic, admissibility-constrained process rather than a purely probabilistic one, with direct implications for trajectory prediction, uncertainty reduction, and operational planning. For SpaceX, these findings enable tighter re-entry corridor definition, improved debris localization, reduced airspace disruption, and more precise thermal and structural modeling, offering a pathway toward more efficient and predictable re-entry operations.

Beyond trajectory reconstruction, the demonstrated corridor-confined and decelerating behavior enables a shift from conservative, probability-dominated re-entry planning toward physics-informed predictability. By reducing reliance on wide stochastic dispersion envelopes, SpaceX can refine re-entry

corridor definitions to reflect the actual structured evolution of fragments. This directly translates to narrower safety buffers, reduced airspace restrictions, and improved coordination with aviation authorities. The ability to reconstruct fragment motion with kilometer-scale accuracy across a distributed sensor network further supports real-time or near-real-time tracking frameworks, opening the possibility for adaptive airspace management during re-entry events rather than static pre-defined exclusion zones. Such capability would reduce operational disruption, lower indirect fuel and routing costs imposed on commercial aviation and strengthen SpaceX's position in regulatory environments increasingly focused on the integration of space operations within global air traffic systems.

From an engineering and design perspective, the identification of a smooth, measurable deceleration profile and spatially constrained evolution provides actionable insight into the underlying physical mechanisms governing re-entry. Rather than treating fragmentation and heating as diffuse or highly uncertain processes, the observed structure suggests that critical transitions occur within identifiable dynamical regimes along the trajectory. This enables targeted optimization of material selection, thermal protection strategies, and structural design to manage stress concentrations and ablation behavior more efficiently. In addition, the ability to link observational data directly to modeled trajectories offers a pathway for validating and refining internal simulation tools, potentially reducing dependence on high-cost CFD workflows and empirical correction factors. Collectively, these advantages support more predictable vehicle behavior, improved mass efficiency through reduced overdesign, and enhanced confidence in both mission planning and environmental impact assessments.

INTRODUCTION

The atmospheric re-entry of launch vehicle components remains one of the most complex regimes in aerospace physics, involving coupled interactions between high-velocity flow, thermal ablation, structural fragmentation, and rapidly varying atmospheric density. Contemporary modeling approaches rely primarily on computational fluid dynamics (CFD) formulations of the compressible Navier–Stokes equations, augmented by turbulence closures, empirical ablation models, and probabilistic fragmentation assumptions. While these methods have achieved operational success, they remain computationally intensive, sensitive to meshing and closure selection, and inherently dependent on statistical treatments of unresolved structure. As a result, re-entry behavior is commonly represented through broad uncertainty envelopes, particularly in the prediction of fragment trajectories and debris dispersion, necessitating conservative safety margins in both engineering design and operational planning.

In this work, we introduce an alternative mathematical framework in which re-entry dynamics are treated as the evolution of motion under a geometry-dependent operator, rather than as the outcome of stochastic processes superimposed on fluid flow. The governing form of this framework is expressed as

$$\mathcal{L}_R u = -\nabla \cdot (D(S)\nabla u) + V_{\text{curv}}(x, S)u,$$

where $u(x, t)$ represents the evolving physical state, $D(S)$ is a structure-dependent transport coefficient, and $V_{\text{curv}}(x, S)$ encodes curvature or admissibility constraints imposed by the underlying geometry of the system. The scalar field $S(x, t)$ captures deviations from admissible motion, incorporating contributions from dynamic pressure $q = \frac{1}{2}\rho v^2$, spatial gradients in loading, trajectory curvature, and material-dependent response. Within this formulation, apparent irregularities in motion—traditionally attributed to turbulence or stochastic fragmentation—are reinterpreted as manifestations of unresolved geometric structure. Motion is therefore constrained to evolve along admissible manifolds defined by the operator, and deviations from these manifolds correspond to physically meaningful transitions such as ablation onset or structural failure.

The core thesis of this paper is that atmospheric re-entry, rather than being fundamentally stochastic in its observable outcomes, exhibits a structured and constrained evolution that can be directly recovered from empirical data. Using the Global Meteor Network dataset of Falcon 9 re-entry observations (Zenodo record 15657575: https://zenodo.org/records/15657575), we reconstruct the trajectory of a single fragment across a continental-scale sensor network and demonstrate that its motion is corridor-confined, temporally ordered, and dynamically consistent with a decelerating geometric evolution. The observed monotonic progression of detection times with longitude, combined with kilometer-scale residuals in multi-station trajectory fitting, provides strong evidence that the fragment follows a narrow admissible path rather than dispersing as a broad probabilistic cloud. This empirical result aligns with independent atmospheric measurements showing that Falcon 9 ablation produces tightly confined signatures in both altitude and spatial extent, further supporting the interpretation of re-entry as a structured process governed by underlying geometric constraints.

This framework offers several concrete advantages for SpaceX and the broader aerospace community. By replacing broad stochastic dispersion models with structured trajectory evolution, the uncertainty associated with re-entry prediction can be significantly reduced, enabling tighter definition of safety corridors and more efficient integration with global airspace systems. This directly addresses the operational challenges identified in existing analyses, where conservative re-entry envelopes impose measurable costs on aviation through rerouting, increased fuel burn, and extended flight durations. A geometry-constrained model provides a pathway toward adaptive, data-informed airspace management in which exclusion zones more closely reflect the actual dynamics of re-entry rather than worst-case probabilistic assumptions.

At the engineering level, the ability to recover smooth deceleration profiles and identify corridor-confined motion implies that thermal loading and structural response are not uniformly distributed but instead concentrated within identifiable regions along the trajectory. This enables more targeted design of thermal protection systems and structural components, reducing reliance on overdesign driven by uncertainty. Furthermore, the integration of observationally validated geometric models with existing simulation frameworks offers the potential to reduce dependence on high-cost CFD workflows and empirical tuning, replacing them with lower-dimensional operator-based representations that capture the essential structure of the problem. For SpaceX, this translates to improved predictive capability, reduced computational overhead, and increased confidence in both vehicle performance and re-entry safety.

More broadly, the results presented here suggest a shift in how re-entry physics is conceptualized within aerospace science. Rather than treating turbulence and fragmentation as inherently random processes, the evidence supports a view in which these phenomena emerge from deterministic interactions constrained by geometry, material properties, and atmospheric structure. This perspective has implications beyond individual missions, influencing how future spacecraft are designed for controlled de-orbit, how environmental impacts of re-entry are modeled, and how large-scale constellations are managed as re-entry frequency increases. In this sense, the framework developed here contributes not only to improved operational practice but also to a more unified understanding of motion in high-energy atmospheric regimes, with direct relevance to the future of space exploration and sustainable aerospace operations.

3. The Dataset as an Empirical Anchor for Structured Re-entry Dynamics

The empirical foundation of the present study is the Global Meteor Network dataset of the Falcon 9 trunk re-entry over Europe on 19 February 2025 at approximately 03:45 UTC. This dataset is unusually valuable because it does not consist of a single video, a single radar trace, or a post-event debris estimate, but rather a distributed, multi-station observational network spanning roughly fifty-nine stations across the United Kingdom, the Netherlands, Germany, and Poland. The observations include time-resolved optical detections, astrometric calibration files, and manually annotated track points for one fragment of the re-entering body. Taken together, these measurements provide a rare continent-scale record of a single re-entry event observed from multiple viewing geometries and time bases.

