A Methodology for Evolving Aerodynamically Optimal Spacecraft in the Free Molecular Flow Conditions of Very Low Earth Orbits

Julie Rolla,[*] Rick Marcusen,[†] Evan Imata,[‡] Emily Dolson,[§] Max Foreback,[§] Kyle Helson,[¶] Marcin Pilinski,[†] Anselmo Pontes,[‖] Jonathan Sy,[‡] and Joey Wagner[§]

ABSTRACT. — Very low Earth orbit (VLEO) enables improved spatial resolution and access to the lower thermosphere, but sustained operations are limited by strong, highly variable aerodynamic drag in the free molecular flow regime. In addition to drag, spacecraft must also meet numerous other requirements, including mass, volume, and power. Together, these constraints motivate design methods that can explore large parameter spaces with multiple competing objectives. This report presents early work toward an end-to-end workflow for VLEO shape optimization that couples the primitive-based Nebulous evolutionary algorithm with the Vehicle Environment Coupling and TrajectOry Response (VECTOR) aerodynamic simulation software. A key challenge is ensuring that the aerodynamic model and optimization representation use a consistent and effective level of geometric detail. We therefore first ran sensitivity studies in VECTOR to estimate the impact of surface feature size and surface resolution on force coefficient. These results guide the minimum feature scale and mesh/shape resolution used in the Nebulous pipeline, so that design changes can result in evolutionary pressure. We then tested the algorithm by minimizing the force coefficient subject to CubeSat volume limits. The evolution shows a steady

---

[*]Tracking System and Applications Section.

[†]University of Colorado.

[‡]University of California, Berkeley.

[§]Michigan State University.

[¶]University of Maryland Baltimore County, NASA Goddard Space Flight Center.

[‖]Autogenetics Research Lab.

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

improvement in aerodynamic performance, demonstrating that the workflow can generate and evaluate diverse shapes and apply selection pressure effectively. This establishes a baseline for follow-on multi-objective optimization that will explicitly trade drag against functional surface features, stability, power, and other mission constraints.

I. Introduction

Very low Earth orbit (VLEO) offers compelling advantages for both science and other practical applications [1, 2, 3, 4, 5]. By flying close to Earth (below, roughly, 400 km altitude), spacecraft achieve sharper ground resolution and higher signal-to-noise ratios for imaging and remote sensing, while reduced path length and higher link margins can improve latency and data throughput for communications. For aeronomy and space-weather investigations, sustained platforms in this regime enable direct probing of the lower thermosphere and ionosphere, where energy and momentum from the lower atmosphere couple into the space environment. There are also operational benefits, including lower mission and launch costs, and reduced risks of orbital collisions and radiation damage. However, these benefits come with significant challenges: At VLEO altitudes the residual atmosphere produces strong aerodynamic drag that dominates the force environment, drives rapid orbital decay, and tightens constraints on mass, cross-sectional area, and available propulsion. Neutral density also varies strongly with solar and geomagnetic activity, making it difficult to predict lifetime and plan station-keeping. This combination of high potential value and severe drag-driven constraints motivates new approaches to spacecraft design in VLEO.

At VLEO altitudes, the residual atmosphere produces relatively strong aerodynamic drag that dominates other orbital perturbations, driving rapid decay [6]. Neutral density in this region is not constant as it changes with the solar cycle, responds to geomagnetic storms, and varies over local time and latitude. As a result, predicted lifetime, station-keeping requirements, and even collision risk carry large uncertainties that must be absorbed by the design. These environmental constraints interact with a crowded design trade space. Reducing drag often pulls against other system drivers, such as the need for large areas for power generation, thermal control surfaces, clear sensor fields of view, and communications apertures. Standard CubeSat and SmallSat form factors further restrict options, since volume, mass, and rail geometries are fixed by deployment standards. In addition, aerodynamic torques couple the external shape and mass distribution to attitude stability and pointing control. Together, these factors mean that spacecraft geometry and aerodynamic behavior must be treated as first-class design variables, not late-stage refinements applied after other subsystems are fixed.

Current spacecraft design practice is not well matched to the demands of VLEO. Many standard tools and processes assume low-drag regimes or rely on continuum-flow intuition, then apply simple ballistic-coefficient estimates or flat-plate approximations

late in the design process, after the overall geometry has already been fixed by power, payload, and structural constraints [7, 8, 9]. High-fidelity free molecular flow models exist and can predict drag and moments on detailed 3D geometries, but they are typically used to analyze a small number of hand-designed configurations rather than to explore a broad topology design space. As a result, there is no systematic way to *optimize* spacecraft topology—including overall shape, small-scale surface features, and deployable structures—for a high-drag, highly variable environment. At the same time, designers have little guidance on which geometric details must be modeled, such as feature size or surface resolution, to capture drag accurately, particularly for CubeSat-scale spacecraft where small features can be comparable to the characteristic dimensions of the bus.

Evolutionary algorithms (EAs) provide a flexible class of heuristic optimizers that are well suited to high-dimensional, nonlinear design spaces with complex constraints and discrete choices [10, 11]. The Nebulous project has developed an EA framework (ECLIPSE [12]) that it uses to optimize hardware and instruments by coupling three-dimensional (3D) representations with science simulation software to determine cost functions [13, 14, 15]. Each candidate design is evaluated with a high-fidelity "digital twin" simulation, for example, a finite-difference time-domain electromagnetic solver, to compute mission-relevant performance metrics such as gain, bandwidth, or science sensitivity. Nebulous employs a steady-state genetic algorithm with features such as age layering [16] and a variety of mutation and crossover operators to efficiently explore a large topology space while maintaining diversity. Together, prior Nebulous results show that primitive-based representations can span a wide range of geometries without predefining a narrow design family, and that tight coupling between the EA and high-fidelity simulation enables optimization directly against mission-level objectives rather than simple geometric surrogates. While there is some limited exploration of evolutionary methods applied to spacecraft aerodynamics [17], these results were not constrained by spacecraft design and manufacturability parameters.

The challenges of VLEO point to a clear but still unfilled opportunity to optimize with EAs. High-fidelity tools for VLEO aerodynamics can already compute forces and moments on detailed 3D geometries, yet they are almost always used to analyze a small set of hand-designed spacecraft rather than to drive automated topology search. At the same time, EAs like Nebulous can generate and refine complex 3D structures, but they have rarely been applied to spacecraft geometries in the free molecular regime. Bridging this gap raises three concrete questions. First, we need to understand geometric sensitivity: what magnitude and scale of changes—such as adding or removing primitive blocks or small protrusions—produce measurable changes to drag in VLEO. Second, we must determine how finely curved surfaces, including cylinders, booms, and rounded bus corners, must be tessellated to keep drag errors within acceptable bounds. Third, we need to tune the primitive representation and mutation operators so that evolutionary steps correspond to physically meaningful, drag-relevant modifications without making the search intractable.

Answering these questions is requisite for designing effective fitness functions for VLEO (e.g., based on ballistic coefficient, lifetime, or stability), for training surrogate models that approximate high-fidelity drag solvers over the relevant design space, and ultimately for performing full multi-objective optimization of real spacecraft geometries under CubeSat-like constraints.

In this paper, we take early, foundational steps toward using EAs for spacecraft designed to fly in VLEO by focusing on how geometry must be represented and modeled for drag. We begin by describing the VLEO environment and the free molecular aerodynamic modeling framework we use, including the VECTOR plate code and the assumptions adopted for our studies. We then introduce the Nebulous primitive-based representation and discuss how CubeSat-like constraints shape the allowable design space. Next, we present two sets of sensitivity studies: one on the effect of small surface features on drag for simple plates and cubes, and one on the surface resolution required to model curved structures such as cylinders with acceptable accuracy. Finally, we discuss how these results inform the choice of primitive sizes, tessellation strategies, and mutation operators for future Nebulous-style optimization of full spacecraft geometries in VLEO, and outline the next steps toward a full multi-objective design framework.

II. Evolutionary Algorithm

We use the Nebulous EA to explore spacecraft geometries in a high-drag VLEO regime by repeatedly generating candidate designs, evaluating their aerodynamic performance, and using those scores to bias future sampling toward improved solutions. A simplified version of the algorithm is shown in Figure 1. The implementation follows a steady-state genetic algorithm, where new individuals are continually introduced into the population so that successful genetic material becomes immediately available [18]. It is coupled with an age-layered population structure (ALPS) to reduce premature convergence by keeping subpopulations of different ages partially separated [19, 16]. In this study we use three ALPS layers with a population capacity of 100 individuals per layer and a linear age progression rule for promotion between layers, and we execute the full run asynchronously with 100 concurrent evaluations. As soon as an individual finishes evaluation it is inserted into the appropriate layer and a new evaluation is launched to maintain a constant level of parallelism.

