Data-Driven Constellation Architecture Design Using Integrated Models

Allen Wautlet, Dave Brown, Greg Thanavaro

BAE Systems Space & Mission Systems Inc
Boulder, CO; 303-939-4000

Abstract

The modern space mission landscape requires consideration of numerous trade variables to deliver optimal mission performance at low cost. Academic methods exist to address such challenges, however, practical deployment of these methods to constellation mission design remains uncommon.

This paper presents a practical space mission constellation architecture approach that employs proven statistical, data science, and machine learning techniques on the products of an integrated cost and engineering modeling framework. When deployed at the early stages of constellation development, this integrated modeling framework and analysis approach provides stakeholders insight into key design parameters that drive mission performance and cost sensitivity. Furthermore, it pinpoints promising design regions in large trade spaces that can be further examined and refined by subject matter experts. This approach leads to better decision making earlier in the acquisition timeline and increases the efficiency of design cycles.

1.0 Introduction

The traditional industry methodology for designing and evaluating satellite constellations for the early phases of a program involve the creation and maturation of a few design concepts with lean budgets and deadlines. In this fast-paced, linear design flow, cost evaluations are often only addressed towards the end of the process once the design concept has stabilized. Any cost decisions made prior to that point are predominantly based on Subject Matter Expert (SME) judgements and may not fully consider potential downstream effects on the larger system. This traditional design methodology has the following issues:

- Design concepts are often predicated on previous programs or designs which may have limited applicability to the mission.
- Few design iterations are performed, resulting in potentially non-optimal concepts for cost and the required mission performance.
- Design teams have low flexibility with regards to proposed concepts often opposing drastic change in response to aggressive deadlines.
- Chronological separation of the technical design from the cost assessment results in low transparency or complete loss of insight into the cost implications of early design decisions.

The approach presented within this paper poses an alternative methodology to supplement early-stage satellite constellation design that parametrically links performance requirements directly to cost through the integration of engineering and cost models currently being employed at BAE Systems. This parameterized linkage captures not only the downstream impacts to cost from design choices but also allows for the determination of optimal concepts that minimize cost for mission requirements. Furthermore, this approach empowers stakeholders to perform on-the-fly "what-if" analyses and sensitivity studies at the earliest stages of design.

2.0 Methodology

To illustrate a real-world scenario where this methodology can be deployed and showcase its powerful capabilities, an example customer request was formulated:

> *Suppose that a three-year remote sensing mission is required at a competitive cost to help firefighters monitor potential fire risk locations and locate active fires across the continental United States at a maximum revisit rate of 6 hours. This mission requires that each vehicle be capable of monitoring at least 30 separate locations per day with an image resolution less than or equal to 8 meters per pixel through smoke for a duration of no less than 2 seconds per location.*

Adequately responding to this request requires multiple types of engineering and analyses to be coordinated at the systems level. This becomes complicated because a point-of-departure for the design must be chosen and if a problematic or irrelevant starting point design is chosen, the optimal design concept may be inaccessible for the program before the first iterations even begin.

Two ways of proactively addressing these potential shortcomings are to: 1) have enough institutional wisdom to know where to begin or 2) deploy a systems-level framework which can capture the mission, subsystem, and cost level interdependencies on each other. Both methods for addressing potential shortcomings can be leveraged together in a data-driven capacity to build parametric surrogate models that directly capture high-level sensitivities within the entire trade space.

The creation of a parametric surrogate model for a complex design space follows four general steps which are depicted in Figure 1. First, the independent parameters influencing the design space are ascertained, constituting the *Continuous Design Space*. Second, the continuous design space is subdivided (this can be accomplished with varying degrees of resolution), creating the *Discretized Design Space*. Third, instead of simulating each possible combination of parameters from the large, discretized design space, a Design of Experiments (DoE) is deployed to maximize design space coverage while minimizing number of simulations. Finally, a *Parametric Surrogate Model* is developed using various model fitting or machine learning techniques that captures the design space sensitivities and approximates, with some quantifiable confidence, designs that have not been simulated.