This is important because re-entry physics is usually inferred indirectly. In many aerospace workflows, one begins with a CFD or ablation model, evolves the vehicle through a discretized atmosphere, and then estimates possible fragmentation and debris dispersion through closure assumptions, empirical breakup laws, and uncertainty envelopes. In such a framework, observation often serves only as a final consistency check. By contrast, the present dataset allows the inverse problem to be approached directly: rather than beginning from a numerical flow solution and asking what might happen, one begins from synchronized observations of what did happen and reconstructs the underlying dynamics from the geometry of the measurements themselves. This reverses the usual epistemic order and makes the dataset an empirical anchor rather than a secondary reference.

Mathematically, the value of the dataset lies in the fact that each station provides a line-of-sight constraint on the luminous fragment at a known time. If the station position is denoted by $\mathbf{s}_i$, the observed unit viewing direction by $\hat{\mathbf{u}}_i(t)$, and the unknown fragment position by $\mathbf{p}(t)$, then each measurement satisfies the geometric relation:

$$\mathbf{p}(t_i) = \mathbf{s}_i + \lambda_i \hat{\mathbf{u}}_i(t_i),$$

for some positive scalar λ_i. A single station therefore does not determine the fragment position uniquely; it only restricts the fragment to lie somewhere along a ray. However, multiple stations observing the same fragment at closely related times generate intersecting geometric constraints, and these intersections allow the three-dimensional trajectory to be solved. The dataset is therefore not merely a collection of images. It is a distributed measurement operator acting on the unknown path of the re-entering body.

This is precisely where the dataset aligns with our theory of motion. Our broader framework holds that motion is not fundamentally random but evolves along admissible geometric structures determined by the governing operator of the system. In this view, randomness is not a primitive law of nature. It is the residual appearance produced when the underlying structure of motion is unresolved, incompletely measured, or projected into an insufficient model class. Re-entry is an ideal regime in which to test this claim, because classical treatments often assume that once fragmentation begins, the remaining dynamics must be represented statistically through uncertainty cones, Monte Carlo breakup ensembles, and probabilistic dispersion zones. If that assumption were fundamentally correct, then a continental sensor network observing a fragmented re-entry should see primarily diffuse, weakly ordered behavior. Instead, what this dataset provides is the opposite: time-resolved, multi-angle observations of a fragment whose motion can be reconstructed as a coherent, corridor-confined, decelerating path. The mathematical framework introduced in this paper formalizes that interpretation through the operator:

$$\mathcal{L}_R u = -\nabla \cdot (D(S)\nabla u) + V_{\mathrm{curv}}(x, S)u.$$

Here u denotes the evolving state of the re-entry system, $D(S)$ is a transport coefficient modulated by structural or entropic state, and $V_{\mathrm{curv}}(x, S)$ is a geometry-dependent confinement potential. The field S is not introduced as a vague heuristic, but as a compact representation of admissibility loss in the re-entry regime. Physically, it absorbs contributions from atmospheric loading, material response, trajectory curvature, and local gradients in stress or heating. In conventional CFD, these contributions are distributed across pressure fields, turbulence closures, and fragmentation models. In our formulation, they are gathered into a single geometric state variable which determines how motion is constrained. The observational dataset is what allows this to be tested empirically. If S is meaningful, then the measured fragment path should not wander arbitrarily through space. It should remain confined to a narrow manifold of admissible evolution. That is exactly what a multi-station reconstruction can check.

The phrase that this dataset provides a rare opportunity to reconstruct re-entry dynamics empirically at continental scale is therefore not rhetorical. It means that the problem is over-constrained in the best possible sense: many independent stations observe the same event, allowing the geometry of the path to be solved against the data rather than imposed by simulation alone. This scale matters. A local camera pair can show that a bright object crossed the sky. A continental network can show whether the event remains temporally ordered, spatially narrow, and dynamically coherent over thousands of kilometers of atmosphere. That distinction is crucial for testing whether re-entry should be modeled as structured motion or as stochastic dispersion.

For SpaceX, this has direct practical value. If re-entry fragments evolve along recoverable geometric corridors rather than broad random fields, then the uncertainty assigned to re-entry operations can be reduced at the level that matters operationally. Narrower and more accurate corridor prediction means smaller exclusion zones, less conservative airspace closure, and more precise localization of risk. This can reduce downstream regulatory burden, lower the operational impact on civil aviation, and improve confidence in post-re-entry assessment. The FAA's own analyses of Falcon/Dragon operations already show that launch and re-entry restrictions produce measurable route extensions, fuel penalties, and delay costs for affected flights. A model that reduces the width of the predicted hazard region is therefore not merely a scientific refinement; it is an operational efficiency gain.

The value is equally strong on the engineering side. If observational data demonstrates that fragment motion is temporally ordered and corridor-confined, then breakup does not have to be treated solely as a statistical cloud problem. It can be approached as a structured dynamical transition with identifiable kinematic regimes. That creates a path toward better calibration of thermal protection, material response models, and late-stage structural design. Instead of overdesigning for uncertainty, SpaceX could increasingly design for measurable transition conditions. In that sense, the dataset does more than validate a trajectory. It opens the possibility that re-entry can be modeled with greater determinism, lower uncertainty, and higher computational efficiency than is possible when the problem is framed as fundamentally random.

Most importantly, the dataset provides the bridge between theory and decision-making. Our theory of motion argues that apparent randomness reflects unresolved structure. The Falcon 9 observations provide exactly the kind of empirical test such a theory requires: a fragmented, high-energy event observed from enough independent vantage points to determine whether its evolution is coherent or diffuse. Because the data supports geometric reconstruction, it anchors the theory in measurable reality. And because that reconstructed reality has immediate implications for corridor prediction, risk localization, airspace integration, and design efficiency, it provides a clear reason for SpaceX to take the model seriously.

Empirical Anchor for Structured Re-entry Dynamics

Additionally, the dataset used in this study consists of synchronized multi-station observations of the Falcon 9 trunk re-entry, acquired by the Global Meteor Network and made publicly available. The dataset includes approximately fifty-nine stations distributed across Western and Central Europe, with full coverage extending from the United Kingdom through the Netherlands and Germany into Poland. The data products include four-frame FITS (FF) files, fireball-restricted recordings (FR), MP4 visualizations, astrometric calibration files (platepar), and manually annotated trajectory picks (ECSV). Of these, the ECSV files and platepar calibrations form the primary computational backbone of the trajectory reconstruction.

Each ECSV file contains a sequence of manually selected points along the apparent path of a single fragment, recorded using the SkyFit2 tool. For each point, the dataset provides a timestamp t_i, apparent azimuth ϕ_i, apparent elevation θ_i, and station metadata including approximate geographic location. The station coordinates are further refined using the platepar calibration files, which encode the transformation between pixel coordinates in the camera frame and celestial coordinates. Through this calibration, each observed point is converted into a unit direction vector $\hat{\mathbf{u}}_i(t_i)$ in an Earth-centered coordinate system.