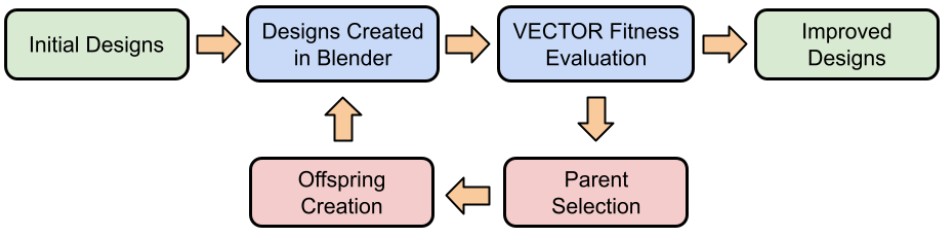

Figure 1. Diagram illustrating a simplified version of the EA loop.

Each candidate spacecraft is encoded using the standard Nebulous constructive representation, where an individual is a rooted tree of connected geometric primitives that "build outward" from an initial seed. Connections in the tree specify both the parent–child connection and the child's position relative to the parent, enabling complex 3D configurations to emerge through incremental growth and modification. To keep the search space aligned with realistic SmallSat architectures, we constrain all generated geometries to fit within a 12U CubeSat bounding volume and enforce a minimum internal cargo volume requirement of 3U; individuals or operations that violate these constraints are regenerated until the constraints are met.

A. Fitness Functions

Fitness is based on the ballistic coefficient, simulated using the VECTOR software, from the aerodynamic force predicted for each geometry under the specified VLEO flow conditions. Individuals are built in Blender [20], and then exported into a format that VECTOR can intake. ballistic coefficient defines how much an an object resists the force of drag, where a higher ballistic coefficient indicates better performance. As mass is undefined in the current representation of Nebulous, a uniform density of 1000 kg/m^3 was assumed to approximate standard CubeSat mass. As it is desirable to maximize ballistic coefficient , we take the reciprocal as the cost function for fitness calculations. Additional details on the fitness are described in Section III.

B. Selection Methods and Genetic Operators

New offspring are generated by selecting parents with tournament selection of size 10 [21]. Selection is performed within the relevant ALPS layer to preserve the intended age stratification, and once an offspring is evaluated, it is inserted back into that same layer based on its age.

Variation is introduced through five genetic operators, each applied with equal probability to produce an offspring from its selected parent. The *grow* operator adds a new primitive by attaching a child to an available face of an existing node, enabling the structure to expand outward in the tree. The *prune* operator removes a leaf node (a primitive with no children), providing a complementary mechanism to simplify geometry and recover feasibility when constraints become tight. The *dimension mutation* operator perturbs one geometric parameter of a selected primitive (e.g., a length or radius) within allowed bounds. Finally the *proportional regeneration* operator replaces a primitive with a different shape whose dimensions are constrained to remain similar in scale to the originals. After any operator is applied, the resulting design is checked against the 12U envelope and minimum cargo-volume constraints; infeasible offspring are rejected and resampled to ensure that the evolutionary search remains within the admissible design space.

III. Fitness Function Details

Spacecraft aerodynamics in VLEO are typically treated in the free molecular flow regime above 180 km altitude. In this regime, standard continuous fluid aerodynamic forces are negligible compared to collisions between gas molecules and the spacecraft surface. Consequently, the flow cannot be treated as a continuous fluid; instead, aerodynamic forces result from the summation of momentum exchanges from individual gas-surface interactions.

When an atmospheric molecule impacts a surface, it may be scattered in an angular distribution with a variety of characteristics [7] or adsorbed (trapped) on the surface before being reemitted. In this work, we use the Vehicle Environment Coupling and TrajectOry Response (VECTOR) software to model these interactions through a gas-surface energy accommodation coefficient, treated here as a constant parameter across the full surface, along with a broad scattering distribution modeled as a cosine function (diffuse scattering). The use of energy accommodation applied globally with values between 0 and 1 combined with a cosine scattering distribution is often referred to as diffuse reflection with incomplete accommodation (DRIA). Additional gas-surface interaction models will be explored in future phases of this project. The energy accommodation coefficient, α, relates the kinetic temperature of incident and reflected molecules to the wall temperature and provides a compact way to represent how much kinetic energy is transferred during reflection. An energy accommodation coefficient of zero indicates that the incident and reflected molecule will have the same energy. Meanwhile, an energy accommodation coefficient of 1 indicates that the molecule leaves the surface with the energy equivalent to the surface thermal energy. The kinetic energy of the scattered molecules can be approximated using a kinetic temperature that assumes energy stored in vibrational and rotational modes is negligible.

$$T_{out} = \frac{m}{3k_b} V_r^2 (1 - \alpha) + \alpha T_w \tag{1}$$

Here, V_r is the magnitude of the spacecraft velocity relative to the atmosphere, or in other words, the free-stream velocity. Meanwhile, T_w is the surface temperature expressed in Kelvin, normally assumed to be approximately 300 K, and m is the molecular mass of the incident molecules. The drag coefficient on each triangular element of the spacecraft geometry mesh is evaluated for species j, using DRIA equations for a flat plate with one side exposed to the flow.

$$\begin{aligned} C_{d,DRIA,j} = {} & \frac{2}{s\sqrt{\pi}} \exp\!\left(-s^2 \sin^2(\theta_{in})\right) \\ & + \frac{\sin(\theta_{in})}{s^2}(1 + 2s^2)\ \mathrm{erf}\!\left(s \sin(\theta_{in})\right) \\ & + \frac{\sqrt{\pi}}{s}\ \sin^2(\theta_{in})\ \sqrt{T_{out}/T_a} \end{aligned} \tag{2}$$

In Equation 2, θ_{in} is the incident angle with respect to the surface normal vector, T_a is the free-stream (atmospheric) temperature, erf() is the error function. s, the speed ratio, is defined as,

$$s = V_r \sqrt{\frac{m}{2k_b T_a}}, \tag{3}$$

For free-stream conditions with multiple species, the species-dependent drag coefficients are computed separately, then combined in a weighted average according to the partial mass density of each species, $n_j m_j$.

$$C_{d,DRIA,pnl} = \frac{\sum C_{d,DRIA,j} n_j m_j}{\sum n_j m_j} \tag{4}$$

The overall force coefficient, $C_d A$, for any shape is the sum of the products of drag coefficient and area of each panel $C_{d,DRIA,pnl} A$.

A. VECTOR

The aerodynamic analysis presented here is based on analytical calculations of drag coefficients applied to individual panels of a spacecraft geometry by the VECTOR simulation software. VECTOR is an open-source code (https://swx-trec.com/vector/) made available by the Laboratory for Atmospheric and Space Physics (LASP) at the University of Colorado, Boulder. The analytical equations used here assume DRIA. The application of DRIA analytical equations is described in a previous section.

Table 1. Simplified VLEO environmental conditions assumed for all aerodynamic evaluations in VECTOR.

Parameter	Value
Atmospheric temperature, T_a	1200.5 K
Free-stream speed, V_r	7800.45 m/s
Number density, n_{O}	10^{11} m^{-3}
Number density, $n_{\mathrm{O_2}}$	0 m^{-3}
Number density, $n_{\mathrm{N_2}}$	0 m^{-3}
Number density, n_{He}	0 m^{-3}
Number density, n_{H}	0 m^{-3}
Energy accommodation coefficient, α	0.93

While simplified plate models, which simulate and then sum the effects of individual surfaces, are computationally efficient, they fail to account for self-shadowing, and multiple reflections, which are critical when analyzing concave geometries. These limitations, particularly self-shadowing, will be addressed in later phases of research.

Given that the EA may generate complex shapes, the VECTOR software may introduce errors as high as 30% due to self-shadowing effects.

VLEO spacecraft operate in the thermosphere, where temperatures are high and neutral composition varies with altitude, geomagnetic activity, and solar activity. For the present study, the aerodynamic simulations are intentionally simplified to isolate geometry-driven trends in drag. Table 1 shows the single set of fixed environmental (free-stream) conditions in VECTOR that all individuals were evaluated under.

B. Drag Metrics and Fitness Formulation

The drag force acting on a spacecraft traveling through an atmosphere is commonly written as

$$F_d = \frac{1}{2}\,\rho\,V_r^2\,C_d\,A, \tag{5}$$

where F_d is the drag force, ρ is atmospheric mass density, v is the free-stream speed, C_d is the drag coefficient, and A is the cross-sectional area projected normal to the flow. Lower drag force generally reduces the propulsive effort required for altitude maintenance. From a topology-optimization perspective, the two geometry-dependent quantities are C_d and A.

To simplify discussion of the geometry-dependent drag response, we define the force coefficient

$$F_c = C_d A, \tag{6}$$

which characterizes the drag-producing capability of a geometry in a given free-stream independent of ρ and v.

The ballistic coefficient is found by dividing the spacecraft mass, m, by the force coefficient:

$$BC = \frac{m}{C_d A} = \frac{m}{F_c}, \tag{7}$$

A higher ballistic coefficient indicates reduced sensitivity to drag for a given mass. We approximate mass using a uniform bulk density of 1000 kg/m^3 and the enclosed solid volume of the evolved geometry,

$$m \approx \rho_s V, \qquad \rho_s = 1000 \text{ kg/m}^3, \tag{8}$$

where V is the geometry volume.