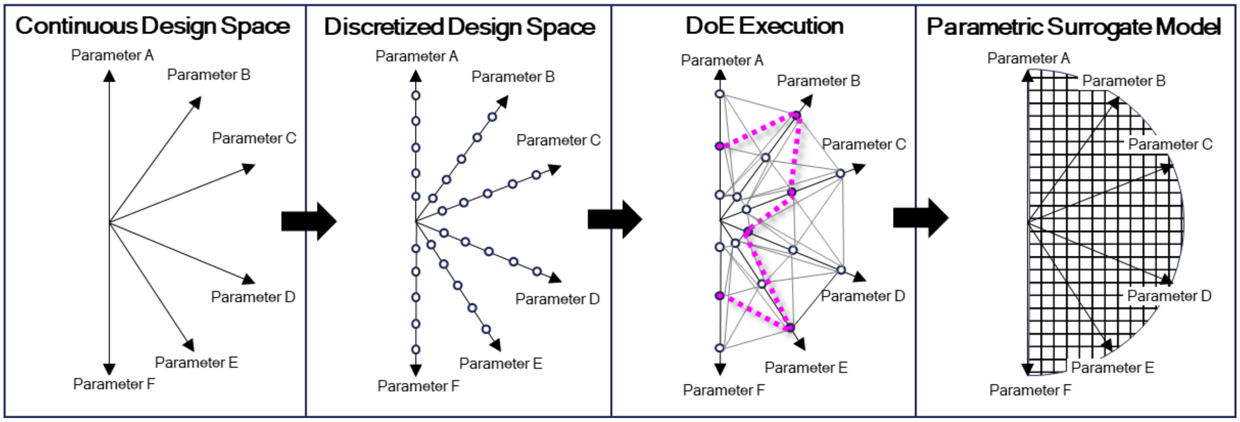

Figure 1: Parametric Surrogate Model Creation Process

2.1 Design Space

For any remote sensing mission three questions are always paramount: 1) What does the customer want to observe, 2) At what rate does this need to be observed, and 3) How much will it cost? Two of these three questions lend themselves well to ascertaining some of the driving independent design parameters in the design space. The requirements specify that the customer wants to identify fires through smoke which can be accomplished using short wave infrared (SWIR) sensors but observing said fires every six hours is not straight forward. Gap time is a mission performance metric that represents the time between consecutive accesses of single point on earth and is a function of complex interactions between sensor parameters, number of vehicles, vehicle spacing, and altitude and is typically arrived at through simulation.

The customer request further requires that the resolution per pixel must be no more than 8 meters. When considered together, certain sensor parameters, altitude, and spherical geometry will yield ground sample distance (GSD) which represents the distance between two consecutive pixels in any two-dimensional direction. However, if the sensor phenomenology is assumed to be static, GSD and altitude can be used to arrive at other sensor properties like physical dimensions. This strategic rearrangement allows GSD to become a driving independent design parameter, although it is conceivable one may want to instead *arrive* at GSD instead of using it as a trade variable.

Finally, the customer requires that at least 30 separate locations be monitored per satellite per day over a region roughly the size of the continental United States. This capacity performance requirement at the satellite-level will be a function of sensor parameters, agility, and altitude and must be determined through simulation because it is non-deterministic and relies upon a scheduling algorithm to compute a result.

From these requirements a picture of linkages between design requirements and independent design space parameters inputs begins to take shape. Table 1 compiles what parameters are expected to influence the customer performance requirements.