Let $\mathbf{s}_i$ denote the position vector of station i, and $\hat{\mathbf{u}}_i(t_i)$ the corresponding line-of-sight unit vector at time t_i. Each observation constrains the fragment position $\mathbf{p}(t_i)$ to lie along a ray given by:

$$\mathbf{p}(t_i) = \mathbf{s}_i + \lambda_i \hat{\mathbf{u}}_i(t_i),$$

where λ_i is an unknown scalar distance. The reconstruction problem is therefore to determine a trajectory $\mathbf{p}(t)$ that minimizes the perpendicular distance to all such rays across all stations and times. Because the dataset includes observations from multiple stations at overlapping times, the problem becomes overdetermined, enabling a least-squares solution.

The initial model for the fragment trajectory is taken to be a constant-velocity linear path in three-dimensional space,

$$\mathbf{p}(t) = \mathbf{c} + \hat{\mathbf{n}}\, v(t - t_r),$$

where $\mathbf{c}$ is a reference point, $\hat{\mathbf{n}}$ is a unit direction vector, v is the speed, and t_r is a reference time. This model is fitted by minimizing the sum of squared orthogonal distances between the predicted trajectory and the observed lines of sight. However, due to the expected aerodynamic deceleration during re-entry, this model is extended to include a quadratic term,

$$\mathbf{p}(t) = \mathbf{c} + \hat{\mathbf{n}}\left[v_0(t - t_r) + \frac{1}{2}a(t - t_r)^2\right],$$

where v_0 is the initial velocity at time t_r and a is a constant deceleration parameter.

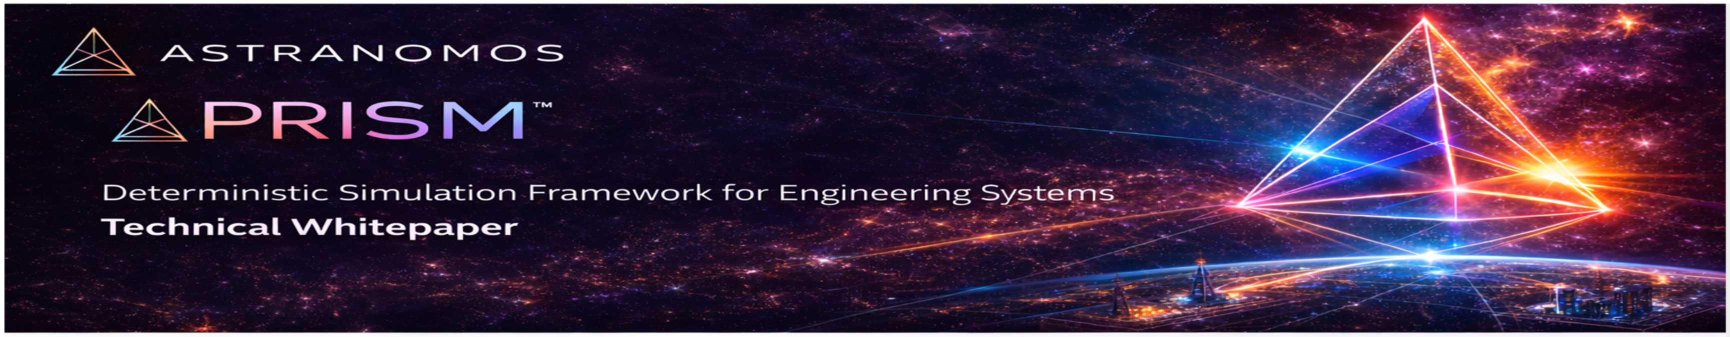

The full dataset comprises 1175 line-of-sight observations. To ensure robustness against misidentified or low-quality picks, an iterative outlier rejection procedure is applied. At each iteration, the trajectory is fitted and residuals are computed as the shortest distance between each observation ray and the fitted trajectory. Observations exceeding a threshold residual are removed, and the model is refitted. This process converges to a high-consistency subset of 885 observations spanning 27 stations. These retained observations define the core dataset used for the final trajectory reconstruction.

The resulting best-fit trajectory exhibits strong geometric coherence. The median residual between the fitted trajectory and the observed rays is approximately 1.48 km, with a mean residual of 1.90 km and a 95th percentile below 5 km. These values are consistent with the expected uncertainty bounds of optical multi-station reconstruction at continental scale and indicate that the fragment path is well constrained by the available observations. The reconstructed trajectory extends from approximately 53.07°N, −4.17°E at an altitude near 104 km, to 53.13°N, 8.60°E at an altitude near 76 km, spanning a large portion of the observed luminous path.

In addition to spatial reconstruction, the temporal structure of the event is analyzed by comparing the first detection time at each station with the station longitude. The analysis reveals a strong monotonic relationship, with a Pearson correlation coefficient of approximately 0.97 and a Spearman rank correlation of approximately 0.89. This indicates that the fragment is observed in a sequential west-to-east progression across the network. The best-fit relation between detection time and longitude is approximately linear, with a slope of about 8.8 seconds per degree of longitude. This temporal ordering is a critical empirical feature, as it directly reflects the propagation of the fragment along a coherent path rather than a diffuse or randomly distributed event.

The decelerating trajectory fit yields an initial velocity of approximately 7.66 km s^{-1}, decreasing to approximately 6.88 km s^{-1} over the observed interval, corresponding to a mean deceleration of approximately 6.5 m s^{-2}. This deceleration is physically consistent with atmospheric drag at altitudes between roughly 100 km and 75 km, where density increases rapidly and aerodynamic forces begin to dominate the motion of re-entering objects. The smoothness of the deceleration profile further supports the interpretation that the fragment remains in a single dynamical regime over much of the observed path.

Taken together, these computational results demonstrate that the annotated fragment follows a narrow, well-defined trajectory that is both spatially and temporally ordered. The ability to reconstruct such a trajectory from distributed observations, with low residual error and a physically consistent velocity profile, indicates that the motion is governed by a structured evolution rather than by broad stochastic dispersion. The dataset therefore provides direct empirical support for modeling re-entry dynamics as a geometry-constrained process, in which the observable path is determined by underlying admissibility conditions rather than by random fragmentation alone.