Nebulous is currently configured as a minimization algorithm. Because it is desirable to maximize BC, we evaluate the reciprocal ballistic coefficient as the fitness (cost) to be minimized:

$$J_{BC} = BC^{-1} = \frac{C_d A}{m} = \frac{F_c}{m}. \tag{9}$$

Both F_c and BC are attitude-dependent: The same geometry can exhibit substantially different drag response depending on which face (or combination of

faces) is presented to the free-stream. To capture both nominal and off-nominal attitudes, two related fitness functions are considered. The first evaluates performance only at the primary (nominal) attitude:

$$J_{\text{nom}} = BC^{-1}(\theta_0, \beta_0), \tag{10}$$

where θ denotes pitch and β denotes sideslip, and (θ_0, β_0) defines the primary attitude.

The second fitness function incorporates robustness to attitude offsets by combining the nominal attitude with a local sweep spanning $\pm 30^\circ$ in both pitch and sideslip about the primary attitude. Let $\mathcal{A}$ denote the set of swept attitudes (excluding the primary attitude), and let $\langle \cdot \rangle_{\mathcal{A}}$ denote the mean over that set. The sweep-augmented fitness is then

$$J_{\text{sweep}} = \frac{1}{2}\, BC^{-1}(\theta_0, \beta_0) \; + \; \left\langle BC^{-1}(\theta, \beta) \right\rangle_{\mathcal{A}}. \tag{11}$$

In this form, the cost penalizes designs that perform well only at a single attitude while encouraging consistently low drag response across nearby orientations.

IV. Sensitivity Analysis

There is limited guidance in the field on the level of geometric detail required to accurately model drag forces acting on a satellite in VLEO. Within the context of this work, having knowledge about geometric sensitivity in aerodynamic modeling will motivate the type of representation that Nebulous can use to evolve a spacecraft design and inform the level of mutation required to produce a noticeable deviation in performance.

The goals of this preliminary study therefore are twofold. Objective 1 is to quantify what magnitude of change to an object's geometry has a meaningful impact on its drag performance. Objective 2 is to determine what level of model resolution is required to accurately represent a spacecraft's aerodynamics in VLEO. This study takes a heuristic approach to addressing these goals, first by evaluating what scale of features on the surface of test geometries significantly alters their drag forces, and then by seeking to quantify the level of resolution needed to accurately model the aerodynamics of curved surfaces by evaluating tessellated cylinder and sphere models of varying fidelity.

While this preliminary study was conducted primarily to inform the current Nebulous representations for spacecraft, it is hoped that these results may also offer some guidance to future researchers about the level of detail required in their own satellite geometries to accurately model drag in VLEO.

A. Objective 1: Surface Features

For the purposes of this initial assessment, we define a surface feature as any geometry that exceeds the main envelope of the spacecraft body, including those that drastically

alter the profile, such as deployable solar panels, booms, or origami structures. These types of features, often requiring a deployment mechanism in space, can be comparable in size to the stowed spacecraft.

In extreme cases, these structures can be significantly larger than the launch configuration of a vehicle. The drag impact of these large deployables can dominate the overall aerodynamic performance of the satellite and therefore must be modeled. The effects of features like whip antennas, which may be long but have a minor impact on the overall cross-sectional area, is not well quantified in the literature. The impact of these kinds of features warrants attention. However, any features on this scale would likely be driven by non-drag-related mission requirements, such as communications or power generation, that are currently not accounted for in Nebulous's spacecraft representation. Therefore, investigating these features remains a future task.

This assessment will focus on relatively small-scale structures on the surface of a satellite body, representing components such as external sensors or patch antennas. These types of features are most accurately modeled by the current primitive-based Nebulous representation and will inform Objective 1 by identifying the minimum size of features that impact aerodynamic performance.

Approaching the surface feature question in this way addresses two problems. In the context of this work, Nebulous's primitive-based representation may choose to append an elementary shape to or remove it from the surface of a previous generation of satellite as a mutation. The size of these basic elements can vary within a predefined range. This study will inform the minimum allowable size of element to be used, so that adding or removing features from the spacecraft geometry will affect drag performance and influence the spacecraft evolution.

More broadly in the satellite community, a drag evaluation like this would likely be a secondary concern to developing the vehicle to perform a scientific or technological activity, which in turn would drive most design requirements. An analysis from this perspective would likely involve starting with a detailed model of a vehicle, then crafting a lower-fidelity model by removing small features that do not have a significant contribution. From this perspective, this study will inform what scale of features can be neglected without significantly altering the satellite's aerodynamics.

1. Surface Feature Heuristic Study: Plate

In order to evaluate the contribution that a population of surface features of similar size has on a parent geometry, a subtractive approach was taken. Geometry models were created with many small-scale surface features, divided into populations based on their size. The drag performance of each of these models is evaluated in VECTOR across an attitude sweep, and F_c will be evaluated. In order of increasing size, populations of features are incrementally removed, and the model is reevaluated. The variation in aerodynamic performance of the reduced model is calculated as a relative

percentage change in F_c compared to the fully detailed model. A threshold of 5% change was decided on for a population of features to be considered significant. This analysis assumes the same environmental conditions in VECTOR as defined in Table 1.

The first test model evaluated in this study was a plate, similar to a drag evaluation performed in [8]. The constructed plate is 1 m square, extruded to a depth of 1 cm for simplicity. On the front face of the plate is a 10 x 10 grid of 5 cm squares. Each row of squares was extruded to the following depths, in decreasing magnitude (5 cm, 4 cm, 3 cm, 2 cm, 1.5 cm, 1 cm, 0.5 cm, 0.4 cm, 0.3 cm, 0.2 cm). An image of the plate can be seen in Figure 2.

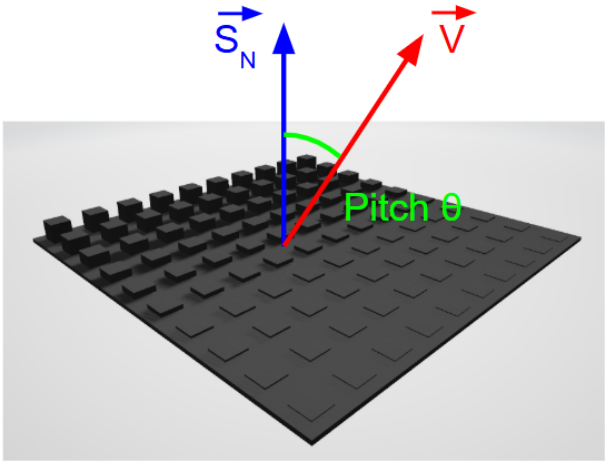

Figure 2. Plate geometry with all surface feature populations. $\tilde{\mathbf{S}}_\mathbf{N}$ represents the surface normal vector, $\tilde{\mathbf{V}}$ is the evaluation velocity vector, and θ is the pitch angle.

The force coefficient of this plate is computed over a one axis pitch sweep between $-90°$ and $90°$ in increments of $5°$. Figure 3 contains the force coefficient outputs of the evaluation in VECTOR and the relative performance.

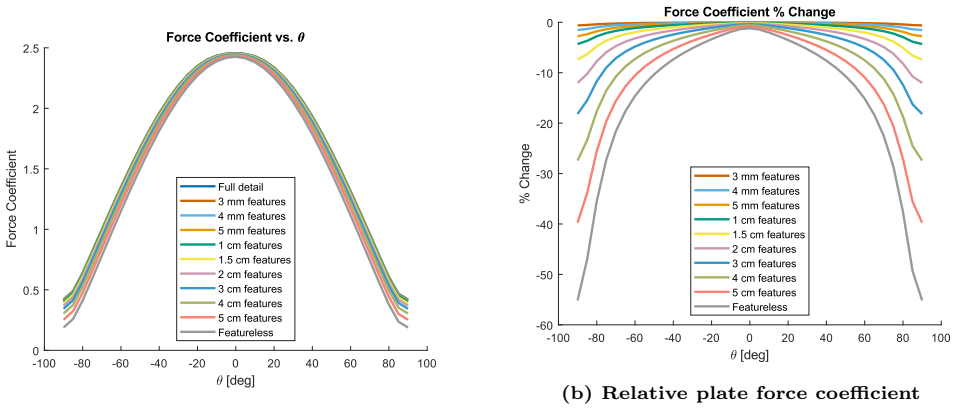

(a) Plate force coefficient evaluation.

(b) Relative plate force coefficient performance.

Figure 3. Drag evaluation for plate geometry shown in Figure 2. Legend indicates the smallest population of features present on the test geometry.

The plate performance in Figure 3 shows F_c being minimized for high pitch angles where the edge of the plate faces into the flow. Figure 3b shows that for the plate geometry, surface features have the greatest impact at high pitch attitudes, and relatively little impact at low pitch. These results also show that the removal of surface features results in a persistent underestimation of force coefficient.

Figure 3b shows that the threshold of 5% change in F_c is reached when $5\,\mathrm{cm}^2$ square features extruded to heights less than 1 cm are removed. However, this is only the case at extreme pitch angles, and the overall impact of each population of features reduces significantly as pitch decreases.