Table 1: Design Parameters Influencing Performance

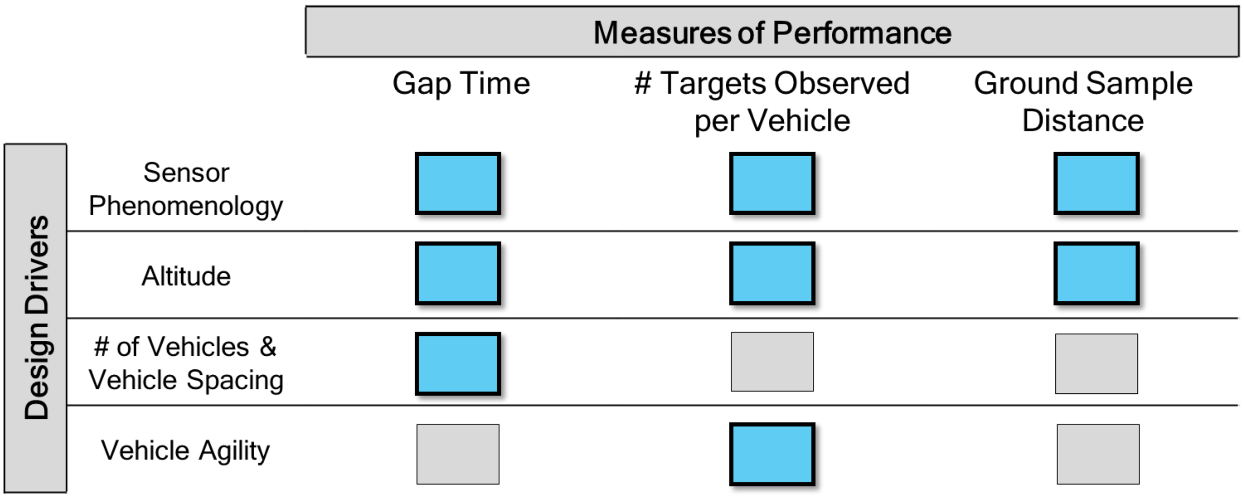

Despite linkages between design parameters and performance requirements being easily formalized, a linkage between those same design parameters and cost is not obvious and must be developed from simulations and analysis. Cost is complex, but for remote sensing missions cost can largely be attributed to the development, production, operations, and launch of space vehicles.

A space vehicle is comprised of one or more payloads hosted by a bus. The parameters that influence space vehicle design are numerous, but at a high level some of the largest drivers are the payload(s), required agility, design life, orbit, and risk-posture. Thus, it follows that cost should be a direct function of mission impacts on space vehicle design whose constituent components have an interdependency on each other’s design.

Certain space vehicle design parameters have been set static as assumptions for the example mission developed within this paper. A few key assumptions are compiled in Table 2.

Table 2: Select Space Vehicle and Mission Assumptions

Assumption	Rationale
24°-45° Latitude Assessment Range	Corresponds to roughly the latitudinal range of the continental United States
Walker Delta Constellation	Symmetric structure and uniform performance
Single-string space vehicle redundancy	Reductions in mass for better multiple manifest launch capabilities and no mission risk posture stated by customer request
Four attitude actuator solution (RWA or CMG)	Permits wider range of agility
Controlled reentry for orbital debris mitigation	NASA/international guideline

2.2 Analysis Flow

Figure 2 depicts the formalized analysis flow developed through the formulation of notional model linkages and inputs so far. At the center of this process is the integrated model framework that captures each individual analysis' interdependency on one another while assessing launch accommodations and cost simultaneously. The inputs from a Design of Experiments (DoE) directly link to the modeling outputs, enabling data exploration and parametric models to be developed.

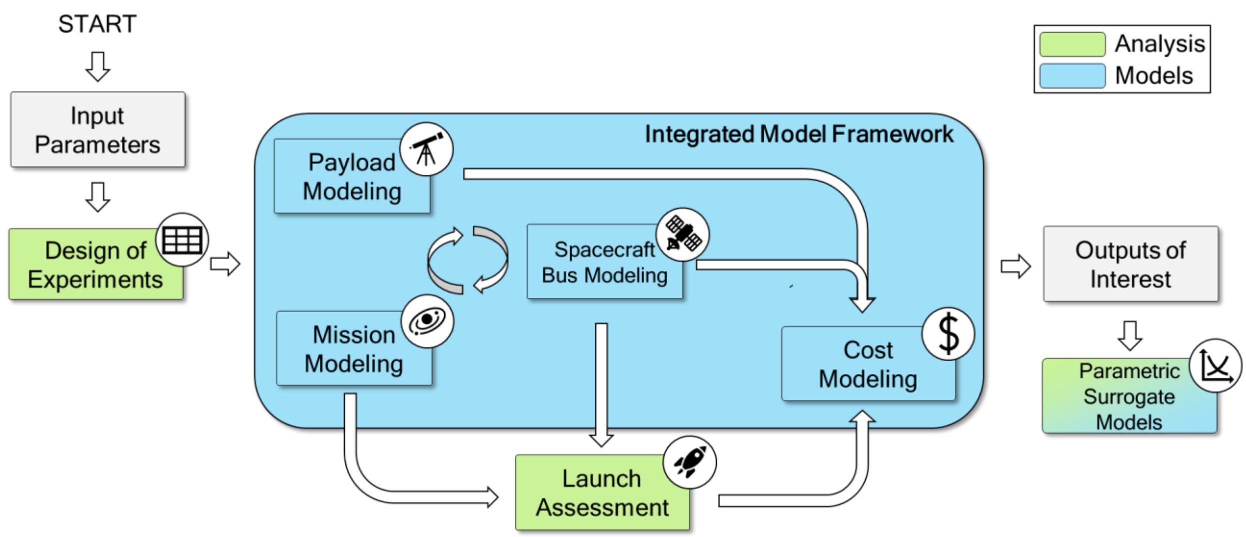

Figure 2: Integrated Model Analysis Flow

The analysis flow shown in Figure 2 provides a process through which large design trade studies can be executed. This data-driven approach yields parametric models, providing decision makers key insights into the sensitives impacting their mission.