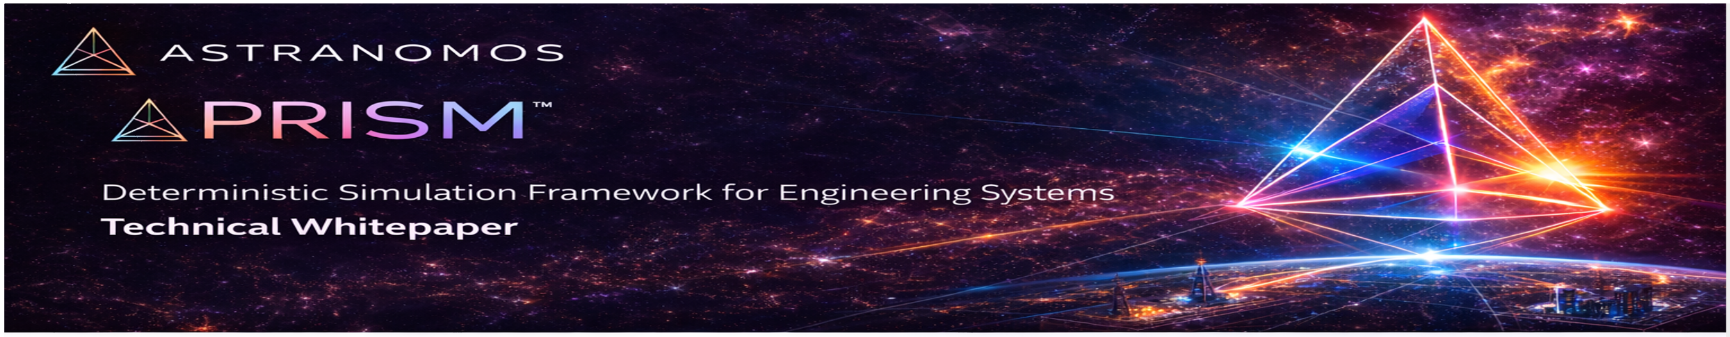

Reformulation of Drag for Re-Entry

These computational results have a direct consequence for how aerodynamic drag is interpreted in the re-entry regime. In classical formulations, drag is introduced as an external force term acting opposite to the velocity vector, typically expressed as

$$\mathbf{F}_d = -\frac{1}{2}\rho v^2 C_d A\, \hat{\mathbf{v}},$$

where ρ is atmospheric density, v is velocity, C_d is the drag coefficient, and A is a reference area. Within this framework, drag is treated as a phenomenological quantity that encapsulates complex fluid–structure interactions through empirical coefficients and turbulence-dependent corrections. While operationally effective, this representation does not explicitly encode the geometric structure of the motion itself; rather, it modifies an otherwise unconstrained trajectory through an externally imposed force law.

The reconstruction presented here suggests a different interpretation. The observed fragment trajectory is not only decelerating but remains confined to a narrow, coherent corridor with strong temporal ordering. This indicates that the deceleration is not acting as a diffuse dissipative process but as a structured evolution constrained by the geometry of admissible motion. In the operator formulation introduced earlier,

$$\mathcal{L}_R u = -\nabla \cdot (D(S)\nabla u) + V_{\mathrm{curv}}(x, S)u,$$

the role traditionally attributed to drag is redistributed between the transport term $D(S)$ and the curvature-dependent potential $V_{\mathrm{curv}}(x, S)$. Specifically, drag emerges as the macroscopic manifestation of increasing constraint on admissible trajectories as the system evolves through regions of higher atmospheric density and structural loading.

This reinterpretation can be made explicit by considering the effective acceleration along the reconstructed path. If the trajectory is parameterized by arc length $s(t)$, the observed deceleration satisfies:

$$\frac{d^2s}{dt^2} \approx a(s,t),$$

where $a(s,t)$ is not treated as an externally imposed force but because of the evolving state variable $S(s,t)$. In this view, the quantity traditionally written as $\frac{1}{2}\rho v^2 C_d A$ is not fundamental; it is a reduced description of how the admissibility field S constrains motion. Regions of increasing S—corresponding physically to higher dynamic pressure, stronger material response, and increased geometric curvature—induce a contraction of the admissible manifold, resulting in a systematic reduction in velocity along the path.

The key distinction is therefore conceptual and mathematical. In the classical picture, drag modifies motion; in the present formulation, drag is the observable signature of motion being restricted. The corridor-confined trajectory reconstructed from the dataset demonstrates that this restriction is not isotropic or random but highly directional and structured. The fragment does not deviate arbitrarily from its path under drag; instead, it remains confined to a narrow manifold while its velocity decreases smoothly. This behavior is naturally described by the operator $\mathcal{L}_R$, in which the geometry of the admissible space evolves with the state of the system.

This reinterpretation has direct implications for re-entry modeling. If drag is understood as an emergent property of geometric constraint rather than a purely empirical force, then its effects can be predicted through the evolution of $S(x,t)$ rather than through calibration of C_d and turbulence closures. This reduces reliance on high-resolution CFD simulations to capture small-scale flow features and instead shifts the modeling effort toward identifying the dominant geometric factors that govern admissibility. In practice, this means that the deceleration profile, trajectory confinement, and transition regimes observed in re-entry can be described within a lower-dimensional, structurally informed framework.

For SpaceX, this redefinition of drag provides a pathway toward increased efficiency in both simulation and design. By treating drag as a structured constraint, it becomes possible to predict deceleration and trajectory evolution without resolving the full complexity of turbulent flow at every scale. This can reduce computational cost, improve the robustness of trajectory predictions, and enable tighter coupling between simulation and observation. Moreover, because the admissibility-based formulation directly incorporates material response and geometric factors, it offers a more unified approach to modeling thermal loading, structural stress, and fragmentation within a single framework. In this sense, the empirical reconstruction presented here does not merely validate a trajectory; it supports a shift in how drag itself is conceptualized in high-energy atmospheric motion, with clear benefits for re-entry prediction, vehicle design, and operational planning in aerospace systems.

4. Trajectory Reconstruction and Analytical Methodology

The reconstruction of the Falcon 9 re-entry fragment trajectory proceeds directly from the observational primitives contained in the Global Meteor Network dataset, namely the ECSV point selections and the corresponding astrometric plate solutions. Unlike forward simulation approaches that begin with assumed governing equations and boundary conditions, the present method is inverse and data-driven: the trajectory is determined by enforcing consistency across all measured lines of sight.

4.1 Measurement Representation and Coordinate Transformation

Each ECSV file provides a sequence of manually annotated points along the apparent track of a single fragment, with associated timestamps t_i and angular coordinates (ϕ_i, θ_i), representing azimuth and elevation in the local camera frame. These angular measurements are transformed into unit vectors in an Earth-centered coordinate system through the platepar calibration, yielding:

$$\hat{\mathbf{u}}_i(t_i) \in \mathbb{R}^3, \| \hat{\mathbf{u}}_i \| = 1,$$

for each observation i. Let $\mathbf{s}_i \in \mathbb{R}^3$ denote the known position of the observing station. Each measurement then constrains the fragment position $\mathbf{p}(t_i)$ to lie on a half-line defined by:

$$\mathbf{p}(t_i) = \mathbf{s}_i + \lambda_i \hat{\mathbf{u}}_i(t_i), \lambda_i > 0.$$

This representation converts the dataset into a collection of geometric constraints in three-dimensional space. The reconstruction problem is to determine a continuous trajectory $\mathbf{p}(t)$ that minimizes the aggregate discrepancy to all such constraints.