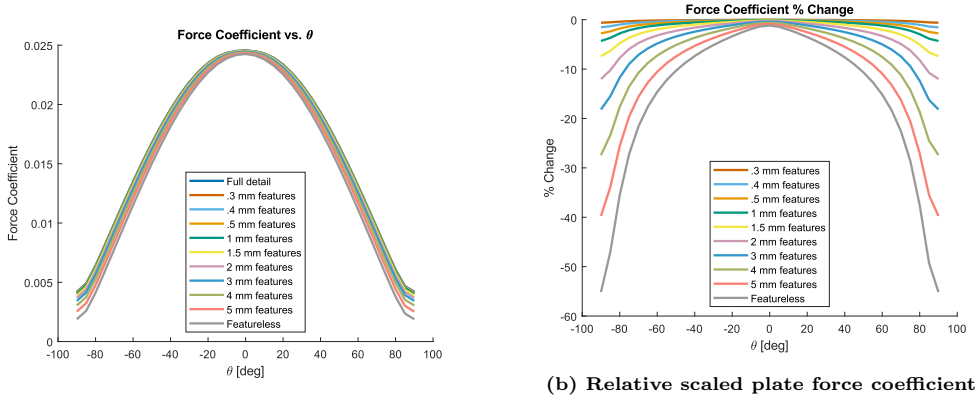

(a) Scaled plate force coefficient evaluation.

(b) Relative scaled plate force coefficient performance.

Figure 4. Drag evaluation for plate geometry shown in Figure 2, scaled to 10% of its original size. Legend indicates the smallest population of features present on the test geometry.

To further evaluate this result, the same plate model contained in Figure 2 was scaled to 10% of its original size and reevaluated under the same conditions to determine if these trends are relative. Figure 4 contains the outputs of the reduced evaluation.

From a high level, the plots in Figure 4 demonstrate the same trends as the full-scale evaluation. Figure 4a shows that at 10% of the original scale, the plate's force coefficient scales at 1% of the original. Figure 4b shows that $5\,\mathrm{mm}^2$ square features taller than 1 mm on the scaled plate are significant to the total performance at the same 5% threshold. Therefore, by comparing these results it is accurate to say that the impact of any population of surface features depends on their relative size to the spacecraft parent body.

2. Surface Feature Heuristic Study: Cube

While useful for evaluating trends, the plate geometry is not a particularly accurate representation of a satellite that might operate in VLEO. Therefore, a new 1 m cube geometry was created, featuring the same surface feature pattern as the full-scale plate, but now including those populations on all six sides of the structure. An image of the cube geometry is shown in Figure 5.

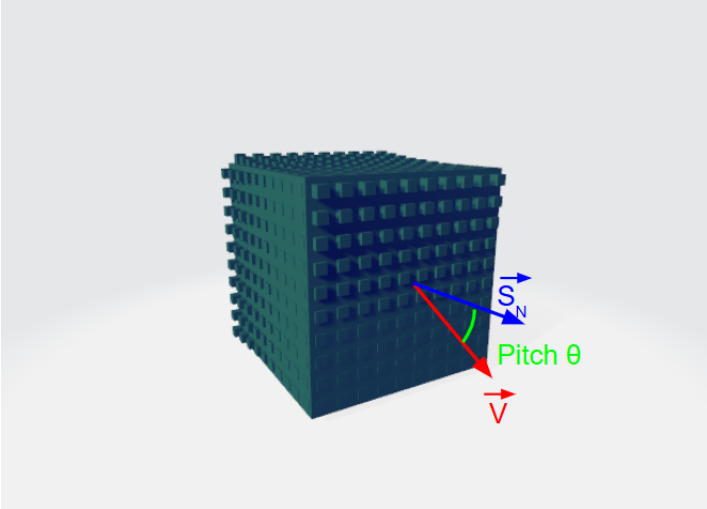

Figure 5. Cube geometry with all surface feature populations. $\tilde{\mathbf{S}}_\mathbf{N}$ represents the surface normal vector, $\tilde{\mathbf{V}}$ is the evaluation velocity vector, and θ is the pitch angle.

The cube is evaluated over a one axis pitch sweep between $-90°$ and $90°$ in increments of $5°$ as shown in Figure 5. As with the plate, populations of features are removed incrementally to evaluate their relative impact. Note that the removal of a population of features removes them from all six sides of the structure.

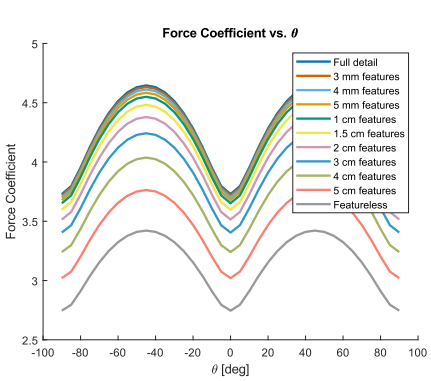

(a) Cube force coefficient evaluation.

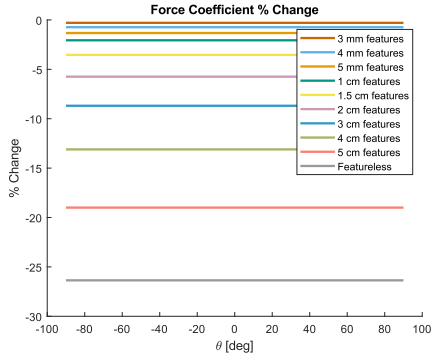

(b) Relative cube force coefficient performance.

Figure 6. Drag evaluation for cube geometry shown in Figure 5. Legend indicates the smallest population of features present on the test geometry.

The results of the cube evaluation are shown in Figure 6. When compared to the plate results, Figure 6a shows periodic trends that align with the rotation of the cube and a much larger dependence on surface features. Despite this variation, however, Figure 6b indicates that the relative change in F_c is independent of pitch. It can also be seen that when compared to the 1 m plate, larger features fall under the 5% change threshold. Linearly interpolating between feature sizes in Figure 6a, it can be determined that for the 1 m cubic structure, populations of features 1.8 cm tall and larger have a notable impact on the overall aerodynamics. Having seen that the

impact of these features is relative to the size of the parent structure, it can therefore be restated that square features 5% of a cube's side length extruded to a height of 1.8% of a side length are significant when modeling drag forces, in comparison to features with heights over 1% with the plate.

However, there are some restrictions that should be acknowledged when interpreting these results, due to limitations with the current drag evaluation method. VECTOR utilizes a panel approach to compute F_c, as discussed in Section III. While this method results in mostly accurate drag metrics, the current implementation of the program does not account for "shadowing" effects on the structure. This effect can be seen in the feature cube projected area output by VECTOR, contained in Figure 7. For highly detailed models such as the feature cube, VECTOR accounts for the cumulative contribution for all surfaces normal to the velocity, regardless of whether they are covered by features in the relative foreground. Physically, this shadowing effect would prevent some particle interactions between features and the incoming flow, and with the analytical method in VECTOR this results in nonphysical projected areas and force coefficients that must be acknowledged.

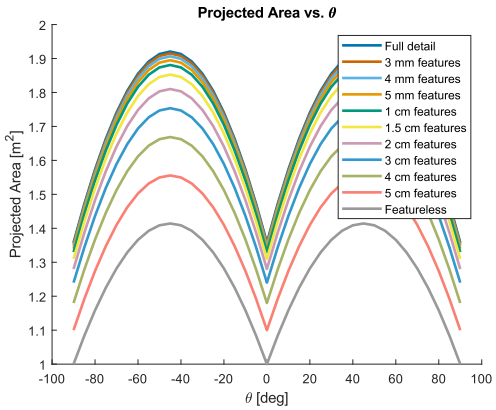

Figure 7. Nonphysical feature cube projected area computed with VECTOR.

To reduce the impact of this shadowing effect on the analysis, an additional cube model was developed. As before, a 10 x 10 grid of surface features was created on each face, but instead of features with constant area and variable heights, each feature on this new geometry is perfectly cubic. This model uses the same base 1 m cube structure, and the cubic features follow the same dimensions as the previous analysis (5 cm, 4 cm, 3 cm, 2 cm, 1.5 cm, 1 cm, 0.5 cm, 0.4 cm, 0.3 cm, 0.2 cm). This new cubic feature model is shown in Figure 8 and evaluated under the same conditions as the original cube.

The simulated results in Figure 9 are broadly similar to those found in Figure 6. However, with these new cubic surface features, the impact of each population is smaller relative to its size. From this new analysis, it can be concluded that for a cubic structure, cubic features larger than 5% the scale of the total body produce a significant impact on overall drag performance.