2.2.1 Design of Experiments

A major bottleneck in the execution of any large trade space is the number of samples or simulations necessary to achieve a desired understanding of that space. Finding a reduced number of samples which adequately describe the design space can be easily achieved using a design of experiments. Many different DoE methods exist, but for surrogate model development, a Latin-Hypercube (LHC) sampling design proves most advantageous for sampling continuous data efficiently.[1] The discrete data will be sampled using straight-forward factorial combinations, although it is conceivable that if the number of factorial combinations were large enough, that a fractional-factorial scheme could be used instead. Table 3 details the DoE inputs, their ranges, and the DoE sample scheme applied to each input.

For each factorial combination of number of constellation planes and number of satellites per planes, an LHC combination of altitude, slew rate, and ground sample distance is associated until there are 45 LHC samples for each possible factorial combination. The resulting number of samples to be processed through the integrated analysis flow is 1800.

Table 3: DoE Inputs and Ranges

Design Input	Lower Range	Upper Range	DoE Type
Altitude (km)	450	800	LHC
Number of Constellation Planes	4	11	Factorial
Number of Satellites per Plane	2	6	Factorial
Slew Rate (deg/s)	0.5	4.0	LHC
Ground Sample Distance (m)	3.0	30.0	LHC

2.3 Engineering Models

With a rough understanding of the factors influencing mission performance and cost, it becomes clear that various analyses and models must be integrated so that design space interdependencies are captured. Within the Digital Engineering group at BAE Systems, a unified, configurable, and rapid mission modeling, space vehicle sizing, and payload estimating framework has been developed called MOSAIC (Mission Optimization and System Architecture Insight Capability).

The mission modeling portion of MOSAIC includes a multitude of features but a few pertinent to the example customer request are orbit propagation, target collection scheduling, and space vehicle maneuver dynamics. This modeling capability is leveraged for two specific types of analysis: access and capacity. Access analyses are predicated on interval-based gap time between satellite lines-of-sight and are used to derive revisit metrics. Capacity analyses leverage satellite agility, sensor fields-of-view, and scheduling algorithms to assess a constellation's target observation capability.

The space vehicle modeling capability in MOSAIC is a physics-first, rapid, iterative, and subsystem integrated tool suite that leverages historic design choices and SME-developed sizing routines. This capability serves as the primary mode through which mission performance is transposed into cost and includes numerous inputs ranging from the mission and payload level to the subsystem and space vehicle level. The results from this capability are passed into the launch assessment analysis and cost modeling portion of MOSAIC.

The payload modeling module in MOSIAC is a configurable first-order sizing capability that leverages physics-based principles and historic regressions. This capability converts mission parameters into size, weight, and power (SwaP) requirements for space vehicle modeling.

2.4 Cost Modeling

In the example scenario, a key element of the customer's request is to design a mission *at a competitive cost*. Cost is a major consideration when analyzing architecture design trades. Examples of how cost can be considered in a mission analysis may be as: 1) a *constraint* (i.e. what is the highest level of performance for a given cost point), 2) a *minimum* (ie. what is the lowest cost solution possible for acceptable performance?), or 3) as an *independent variable* (ie. what is the optimal solution if both cost and performance are weighted equivalently in importance—the best value approach). In wide-ranging trade spaces, there is a fundamental need for reliable cost correlations and sensitivity predictions in relation to modeled performance for each of these cases.

Cost modeling capabilities are therefore an essential component of the integrated model framework.