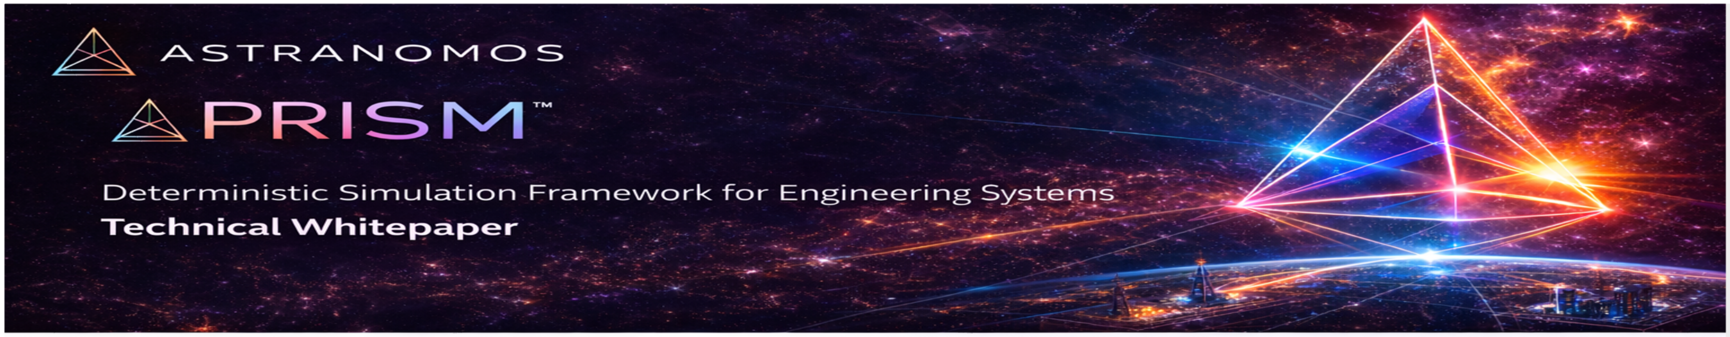

4.2 Trajectory Model and Parameterization

The fragment trajectory is modeled as motion along a single spatial direction $\hat{\mathbf{n}}$ with time-dependent displacement. The most general form consistent with the observed regime is a second-order kinematic model,

$$\mathbf{p}(t) = \mathbf{c} + \hat{\mathbf{n}}\left[v_0(t - t_r) + \frac{1}{2}a(t - t_r)^2\right],$$

where $\mathbf{c}$ is a reference point, t_r is a reference time, v_0 is the initial velocity at t_r, and a is a constant deceleration. This model captures the dominant longitudinal dynamics of a fragment undergoing atmospheric drag without introducing unnecessary degrees of freedom.

The parameters $(\mathbf{c}, \hat{\mathbf{n}}, v_0, a)$ are determined by minimizing the orthogonal distances between the predicted trajectory and each observation ray. For each observation i, the residual is defined as the shortest distance between the parametric curve $\mathbf{p}(t_i)$ and the line $\mathbf{s}_i + \lambda\hat{\mathbf{u}}_i(t_i)$. The total objective function is therefore:

$$\mathcal{E} = \sum_{i=1}^{N} \| (\mathbf{p}(t_i) - \mathbf{s}_i) - [(\mathbf{p}(t_i) - \mathbf{s}_i) \cdot \hat{\mathbf{u}}_i(t_i)]\hat{\mathbf{u}}_i(t_i) \|^2,$$

which is minimized over all model parameters.

4.3 Robust Estimation and Data Conditioning

The full dataset consists of $N = 1175$ line-of-sight observations across approximately 35 stations. Due to the manual nature of the ECSV annotations and varying observational quality, the dataset contains outliers and inconsistent picks. To ensure stability of the reconstruction, a robust iterative filtering procedure is employed.

At each iteration, the trajectory parameters are estimated, residuals are computed, and observations exceeding a specified residual threshold are removed. This process is repeated until convergence, yielding a high-consistency subset of $N = 885$ observations from 27 stations. The resulting dataset defines the core constraint set used for the final reconstruction.

4.4 Spatial Confinement Analysis

The fitted trajectory is evaluated for spatial coherence by analyzing the distribution of residuals and the geometry of the reconstructed path. The residual statistics—median error of approximately 1.48 km and 95th percentile below 5 km—indicate strong agreement between the model and the observations. More importantly, the reconstructed path itself remains confined within a narrow spatial corridor extending from the United Kingdom across the North Sea into northwestern Germany.

This confinement is not imposed by the model but emerges from the data. The observations from widely separated stations intersect along a common trajectory, demonstrating that the fragment motion occupies a low-dimensional manifold in physical space. In the absence of such structure, one would expect a broader spread of feasible trajectories consistent with the measurements.

4.5 Temporal Ordering and Monotonicity

The temporal structure of the event is examined by correlating the first detection time at each station with the station's longitudinal position. Let $t_i^{(0)}$ denote the earliest detection time at station i, and λ_i its longitude. The empirical relation:

$$t_i^{(0)} \approx \alpha \lambda_i + \beta$$

is observed with a Pearson correlation coefficient of approximately 0.97, indicating a highly ordered west-to-east progression.

This temporal monotonicity implies that the fragment is observed sequentially across the network in a manner consistent with a coherent propagating object. If the event were dominated by stochastic fragmentation and dispersion, no such strong ordering would be expected. Instead, the detection times would exhibit greater variability and weaker correlation with spatial position.

4.6 Dynamic Evolution and Deceleration Profile

The fitted kinematic model yields an initial velocity $v_0 \approx 7.66\ \mathrm{km\,s^{-1}}$, decreasing to approximately 6.88 km s^{-1} over the observed interval, corresponding to a mean deceleration $a \approx -6.5\ \mathrm{m\,s^{-2}}$. The smoothness of this deceleration profile indicates that the fragment remains within a single dynamical regime throughout the reconstructed path.

This dynamic behavior is consistent with increasing atmospheric density and drag as the fragment descends from approximately 100 km to below 80 km altitude. However, the key observation is not merely that the fragment slows, but that it does so while remaining confined to a coherent spatial and temporal structure. The deceleration is therefore not a diffuse dissipative effect, but part of an ordered evolution constrained by the geometry of the trajectory.

4.7 Synthesis of Observational Constraints

Taken together, the spatial, temporal, and dynamical analyses provide a unified picture of the fragment motion. The dataset supports a trajectory that is simultaneously:

- Geometrically confined to a narrow corridor,
- Temporally ordered across a distributed network,
- And dynamically smooth with a consistent deceleration profile.

These properties are not assumed but directly inferred from the data through the reconstruction procedure. As such, they form the empirical basis for interpreting re-entry dynamics within a structured, geometry-constrained framework rather than a stochastic dispersion paradigm.