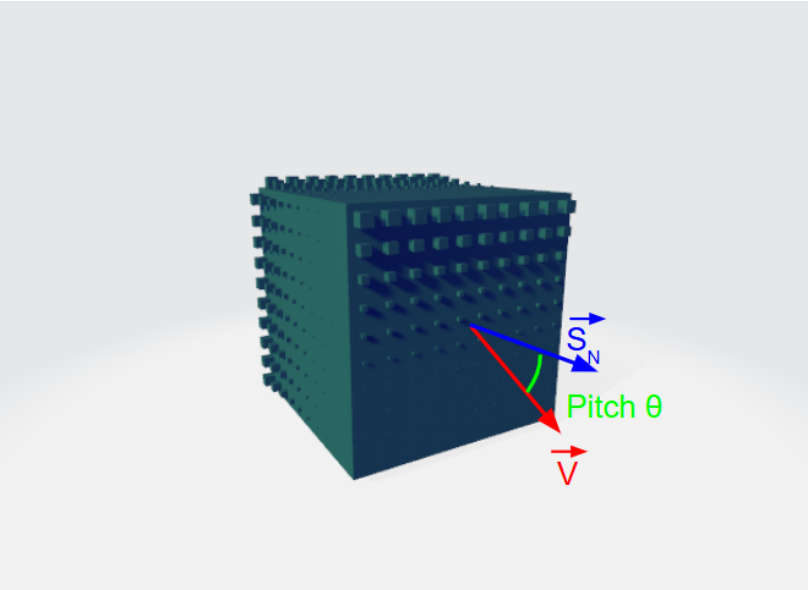

Figure 8. New feature cube with cubic surface features. $\tilde{S}_N$ represents the surface normal vector, $\tilde{V}$ is the evaluation velocity vector, and θ is the pitch angle.

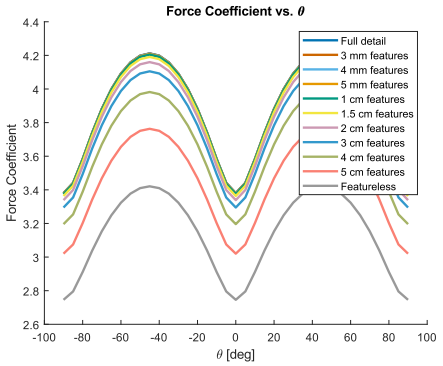

(a) Cube with cubic features force coefficient evaluation.

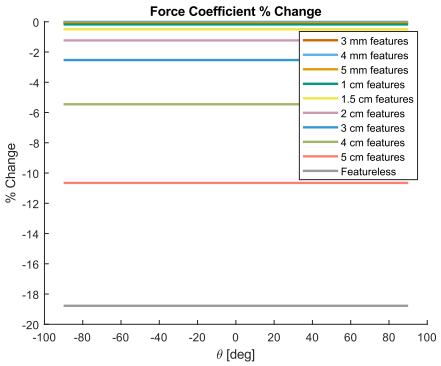

(b) Cube with cubic features relative force coefficient performance.

Figure 9. Drag evaluation for cube geometry shown in Figure 8. Legend indicates the smallest population of features present on the test geometry.

3. Surface Feature Heuristic Study: Discussion

With the current evolution in Nebulous, CubeSat standards are being used as a comparison for overall volume constants, with a 1U CubeSat being roughly 10 cm x 10 cm x 11 cm. Therefore, based on this preliminary analysis, the smallest primitive size allowed for Nebulous's primitive representation should be a 5 mm cube, or equivalent for other primitive shapes, so that a population of features of this scale can tangibly impact the spacecraft's aerodynamic performance. It should also be restated that the removal of any population of features results in an underestimation of force coefficient and thus a bias toward improved aerodynamic performance in VLEO.

While these initial results can inform the current representation choices for Nebulous, this analysis is not comprehensive. As future work, other primitive types, such as cylinders and spheres, should be evaluated in a similar manner, and more comprehensive attitude sweeps should be incorporated. Furthermore, it is intended that in future work a similar analysis will be performed with geometry from more representative CubeSat geometries, such as a 3U with realistic sparsity of features, to validate the conclusions drawn from this basic analysis.

B. Objective 2: Surface Resolution

An additional question when evaluating low-fidelity drag models is what level of surface resolution is required to model curved structures in free molecular flow. How coarse a curved mesh is will both dictate how realistic of a representation can be achieved for a curved geometry and dictate how computationally expensive a drag analysis is. For Nebulous, both of these factors need to be considered, so there is a significant incentive to utilize the lowest-resolution geometries necessary to accurately represent a surface.

1. Cylinder Tessellation

As a preliminary approach to answering these questions, VECTOR was used to evaluate the drag profiles of various extruded polygons roughly representing a cylinder with higher and higher resolution. Figure 10 depicts each of the models created for this test. Polygons defined by their number of sides (N) were created, with a maximum dimensions of 1 m. To simulate a cylinder, each of these polygons were extruded to a depth of 1 m, such that the maximum cross-sectional area for each cylinder is 1 m^2. Note that while Figure 10 contains a model of a cylinder, resolution limitations of 3D modeling software still tessellate this cylinder to some extent. Therefore, the ideal cylinder option in VECTOR, which computes the exact solution to the drag of a cylinder under DRIA assumptions, was used as reference rather than a high-resolution cylindrical geometry. Each cylinder was rotated about its primary axis from angles of 0° to 180° in increments of 1° to account for symmetry. This rotation angle is defined as the pitch angle, or θ, as before.

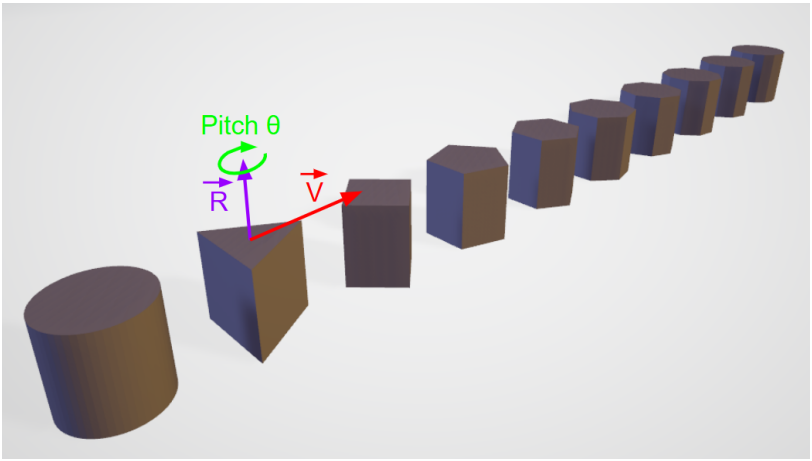

Figure 10. Cylinder tessellation test models. With the N=3 geometry as an example, vector $\tilde{\mathbf{V}}$ represents the nominal velocity vector for the shapes, $\tilde{\mathbf{R}}$ is the rotation vector, and θ is the rotation angle.

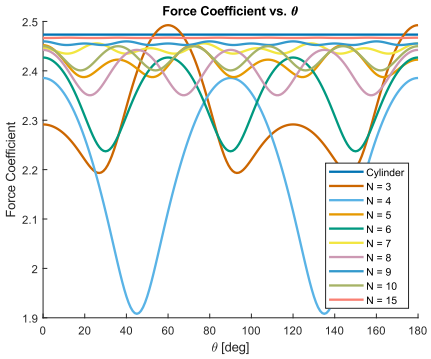

(a) Cylinder tessellation force coefficients

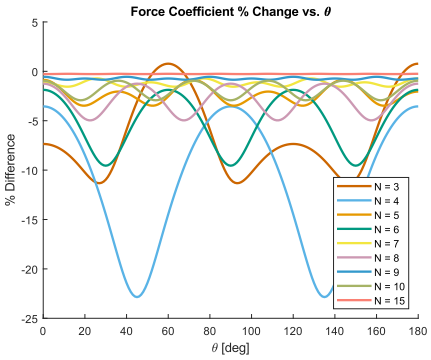

(b) Cylinder tessellation force coefficient change relative to a an actual cylinder.

Figure 11. Drag evaluation for cylinder tessellation geometries shown in Figure 10. Legend N number indicates the number of sides.

From Figure 11, it can be seen that within 5% variation of force coefficient, an eight-sided shape can roughly be used to approximate a cylinder. This corresponds to a mesh size of around 76% of the approximated radius of curvature and would likely be the simplest applied method for constructing a low-fidelity cylinder due to its symmetry. However, the performance of any shape with more than seven sides appears to be sufficient in approximating a cylinder, and the minimum faces to consistently stay below 5% is $N = 5$, though improvements are not consistent as the number of faces increases. For example, the $N = 7$ cylinder has better performance than the $N = 8$ cylinder. Consistently, objects with an odd number of faces tend to perform better than their $N + 1$ counterparts. Understanding why this occurs warrants future investigation.

Figure 11a also shows that that for the $N = 3$ geometry (triangular prism), the force coefficient is enhanced over that of the ideal cylinder at angles where a flat side faces directly into the flow. This is most likely due to the large surface area that is facing into the flow maximizing the overall possible drag coefficient. A flat face with its surface normal aligned with the flow will have a large drag coefficient while also being the only surface that is interacting with the flow. In other words, the C_d of a flat plate in this configuration is always larger than that of a cylinder. All other cylinder approximations at all attitudes produce an underestimation in force coefficient. As an additional observation from Figure 11b, the $N = 4$ geometry (rectangular prism), has the worst overall performance in approximating a cylinder, with error over 20% at certain attitudes.