2.4.1 Parametric Estimating

Projecting cost can be difficult while the design is still in flux during the early phases of a program. This can be particularly challenging in cases where an Analysis of Alternatives (AoA) is being performed across a broad range of potential technical solutions. For a given technical baseline, there may be a notion of what materials and how many labor hours are required, but without supplier quotes and inputs from SMEs across the scope of a project, cost analysts must use their best judgement to field accurate projections. Even in cases where a design baseline is well established, the time and resources required to perform a detailed cost estimate can be significant. For this reason and others, *Parametric Estimating* is commonly deployed during the early phases of a program's lifecycle—primarily Phase A Concept Development and Phase B/C Design efforts.[3]

Parametric Estimating methodologies often employ Cost Estimating Relationships (CERs) derived from regression analyses that use independent design parameters to predict cost as a dependent variable. CERs can range from complex equations to a simple multiplier such as a *rate* (eg. labor rate or dollar per pound), *factor* (eg. System's engineering is X% of underlying prime mission equipment, or a *ratio* (eg. NR/T1 where non-recurring is a ratio of the recurring first unit cost).[3]

CERs prove beneficial in that they can adapt well to changing design parameters and provide statistical results offering insights into the quality and uncertainty of an estimate. CERs however, can be time-consuming to develop and require the use of cleaned relevant and reliable historical data which may be difficult to obtain.[3] Data relevancy can be objective, as certain design parameters being estimated may simply fall outside the bounds of the CER's underlying data. Relevancy can also be subjective in cases where design inputs are within the bounds of a CER's underlying dataset, but other factors and assumptions associated with the estimated design can add uncertainty to the CER's applicability. Providing crosschecks and offering insights into a CER's construction and assumptions can help mitigate some of these challenges.

Utilizing parametric cost estimating in an integrated model framework requires two essential ingredients: 1) having relevant CERs coverage for the scope of mission architectures and 2) having reliable independent input variables for the CERs.

2.4.2 Ball Estimating Relationships (BALLER)

There are a small number of widely accepted and utilized CERs within the cost community, but for the purposes of the analysis presented herein the Cost Modeling approach employs custom BAE Systems specific CERs called BALLER, whose namesake is derived from legacy Ball [Aerospace] Estimating Relationships. BALLER models both non-recurring (NR) and recurring (RE) costs separately and are derived from over 30 years of relevant historic program costs. The CER development process follows the principles of the International Cost Estimating & Analysis Association's (ICEAA) Cost Estimating Body of Knowledge (CEBoK®), which includes three distinct efforts outlined in Figure 3.

Figure 3: CER Development & Integration Process

Data Normalization – *Clean the data needed for CERs*

1. Develop a Standard Work Breakdown Structure (SWBS).
2. Map SWBS to all historical programs' cost data *(map tech. data to SWBS in step 5)*.
3. Normalize cost data: allocate NRE vs RE charges and adjust costs to a common base-year.
4. Identify candidate variables for data collection (internal and external assessment).
5. Collect and review technical and programmatic data.
6. Organize and consolidate data to support efficient regression analysis.

Regression Analysis – *Analyze the data and develop the CERs*

1. Data exploration: assess correlations of independent variables (*to each other, and to cost*).
2. Review preliminary data analysis with technical SMEs (*revisit tech data as required*).
3. Perform regression analysis and develop CERs.
4. Review preliminary CERs with technical SMEs (*revisit previous steps as necessary*).
5. Publish and document CERs once validation is complete.

Integration into MOSAIC – *Utilize the data by integrating CERs into model framework*

1. Review CERs and assess all independent variables.
2. Augment engineering models to ensure all CER inputs are represented if necessary.
3. Integrate CERs into automated mission modeling tools.
4. Develop software unit tests to validate CER integrations against independent calculations.

Once the CER development and integration process is completed, a comparison of the BALLER CERs in relation to other industry standard CERs also integrated into the modeling framework is performed. An illustration of how a variety of T1 satellite bus costs were estimated utilizing each of the CER alternatives, relative to an arbitrary baseline, is visualized in Figure 4.

Figure 4 demonstrates that, for a wide range of vehicle designs, the BALLER cost models generally project lower costs than some other industry standard CERs. Many of these industry standard CERs are built from a multitude of historic programs across many different eras, customers, suppliers, and mission classes resulting in projections that are less representative of an individual supplier's specific cost and more representative of broader market costs that have been historically realized. Procuring U.S. government agencies and commercial buyers each have their own unique influences on the costs of their space programs as well based on fundamental differences in mission requirements, oversight, and acquisition processes. Interestingly in all cases, the relative T1 bus costs follow a log-normal right skewed shape, a distribution commonly associated with space vehicle costs.