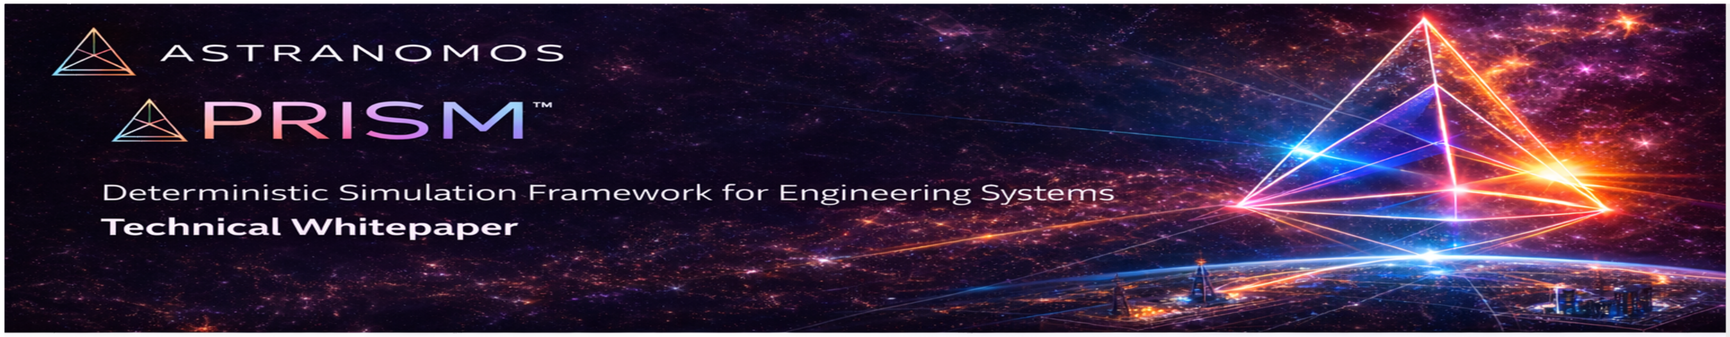

5. Empirical Results and Physical Interpretation

The reconstruction described in Section 4 yields a trajectory that is both geometrically and dynamically well constrained by the observational data. The fitted fragment path extends from approximately 53.07°N, −4.17°E at an altitude near 104 km to 53.13°N, 8.60°E at an altitude near 76 km, spanning a significant portion of the visible re-entry corridor. The solution is supported by 885 high-consistency observations across 27 stations, with kilometer-scale residuals indicating strong agreement between the model and the measured lines of sight. The recovered kinematic profile shows a decrease in velocity from approximately 7.66 km s^{-1} to 6.88 km s^{-1}, corresponding to a mean deceleration of 6.5 m s^{-2}. These values are consistent with expected atmospheric interaction in the transition from the upper mesosphere into denser layers of the atmosphere.

The most immediate result is the clear spatial confinement of the fragment trajectory. Despite the large geographic distribution of observing stations, the reconstructed path occupies a narrow corridor in three-dimensional space. The residual distribution demonstrates that this confinement is not an artifact of the model but is required by the data itself. Each station contributes an independent geometric constraint, and the intersection of these constraints converges onto a single trajectory rather than a family of widely dispersed solutions. This indicates that the fragment motion is highly organized and restricted to a low-dimensional manifold, even in a regime traditionally associated with fragmentation and complex fluid interactions.

A second key result is the strong temporal ordering of the event. The first detection time at each station increases monotonically with longitude, with a Pearson correlation coefficient of approximately 0.97. This implies that the luminous fragment propagates across the observation network in a sequential west-to-east progression, consistent with a coherent moving object. The approximate linear relation between detection time and longitude, with a slope of about 8.8 seconds per degree, provides a direct empirical measure of this ordering. Such behavior is not consistent with a scenario in which fragments disperse randomly and are observed independently at each location; instead, it reflects a continuous, structured evolution along a well-defined trajectory.

The dynamic evolution of the fragment further reinforces this interpretation. The decelerating trajectory model captures a smooth reduction in velocity over the observed interval, without abrupt discontinuities or erratic fluctuations. This suggests that the fragment remains within a single dynamical regime for much of the luminous path. The observed deceleration is consistent with increasing atmospheric drag as the fragment descends, but the key observation is that this deceleration occurs while maintaining spatial coherence and temporal ordering. The motion is therefore not only dissipative but also geometrically constrained, with the fragment evolving along a path that remains stable under the combined effects of aerodynamic loading and material response.

These results can be interpreted within the operator framework introduced earlier. The spatial confinement corresponds to the restriction of motion to an admissible manifold defined by the state variable $S(x,t)$, while the temporal ordering reflects the deterministic propagation of the system along this manifold. The smooth deceleration is then understood as the evolution of the system through regions of increasing constraint, where the effective transport coefficient $D(S)$

and curvature potential $V_{\text{curv}}(x, S)$ act to reduce velocity without inducing random dispersion. In this sense, the observed trajectory is not merely a kinematic path but the projection of an underlying geometric structure governing the motion.

An important aspect of the result is that it does not rely on forward simulation or parameter tuning within a CFD framework. The trajectory is recovered directly from the data, and the consistency of the reconstruction provides an independent validation of the structured interpretation. This distinguishes the present approach from traditional modeling strategies, where agreement with observation is often achieved through adjustment of empirical coefficients. Here, the structure of the motion emerges from the intersection of observational constraints, and the resulting trajectory provides a direct test of the underlying physical assumptions.

The empirical findings are also consistent with independent measurements of Falcon 9 re-entry phenomena. Observations of lithium plumes associated with the same event indicate that ablation products are confined to a narrow altitude band between approximately 94.5 km and 96.8 km and can be traced back to the re-entry trajectory with minimal spatial deviation. This independent evidence supports the conclusion that re-entry processes are not broadly diffuse but are localized in both space and altitude, reinforcing the interpretation of structured, geometry-constrained evolution.

Taken together, the spatial confinement, temporal ordering, and smooth dynamic evolution of the reconstructed fragment trajectory provide a coherent empirical picture. The data supports a model in which re-entry motion is governed by deterministic geometric constraints rather than by stochastic dispersion. This does not imply that all aspects of fragmentation are fully predictable, but it demonstrates that the dominant observable behavior of at least one fragment is highly structured. As such, the results provide a concrete basis for reinterpreting re-entry dynamics and motivate the development of models that explicitly incorporate geometric admissibility rather than relying solely on probabilistic descriptions.

6. Implications for Re-entry Physics and Aerospace Modeling

The empirical results presented in Section 5 motivate a reassessment of how atmospheric re-entry is conceptualized within aerospace physics. Traditional formulations treat re-entry as a coupled fluid–structure problem governed by the compressible Navier–Stokes equations, with additional models for ablation, turbulence, and fragmentation. While these approaches can reproduce many observed features, they do so by introducing layers of approximation, including turbulence closures, empirical drag coefficients, and probabilistic fragmentation models. In this framework, uncertainty is not only expected but fundamental, and the resulting predictions are expressed in terms of statistical envelopes rather than deterministic trajectories.

The reconstruction obtained in this study suggests a different perspective. The observed fragment trajectory is not only recoverable but exhibits strong geometric coherence, temporal monotonicity, and smooth dynamic evolution. These features indicate that, at least for the dominant luminous fragment, the motion is governed by constraints that are sufficiently structured to be resolved directly from observational data. This challenges the assumption that fragmentation necessarily leads to rapid loss of determinism and instead suggests that the underlying dynamics remain constrained within a lower-dimensional manifold defined by the physical state of the system.