2. Sphere Tessellation

Following the cylinder study, a similar evaluation was performed with spherical geometries. As seen in Figure 12, six spherical models were created, each with a different number of triangular faces that create the mesh.

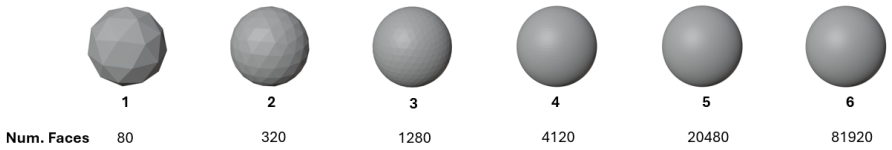

Figure 12. Sphere tessellation geometries.

Each sphere in Figure 12 is designed to approximate an ideal sphere with a diameter of 1.212 m. However, due to the complex reticulation of these spherical approximations, these models were not normalized in maximum cross-sectional area, meaning both C_d and A need to be considered when interpreting F_c. As an initial evaluation for a single attitude, Figure 13 shows how increasing the number of faces on the sphere allows the model to converge on the performance of an ideal sphere. To determine attitude dependence, these geometries were evaluated in VECTOR over a single axis rotation from $-90°$ to $90°$ sweep in pitch angle, in increments of $5°$.

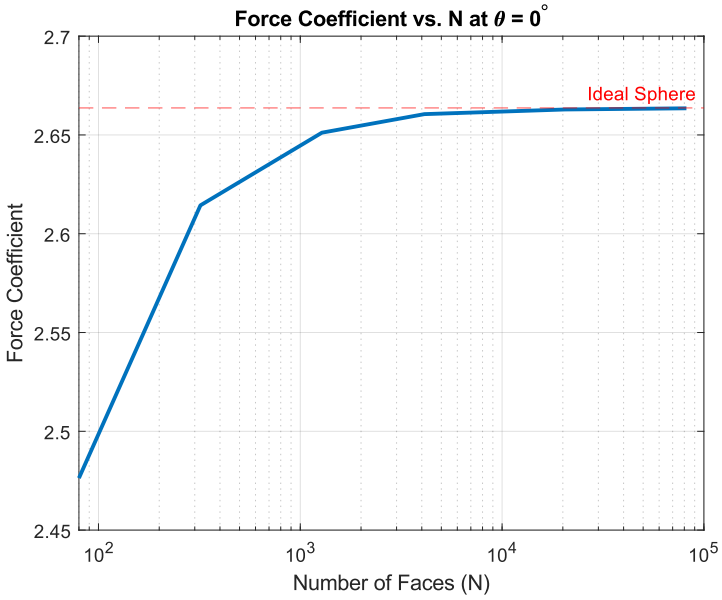

Figure 13. Sphere $\mathbf{F_c}$ vs. N.

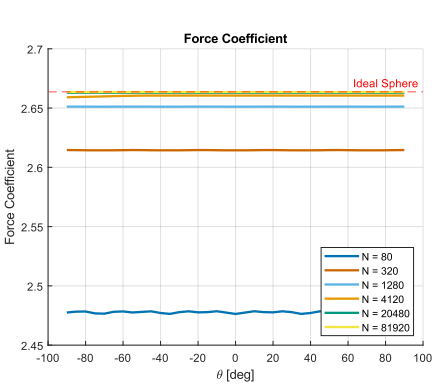

(a) Sphere force coefficient vs. pitch angle θ

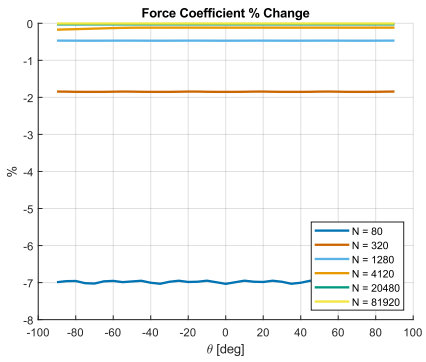

(b) Percentage change of spherical force coefficient vs. pitch angle θ

Figure 14. Drag evaluation for spherical tessellation geometries shown in Figure 12. Legend N number indicates the number of faces on the sphere.

Performance is overall consistent across the pitch angle sweep, as expected from a sphere. However, Figure 13 shows that there is still a noticeable amount of variation in the calculated F_c for the $N = 80$ sphere, which implies some level of attitude dependence. This behavior is less prevalent at higher resolutions. Again using a >5% change in force coefficient as our threshold, it can be determined from Figure 14b that the ideal number of faces to approximate a sphere must lie between 80 and 320.

3. Surface Resolution Heuristic Study: Discussion

From the cylinder analysis, it was expected that a maximum element length corresponding to roughly 76% of the approximated radius of curvature would provide an acceptable level of accuracy. However, in the sphere study the $N = 80$ shape has a maximum element length of 0.375 m, corresponding to 61.8% of the approximated radius. With a finer surface resolution, better performance should be expected. However, with this element size the $N = 80$ shape consistently has around -7% error compared to the ideal sphere, above the threshold. This implies a level of nuance in selecting a mesh size to approximate more complicated curved geometries that requires further investigation.

Future work on this topic will include identifying the ideal number of faces to approximate a sphere within 5% error, investigating rotational dependence in both pitch and sideslip angles, and further investigating the performance variability as a result of inconsistent areas between the sphere approximations.

V. Results

We applied the primitive-based Nebulous representation to reduce aerodynamic drag for spacecraft operating in VLEO, and integrated the VECTOR aerodynamic model into the Nebulous pipeline to evaluate each candidate geometry's fitness.

A. Evolution with Nebulous

To test end-to-end feasibility, we initialized the evolutionary process from a population of randomly generated geometries that satisfy the representation and volume constraints described in Section II. The parameters of the run are provided in Table 5 in the Appendix. Figure 15 shows the fitness of evaluated individuals over the course of a representative run. Despite starting from unstructured random assemblies, the population exhibits a clear trend toward improved performance: As the run progresses, newly generated individuals achieve increasingly higher ballistic coefficient (equivalently, lower force coefficient). This behavior indicates that (i) the primitive-based encoding spans aerodynamically distinct shapes and (ii) the coupled Nebulous+VECTOR pipeline provides sufficient selection pressure to drive measurable drag improvements over time.

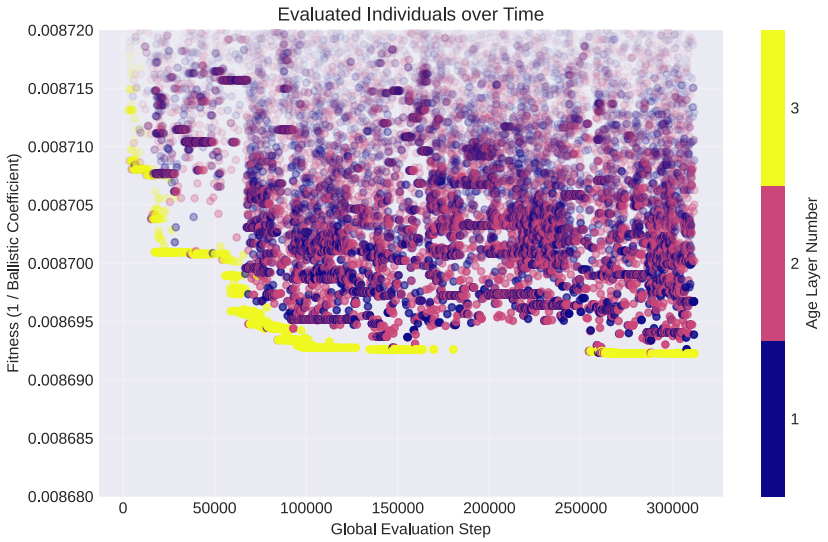

Figure 15. Fitness of population over time.

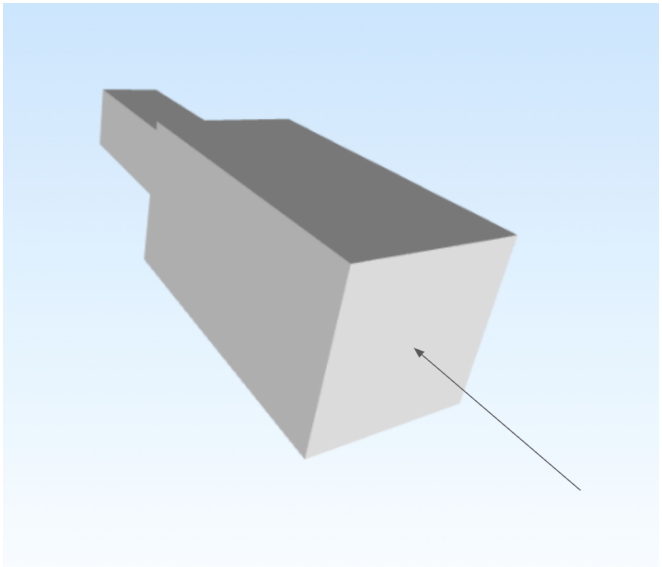

Figure 16. Individual with the best fitness (ID 309124). The flow of particles is indicated by the arrow.