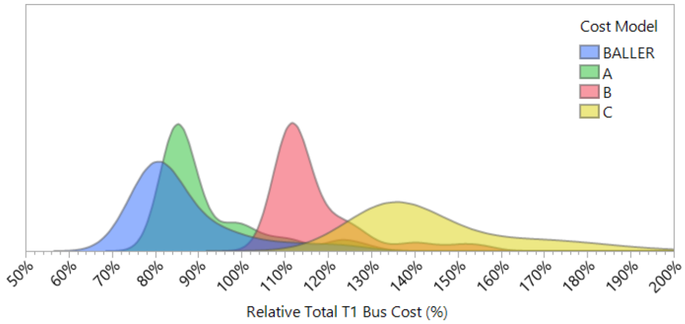

Figure 4: CER Comparison-Satellite Relative T1 Bus Cost Distributions

2.4.3 Cost Estimating Methodology

The architecture costs assessed for the mission example developed in this paper are the relative total space segment development and production costs. These costs are modeled using the custom BAE Systems specific CERs and are relative to an arbitrary baseline for sensitive information protection. The arbitrary baseline cost used to develop relative space vehicle development and production costs is different than the one used to develop the T1 bus cost distributions in Figure 4. Relative total space vehicle development and production costs include the NR and RE costs associated with space vehicle development and material procurement, systems engineering, program management, assembly integration & test (AI&T), launch operations, and launch procurement. Additionally, a cost improvement curve is applied to multiple unit builds, capturing reductions in RE associated with multi-unit procurements. The cost improvement curve is applied assuming a 90% learning curve slope with an arbitrary applicability up to 12 build units. In the cases where less than 12 build units are modeled, the learning curve applicability number tracks along with the maximum number of builds.

2.5 Launch Vehicle Assessment

Launch Vehicle accommodations are assessed for each design using modeled space vehicle SWaP and mission altitude inputs. Each design modeled through the integrated framework is assigned a launch vehicle using a set of rules that minimizes launch cost while maximizing number of satellites capable of being launched from a single launch vehicle. These rules are the following:

1. The number of satellites that can be launched by a single vehicle is the number of satellites that, when summed over mass and volume, can meet the launch vehicle's mass and volume constraints.
2. The number of launch vehicles required is the number of constellation planes multiplied by the number of launch vehicles required to populate that plane given the maximum number of satellites that can fit within a single launch vehicle as determined by the previous rule.
3. The chosen launch accommodation is the minimum cost solution arrived at by multiplying the number of launches necessary by estimated launch vehicle cost.

3.0 Analysis

Before proceeding straight to building parametric surrogate models it is prudent to gain an understanding of how the data generated behaves through data exploration. When visualized creatively, data inspection alone may uncover insights into the nature of the relationships that need to be synthesized. The more simply a relationship can be captured the better it may be extrapolated on or understood for development of institutional wisdom.

Various combinations of visualizations were explored for the example scenario analysis but a few of note are presented here as they clearly and directly capture intriguing major relationships and help build confidence in the results. The first figure of interest is Figure 5, which depicts the average targets observed per satellite per day vs the average gap time and is colored by altitude.