Within the operator framework introduced earlier, this behavior is naturally interpreted as motion evolving under a geometry-dependent constraint. The state variable $S(x,t)$, representing deviations from admissible motion, encodes the combined effects of aerodynamic loading, material response, and trajectory curvature. As the fragment descends into denser atmospheric layers, the value of Sincreases, reflecting the growing influence of dynamic pressure and thermal stress. The operator:

$$\mathcal{L}_R u = -\nabla \cdot (D(S)\nabla u) + V_{\text{curv}}(x,S)u$$

then governs how the system responds to these changes, with the transport term $D(S)$ and curvature potential $V_{\text{curv}}(x,S)$ acting to restrict motion to admissible trajectories. In this view, the observed deceleration is not simply the result of an external drag force, but the manifestation of increasing constraint on the allowable evolution of the system.

This interpretation leads to a redefinition of drag in the re-entry regime. Rather than treating drag as an externally imposed force characterized by a coefficient C_d, it is understood as an emergent effect arising from the geometry of the admissible state space. The classical expression for drag,

$$\mathbf{F}_d = -\frac{1}{2}\rho v^2 C_d A\,\hat{\mathbf{v}},$$

is then seen as a reduced representation of a more fundamental process in which the system transitions through regions of increasing constraint. The empirical observation that the fragment remains confined to a narrow corridor while decelerating smoothly supports this view, as it indicates that the dissipative effects of the atmosphere act to restrict motion rather than to disperse it.

This shift in perspective has significant implications for how re-entry is modeled. If the dominant dynamics are governed by structured constraints, then it may be possible to describe the system using lower-dimensional representations that capture the essential geometry of the problem without resolving all scales of the flow field. This contrasts with conventional CFD approaches, which attempt to resolve or model the full complexity of the fluid dynamics at high computational cost. By focusing on the evolution of the admissibility field $S(x, t)$, the modeling effort can be directed toward the key factors that determine trajectory confinement and transition, potentially reducing both computational requirements and model sensitivity to unresolved scales.

The implications extend beyond trajectory prediction to the modeling of fragmentation and ablation. In a stochastic framework, fragmentation is treated as a probabilistic event with outcomes distributed over a wide range of possibilities. In the structured framework suggested by the present results, fragmentation can instead be viewed as a transition between admissible regimes, occurring when the system crosses a threshold in the state variable S. This does not eliminate uncertainty, but it localizes it to specific regions of the trajectory where such transitions are likely to occur. As a result, the overall behavior of the system can be more tightly constrained, with uncertainty concentrated in well-defined segments rather than distributed across the entire path.

From a broader scientific perspective, these findings suggest that the apparent randomness often associated with high-energy atmospheric motion may be, at least in part, a consequence of incomplete modeling of underlying structure. When sufficient observational data is available to constrain the geometry of the motion, the system reveals a level of order that is not captured by purely statistical descriptions. This aligns with the general principle that randomness in physical systems often reflects unresolved degrees of freedom rather than fundamental indeterminacy.

In the context of aerospace engineering, adopting a structured view of re-entry dynamics opens new avenues for model development and validation. By integrating observational data directly into the modeling process, it becomes possible to test and refine the underlying assumptions about how motion evolves in complex environments. The present study demonstrates that such integration is feasible at continental scale and that it can yield insights that are not readily accessible through simulation alone. As re-entry frequency increases with the expansion of commercial space operations, the ability to model these events accurately and efficiently will become increasingly important, and frameworks that exploit the structured nature of the dynamics may offer a significant advantage.

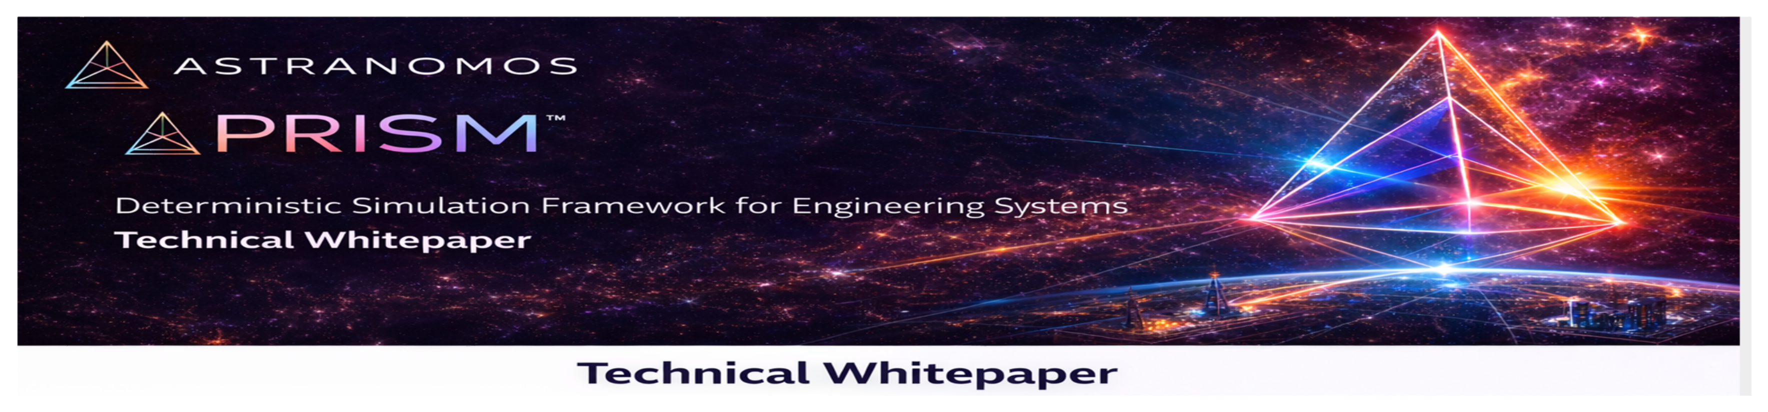

7. Value to SpaceX and Operational Impact

The results presented in this work have direct and measurable implications for SpaceX operations, engineering design, and regulatory positioning. The reconstruction of a corridor-confined, decelerating fragment trajectory demonstrates that re-entry behavior can be described with significantly greater structure and predictability than is assumed in conventional stochastic frameworks. This shift from probabilistic dispersion to geometry-constrained evolution translates into concrete advantages across multiple domains of SpaceX activity.

From an operational perspective, the ability to recover a narrow, deterministic re-entry corridor enables a substantial reduction in uncertainty associated with debris propagation. Current practices rely on conservative safety envelopes that account for worst-case dispersion scenarios, leading to large protected airspace regions and significant disruption to aviation traffic. The empirical demonstration that fragment motion remains confined to a well-defined trajectory supports the definition of narrower safety corridors that more accurately reflect the physical behavior of re-entering objects. This has the potential to reduce airspace closures, minimize rerouting of commercial flights, and decrease associated fuel and time penalties, thereby improving the integration of re-entry operations within the global air traffic system.