The best individual from one of these evolutionary runs is shown in Figure 16. Qualitatively, the geometry is composed of two separate primitives (i.e., two disconnected components in the encoded design). While this configuration is permitted under the current representation rules, it exposes an interface issue between the representation and evaluator. When the two primitives are passed to VECTOR as separate meshes, the resulting total ballistic coefficient is lower than the value obtained when those components are combined into a single, watertight mesh. This difference is driven by how disconnected components are treated during evaluation, including the possibility of additional exposed surface area (e.g., coincident or opposing faces that would be internal after merging), which biases the drag estimate and can distort selection.

B. Primitive-to-Mesh Conversion and Its Impact on Evolvable Complexity

VECTOR requires geometry inputs as surface meshes, whereas the current Nebulous spacecraft representation encodes designs as sets of analytic primitives. Converting a set of primitives into a single high-quality mesh suitable for aerodynamic evaluation is not trivial, and we found this conversion step to be a primary limiter on the geometric complexity that could be evolved reliably in the present study.

A major difficulty occurs when primitives touch or nearly touch. Depending on the configuration, two primitives may produce coplanar faces, share an edge, or intersect at only a single point. These degeneracies are common outcomes of mutation and crossover in a block/primitive construction grammar, but they are challenging for boolean union and meshing operations. In our experiments, these cases frequently produced meshing artifacts (e.g., nonmanifold edges, sliver triangles, self-intersections, or holes), which can change the effective area and surface orientation distribution seen by VECTOR. As a result, the optimization is implicitly biased toward simple assemblies that remain easy to mesh and evaluate, rather than toward the full range of topologies that the primitive grammar could otherwise express.

To mitigate these issues, we developed a separate representation that directly encodes each candidate design as a single 3D mesh, eliminating the need for downstream mesh union prior to VECTOR evaluation. This mesh-native representation is described in Section VII and is intended to preserve Nebulous' ability to explore diverse geometry while providing consistent, watertight inputs to the aerodynamic evaluator across the full range of candidate designs.

C. Pipeline Integration and Computational Performance

These experiments demonstrate that we have produced a functioning, end-to-end evolutionary optimization pipeline for VLEO aerodynamics. The pipeline cleanly separates (i) representation and genetic operators, (ii) simulation-based evaluation (VECTOR), and (iii) run control and population management. This modularity

enables straightforward substitution of representations (as in Section VII) and supports integration of additional simulation tools for future optimization objectives. In addition, we implemented and validated several computational optimizations across the evaluation workflow. Collectively, these changes reduced overall computation time by a factor of 13 for the same evaluation workload, improving throughput and making larger evolutionary studies feasible on available computing resources.

VI. Discussion

VLEO environments present unique challenges for sustainable spacecraft design, where aerodynamic drag can dominate lifetime and operations for CubeSats and SmallSats. The work presented here provides initial guidance on (i) what surface feature scales must be resolved to produce meaningful changes in modeled drag and (ii) what mesh fidelity is needed to represent curved surfaces with acceptable error in force coefficient.

Our surface-feature sensitivity studies indicate that the minimum effective feature size depends on the parent geometry scale. For flat plates, the smallest features that produce a $\sim$5% change in force coefficient occur primarily at extreme attitudes and shift with the overall plate dimensions. For cubic bodies, surface features have a stronger and more uniform effect across attitude, and we find that features on the order of 5% of a cube's side length are significant for the force-coefficient threshold used here. These heuristics are directly relevant to evolutionary optimization: They inform minimum primitive sizes and mutation magnitudes so that geometric edits are large enough to be selectable, while avoiding overly fine features that primarily add complexity without aerodynamic benefit.

Mesh resolution is likewise a controlling factor in simulation accuracy and runtime. In our curved-surface resolution studies, we compared a true cylinder to polygonal prism approximations and found that an octagonal prism yields an approximately 5% variation in C_D relative to a cylinder. Interpreted as a geometric guideline, this corresponds to representing curvature with a facet scale on the order of the radius of curvature (within a factor of order unity) for the error levels considered here. Our analogous sphere tessellation study showed somewhat larger error at the same nominal resolution fraction, reinforcing that ideal mesh resolution is geometry-dependent and may need to be spatially variable, similar to how finite element analysis meshes are refined near stress risers while remaining coarser elsewhere.

Finally, we note that EAs can exploit evaluator-specific artifacts and nonphysical configurations. In this work, two such examples were evident: (i) limitations associated with primitive-to-mesh conversion (including disconnected components and meshing degeneracies) and (ii) systematic effects tied to the aerodynamic evaluation approach (e.g., neglected shadowing). These considerations motivate mesh-native representations and continued effort to ensure that evaluator assumptions and geometric encodings remain aligned as evolvable complexity increases.

VII. Future Work

Given the limitations of the primitive-based representation (Section 5), we developed a mesh-native representation better suited to drag optimization with VECTOR. In this model, the external geometry of the spacecraft is represented directly as a single 3D mesh, eliminating primitive union and reducing evaluator sensitivity to disconnected components and meshing artifacts.

For computational efficiency and operator simplicity, we adopted a lightweight representation based on point clouds that define a surface that can be triangulated into a watertight mesh for evaluation. In this representation, mutations act as local perturbations of point positions (introducing bumps or dimples that can accumulate over generations), while recombination exchanges contiguous regions between parent individuals to combine larger-scale features evolved in different lineages. Key configuration parameters, including mutation magnitude, point density (mesh resolution), and population settings, must be selected together to balance selection pressure, convergence rate, and diversity.

The surface-feature and curved-surface studies described earlier provide practical guidance for initial parameter choices. In particular, they bound the geometric scales required to exceed the 5% threshold used to define aerodynamically significant changes. Based on these results, we are conducting integration experiments using initial parameter values drawn from the indicated ranges (Figure 17). Future work will refine these parameters systematically, expand the design classes considered, and improve consistency between the evaluator physics and the geometric degrees of freedom (including potential treatment of aerodynamic shadowing, if supported by the evaluation model).

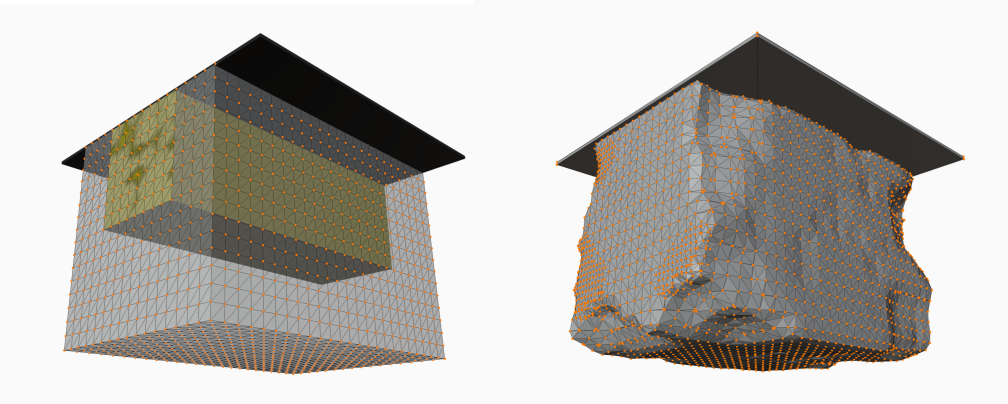

Figure 17. Point-cloud representation of an evolvable fairing surrounding a 3U CubeSat, constrained to a maximum launch volume of 12U. On the left a seeding design. On the right, a design undergoing evolution.

Future work will also extend the current drag-focused studies to a multi-objective formulation that better reflects real spacecraft design trade spaces. In addition to minimizing drag, relevant objectives include attitude stability and control,

maneuverability, power generation and thermal margins, packaging within CubeSat-class volume constraints, manufacturability and cost, and operational considerations such as robust end-of-life deorbit behavior. The Nebulous framework is already structured to support this shift via implementation of the nondominated sorting genetic algorithm II (NSGA-II), enabling the evolution of a Pareto front of nondominated designs rather than a single "best" geometry, and allowing direct exploration of tradeoffs across competing objectives. In parallel, we will pursue additional computational speedups through approaches such as surrogate-assisted or multi-fidelity evaluation, caching and reuse of repeated drag calculations, and adaptive allocation of simulation effort to the most informative candidates.

Exploring deployable components (e.g., solar panels, instruments, or one-time shape modifications for aerodynamics) and incorporating attitude dynamics and additional constraints are further directions that can expand the realism and applicability of evolved VLEO spacecraft concepts.

Acknowledgments

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). The authors are grateful to the entire GENETIS collaboration for their prior work and continued assistance, especially Prof. Wolfgang Banzhaf. We would also like to thank Remcom for their support, especially Tarun Chawla, Walter Janusz, and Benjamin Hardy.