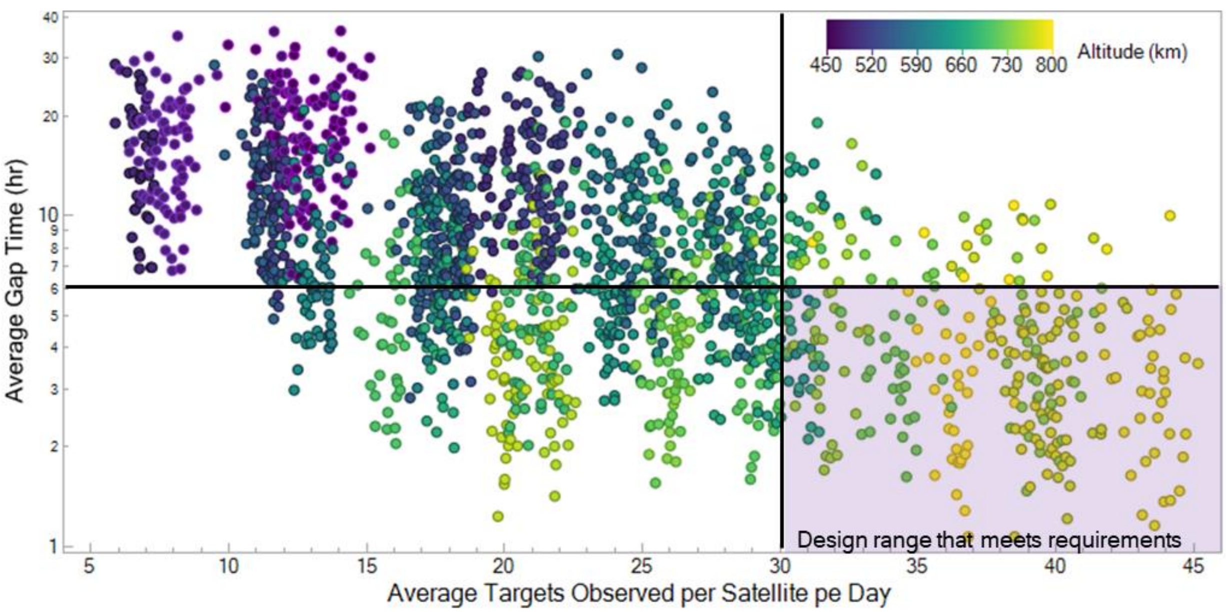

Figure 5: Average Targets Observed vs Average Gap Time by Altitude

Figure 5 offers a few key insights into the two performance metrics that were arrived at through simulation in the mission modeling part of the integrated engineering models. This figure shows that average gap time and average targets observed have an exponential decay relationship. This relationship is intuitive suggesting that, as gap time decreases (i.e., revisit rate increases) more targets can be imaged. The altitude coloring indicates what should be expected, but informs which *direction* matters in altitude. The major takeaway here is that as altitude increases, both performances metrics improve—gap time goes down and the number of targets observed goes up. These leads to a question: If increasing altitude leads to better mission performance, what is the impact to cost? This question requires the consideration of many more variables and will best be addressed through the development of a parametric surrogate model.

The next visualization of Interest Is Figure 6, which shows the relative cost vs average targets observed per satellite and is colored by GSD. The average number of targets observed per day by each vehicle improves gradually linearly with respect to cost. Depending on the altitude and sensor phenomenology, GSD requirements can often become a direct representation of sensor exquisiteness which is captured well when plotted against cost. It is obvious from Figure 6 that both GSD and average number of targets observed per day by each satellite both drive cost, but to what extent GSD drives cost is not fully apparent.