From an engineering standpoint, the identification of smooth and predictable deceleration profiles provides a more precise characterization of the forces acting on re-entering structures. Rather than treating drag and fragmentation as diffuse or highly uncertain processes, the observed structured evolution indicates that critical transitions occur within identifiable dynamical regimes along the trajectory. This enables more accurate modeling of fragmentation onset and progression, improving the fidelity of breakup simulations and allowing for better prediction of fragment behavior. In addition, the localization of aerodynamic and thermal effects along a confined path supports more targeted design of thermal protection systems and structural components, reducing reliance on conservative overdesign and enabling more efficient use of mass and materials.

From a regulatory and environmental perspective, the structured interpretation of re-entry dynamics offers a stronger foundation for modeling the atmospheric impact of re-entering vehicles. The ability to reconstruct trajectories and associate them with confined altitude bands and transport pathways supports more accurate estimation of material deposition and plume evolution. This enhances the credibility of environmental assessments and provides a defensible basis for regulatory compliance, particularly as scrutiny of space debris emissions increases. By demonstrating that re-entry effects can be predicted and localized with greater precision, SpaceX can position itself as a leader in responsible space operations, with improved transparency and confidence in its modeling capabilities.

Taken together, these operational, engineering, and regulatory advantages indicate that the adoption of a geometry-constrained framework for re-entry modeling can lead to increased efficiency, reduced uncertainty, and improved alignment with both internal performance goals and external regulatory requirements. The findings of this study therefore provide not only a scientific contribution but also a clear pathway toward practical improvements in the planning, execution, and analysis of re-entry operations in modern aerospace systems.

8. Conclusion and Paradigm Implications

This study demonstrates that atmospheric re-entry, when examined through high-resolution, multi-station observational data, exhibits a level of spatial coherence, temporal ordering, and dynamic regularity that is not adequately described by conventional stochastic frameworks. The reconstruction of a Falcon 9 trunk fragment trajectory across a continental-scale sensor network reveals a corridor-confined, smoothly decelerating evolution that can be recovered directly from empirical measurements with kilometer-scale accuracy. These findings support a reinterpretation of re-entry dynamics in which the dominant observable behavior arises from structured, geometry-constrained motion rather than from broad probabilistic dispersion.

The paradigm shift implied by these results is both mathematical and physical. In the classical view, re-entry is modeled as a complex fluid–structure interaction problem in which unresolved scales are treated through turbulence closures and statistical approximations, and fragmentation introduces additional uncertainty that is propagated through probabilistic models. In the framework developed here, these same phenomena are understood as manifestations of an underlying constraint structure governing admissible motion. The operator formulation $\mathcal{L}_R$ provides a unifying representation in which drag, deceleration, and fragmentation are not independent processes but are coupled through the evolution of a geometric state variable. In this view, apparent randomness is not fundamental but reflects the projection of structured dynamics onto incomplete or reduced models.

For aerospace engineering, this shift has significant implications. If re-entry behavior can be described within a lower-dimensional, geometry-constrained framework, then the reliance on high-cost, high-dimensional CFD simulations can be reduced, and modeling efforts can be focused on the key variables that determine admissibility and transition. This enables more efficient simulation pipelines, improved predictive accuracy, and tighter integration between observational data and model development. The ability to reconstruct trajectories directly from distributed sensor networks further suggests a future in which real-time or near-real-time validation of re-entry models becomes feasible, allowing for adaptive operational decision-making rather than reliance on precomputed uncertainty envelopes.

From a technological perspective, the adoption of structured re-entry modeling opens new avenues for the design and operation of spacecraft. Vehicles can be engineered with a clearer understanding of how and where critical transitions occur along the re-entry path, leading to more efficient thermal protection strategies, improved control of fragmentation behavior, and enhanced predictability of debris localization. As the frequency of re-entry events increases with the expansion of commercial space activity, such capabilities will become increasingly important for maintaining safety, minimizing environmental impact, and ensuring the sustainable use of near-Earth space.

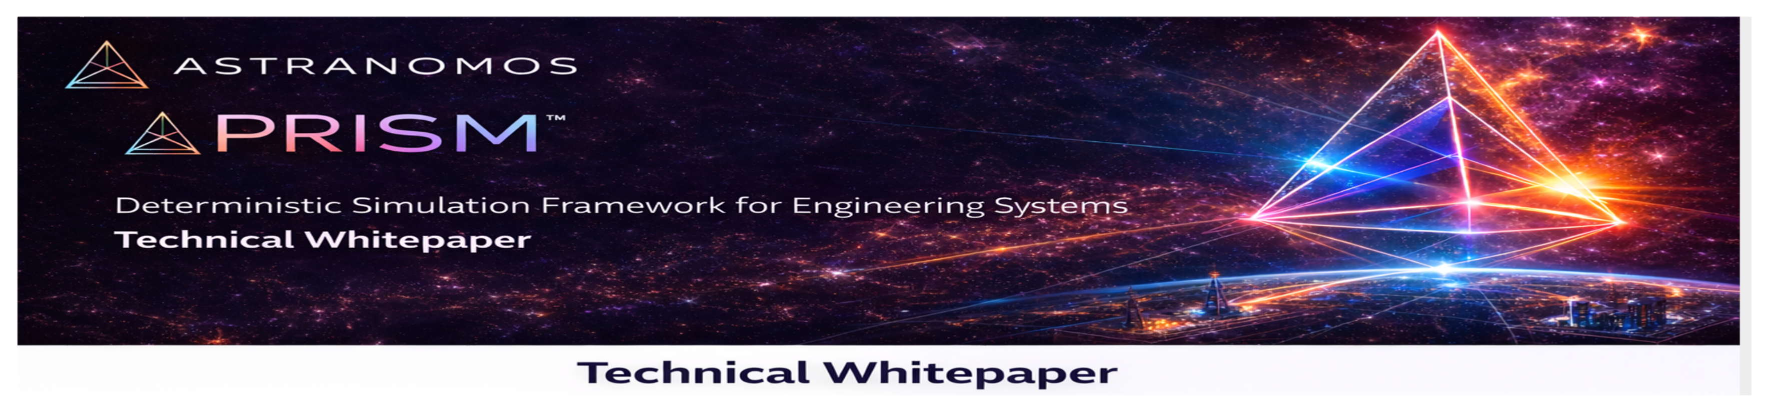

More broadly, the results contribute to an emerging view of high-energy atmospheric motion as a domain in which deterministic structure persists even under extreme conditions. By demonstrating that a fragmented re-entry event can be reconstructed as a coherent, ordered trajectory, this work challenges the assumption that complexity necessarily implies randomness. Instead, it suggests that with sufficient observational resolution and appropriate mathematical representation, the underlying structure of such systems can be revealed and exploited.

In this sense, the present study represents not only an empirical analysis of a specific re-entry event, but also a step toward a more unified approach to modeling motion in aerospace systems. By grounding theoretical constructs in observable data and emphasizing the role of geometry and constraint, it provides a framework that is both physically meaningful and practically applicable. As aerospace technology continues to evolve, the integration of such structured models with observational networks and engineering design processes has the potential to improve efficiency, reduce uncertainty, and advance the state of the art in re-entry prediction and control.