References

[1] N. Crisp, P. Roberts, S. Livadiotti, V. Oiko, S. Edmondson, S. Haigh, C. Huyton, L. Sinpetru, K. Smith, S. Worrall, J. Becedas, R. Domínguez, D. González, V. Hanessian, A. Mølgaard, J. Nielsen, M. Bisgaard, Y.-A. Chan, S. Fasoulas, G. Herdrich, F. Romano, C. Traub, D. García-Almiñana, S. Rodríguez-Donaire, M. Sureda, D. Kataria, R. Outlaw, B. Belkouchi, A. Conte, J. Perez, R. Villain, B. Heißerer, and A. Schwalber, "The benefits of very low earth orbit for earth observation missions," *Progress in Aerospace Sciences*, vol. 117, p. 100619, 2020. https://www.sciencedirect.com/science/article/pii/S0376042120300312

[2] R. Nyamukondinawa, W. Peeters, and S. Udayakumar, "VLEO Satellite Development and Remote Sensing: A Multidomain Review of Engineering, Commercial and Regulatory Solutions," *Preprints*, December 2025. https://doi.org/10.20944/preprints202512.1721.v1

[3] J. Llop, P. Roberts, Z. Hao, L. Tomas, and V. Beauplet, "Very low earth orbit mission concepts for earth observation: Benefits and challenges." In *Reinventing Space Conference*, November 2014, reinventing Space Conference; Conference date: 18-11-2014 Through 20-11-2014.

[4] L. Berthoud, R. Hills, A. Bacon *et al.*, "Are Very Low Earth Orbit (VLEO) satellites a solution for tomorrow's telecommunication needs?" *CEAS Space Journal*, vol. 14, pp. 609–623, 2022.

[5] P. C. E. Roberts, "VLEO and space sustainability: how does Very Low Earth Orbit contribute to the shift towards more sustainable space activities?" *CEAS Space Journal*, 2025, published online 25 November 2025. https://doi.org/10.1007/s12567-025-00677-w

[6] J. P. Thayer, W. K. Tobiska, M. D. Pilinski, and E. K. Sutton, *Remaining Issues in Upper Atmosphere Satellite Drag.* American Geophysical Union (AGU), 2021, ch. 5, pp. 111–140. https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1002/9781119815570.ch5

[7] V. Bernstein and M. Pilinski, "Drag coefficient constraints for space weather observations in the upper thermosphere," *Space Weather*, vol. 20, no. 5, p. e2021SW002977, 2022, e2021SW002977 2021SW002977. https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2021SW002977

[8] M. D. Pilinski, B. M. Argrow, and S. E. Palo, "Drag coefficients of satellites with concave geometries: Comparing models and observations," *Journal of Spacecraft and Rockets*, vol. 48, no. 2, pp. 312–325, 2011. https://doi.org/10.2514/1.50915

[9] N. Crisp, P. Roberts, F. Romano, K. Smith, V. Oiko, V. Sulliotti-Linner, V. Hanessian, G. Herdrich, D. García-Almiñana, D. Kataria, and S. Seminari, "System modelling of very low Earth orbit satellites for Earth observation," *Acta Astronautica*, vol. 187, pp. 475–491, 2021. https://www.sciencedirect.com/science/article/pii/S0094576521003519

[10] J. Frenzel, "Genetic algorithms," *IEEE Potentials*, vol. 12(3), pp. 21–24, 1993.

[11] B. Alhijawi and A. Awajan, "Genetic algorithms: Theory, genetic operators, solutions, and applications," *Evolutionary Intelligence*, pp. 1–12, 2023.

[12] M. Foreback, E. Imata, V. Ragusa, J. Weiler, C. Shao, J. Wagner, K. G. Skocelas, J. Sy, A. Hafez, W. Banzhaf *et al.*, "Eclipse: An evolutionary computation library for instrumentation prototyping in scientific engineering," *arXiv preprint arXiv:2601.05098*, 2026.

[13] J. Rolla, B. Reynolds, J. Weiler, A. Connolly, R. Debolt, A. Machtay, B. Sipe, and D. Wells, "Design of 3D Antenna Geometries Using Genetic Algorithms," *The Interplanetary Network Progress Report*, vol. 42-234, pp. 1–26, August 2023.

[14] J. Rolla, B. Reynolds, D. Wells, J. Weiler, A. Connolly, and R. Debolt, "Design of antennas from primitive shapes using genetic algorithms," *The Interplanetary Network Progress Report*, vol. 42-237, Jet Propulsion Laboratory, Pasadena, California, pp. 1–47, May 2024.

[15] J. Rolla, B. Reynolds, J. Weiler, D. Wells, M. Foreback, A. Connolly, E. Dolson, and C. Ofria, "Designing optimized antennas for science applications using evolutionary algorithms," *The Interplanetary Network Progress Report*, vol. 42-242, Jet Propulsion Laboratory, Pasadena, California, pp. 1–29, August 2025.

[16] G. S. Hornby, "ALPS: The age-layered population structure for reducing the problem of premature convergence," *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation*, pp. 815–822, July 2006.

[17] M. D. Pilinski, "Dynamic gas-surface interaction modeling for satellite aerodynamic computations," Ph.D. dissertation, University of Colorado at Boulder, Boulder, CO, USA, 2011, ph.D. thesis. https://scholar.colorado.edu/asen_gradetds/37/

[18] F. Vavak and T. Fogarty, "Comparison of steady state and generational genetic algorithms for use in nonstationary environments," In *Proceedings of IEEE International Conference on Evolutionary Computation*, 1996, pp. 192–195.

[19] G. S. Hornby, "Steady-state ALPS for real-valued problems," *Proceedings of the 11th Annual Conference on Genetic and Evolutionary Computation*, pp. 795–802, 2009. https://doi.org/10.1145/1569901.1570011

[20] B. O. Community, *Blender - a 3D modelling and rendering package*, Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018. http://www.blender.org

[21] J. Zhong, X. Hu, J. Zhang, and M. Gu, "Comparison of performance between different selection strategies on simple genetic algorithms," *International Conference on Computational Intelligence for Modelling, Control and Automation and International Conference on Intelligent Agents, Web Technologies and Internet Commerce (CIMCA-IAWTIC'06)*, vol. 2, pp. 1115–1121, 2005.

APPENDIX

I. Individual Genes and Genetic Algorithm Parameters

Table 2. Parameters used in genetic algorithm.

Category	Type	Parameters
General	Run	Num. Individuals, Num. Generations, Fitness Function, Target pattern, Steady-state flag, Forced diversity flag
	Individual	Max. Tree Depth, Max. Shapes, Shell thickness
	XFdtd	Frequency range, Frequency step
Shape	Cuboid	Shape allowed, Length, Width, and Height Ranges
	Sphere	Shape allowed, Radius Range
	Cylinder	Shape allowed, Radius, and Height Range
	Cone	Shape allowed, Radius 1, Radius 2, and Height Range
Selection Methods	Tournament	Percent of parents, Group Size
	Roulette	Percent of parents
	Rank	Percent of parents
Genetic Operators	Dim. Mutation	Percent of children, St. Dev. %
	Rotation Mutation	Percent of children, St. Dev. %
	Location Mutation	Percent of children, St. Dev. %
	Grow Mutation	Percent of children
	Prune Mutation	Percent of children
	Regen. Mutation	Percent of children
	Side Switch	Percent of children
	Gene Crossover	Percent of children
	Branch Crossover	Percent of children
	Reproduction	Percent of children
	Injection	Percent of children

Table 3. Primitive shape genes.

Gene	Values
Shape Type	Cuboid, Cylinder, Cone, Sphere
Dimensions	Varies by shape type (See Table 4)
Location	Cartesian coordinates of shape midpoint
Rotation	θ, ϕ
Connected From	Shape built from
Connected To	Side that shapes are attached to

Table 4. Dimensions for each shape type.

Shape Type	Dimension Genes
Cuboid	Length, Width, Height
Cylinder	Radius, Height
Cone	Inner Radius, Outer Radius, Height
Sphere	Radius

Table 5. Run parameters.

Parameter	Value	Parameter	Value
Alps Layer Size	100	Cube Min Height (m)	0.007
Alps Num Layers	3	Sphere Max Radius (m)	0.113
Alps Tournament Size	10	Sphere Min Radius (m)	0.003
Alps Age Gap	5	Cylinder Max Radius (m)	0.113
Alps Num Evaluations	200000	Cylinder Min Radius (m)	0.003
Max Workers	100	Cylinder Max Height (m)	0.366
Percent Rotate	0.20	Cylinder Min Height (m)	0.007
Percent Dimension	0.20	Phi (deg)	0.0
Percent Material	0.0	Theta (deg)	0.0
Percent Regen Proportion	0.20	Temperature (K)	1200.5
Percent Regen Random	0.0	Velocity (m/s)	7800.45
Percent Grow	0.20	Oxygen Composition (m^{-3})	$1e11$
Percent Prune	0.20	O_2 Composition (m^{-3})	0.0
Max Tree Depth	20	N_2 Composition (m^{-3})	0.0
Max Shapes	15	Helium Composition (m^{-3})	0.0
Cube Max Length (m)	0.366	Hydrogen Composition (m^{-3})	0.0
Cube Min Length (m)	0.006	GSI Model	Fixed
Cube Max Width (m)	0.226	Surface Mass (amu)	65.0
Cube Min Width (m)	0.006	Alpha	0.93
Cube Max Height (m)	0.226	Max Spacecraft Volume	0.0172