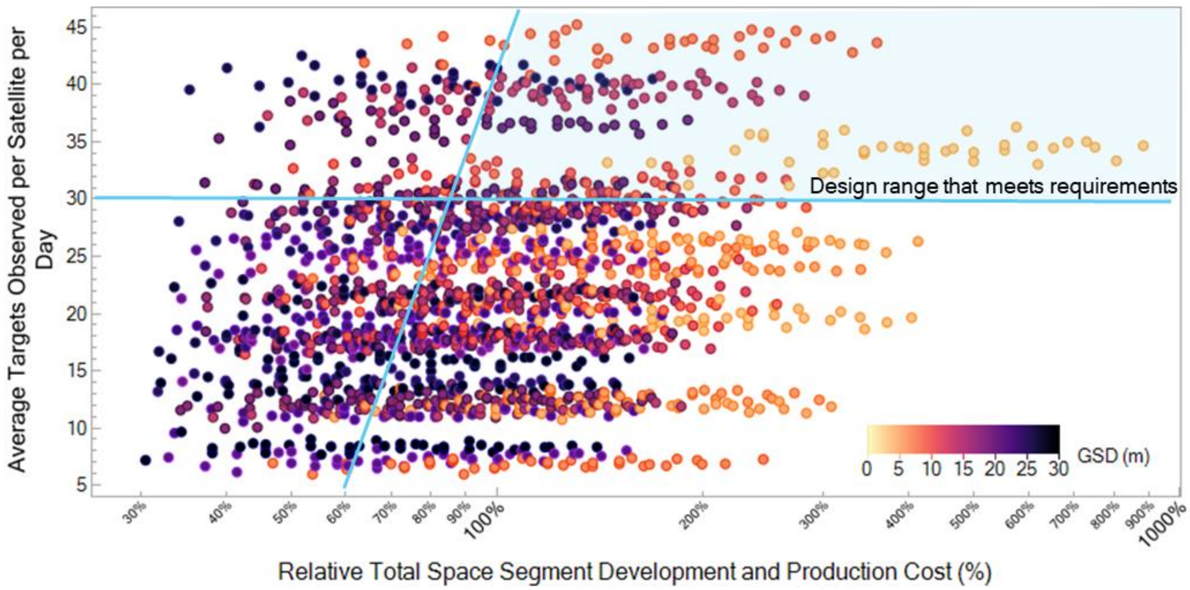

Figure 6: BALLER-Based Relative Cost vs Average Targets Observed by GSD

The last visualization of interest for this data exploration exercise presents the clearest link between performance and cost illuminating much of the architecture design space. Figure 7 shows the relative cost vs average gap time colored by GSD. GSD presents as having no effect on average gap time which is expected as resolution has no bearing on revisit metrics. However, the primary takeaway in Figure 7 is that cost trends much more strongly with the average gap than it does with average targets observed per satellite per day (Shown in figure 6).

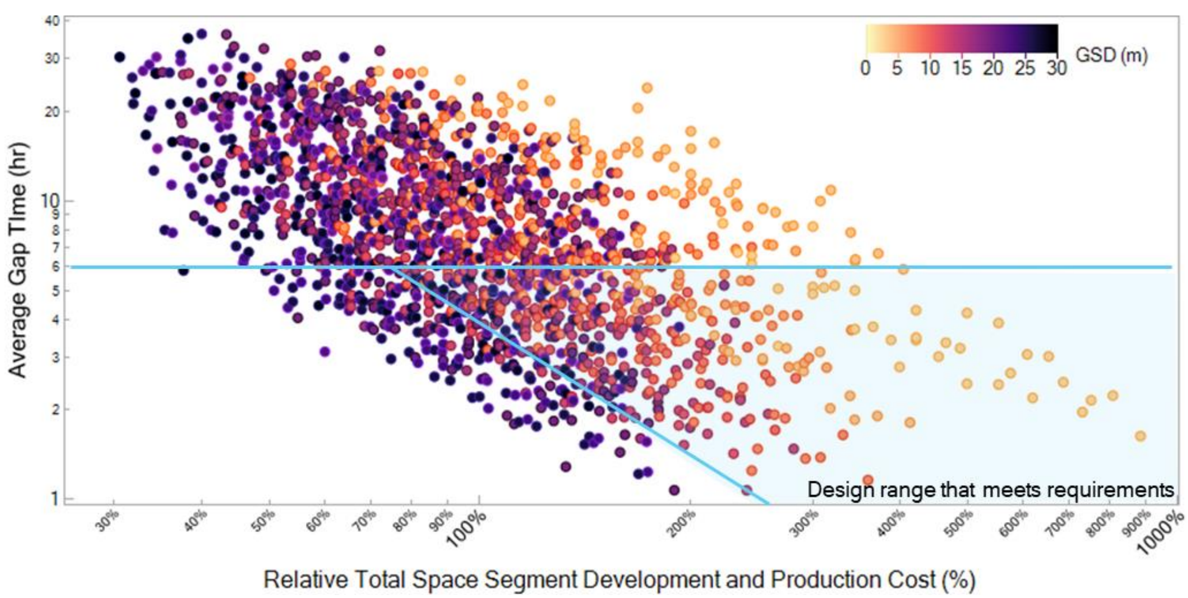

Figure 7: BALLER-Based CER Modeled Relative Cost vs Average Gap Time by GSD

The visualizations shown in Figures 6 and 7 indicate that the optimal cost for the performance requirements given is somewhere in the range between 80% and 100% relative total space segment development and production cost. What this optimal cost is and what design parameters yield it can be ascertained from a parametric surrogate model.

4.0 Parametric Surrogate Modeling

Surrogate models can be created in a more directed fashion using some of the insights gained through the exploration process. Due to the phenomena being captured through the surrogate models being fundamentally grounded in physics and geometry, it is preferable to use more simplistic and easily interpretable approaches like multiple-linear regressions instead of models like neural networks or random forests. The rationale for building the most easily interpretable models in this domain is that the end goal for these parametric surrogate models is not just to create tools for architecture design, but to inform on the broader trends at work within the design space and possibly beyond.

The parametric surrogate models are created using stepwise multiple-linear regressions for relative total space segment cost, average gap time, and average targets observed per satellite per day using the design inputs specified in Table 3. Each model is created using varying degrees of polynomial and factorial combinations of inputs and occasional logarithmic transforms. The prediction model is created using a logistic regression for the same combinations of inputs. An interactive multi-trend matrix visualization of the parametric surrogate models is shown in Figure 8.

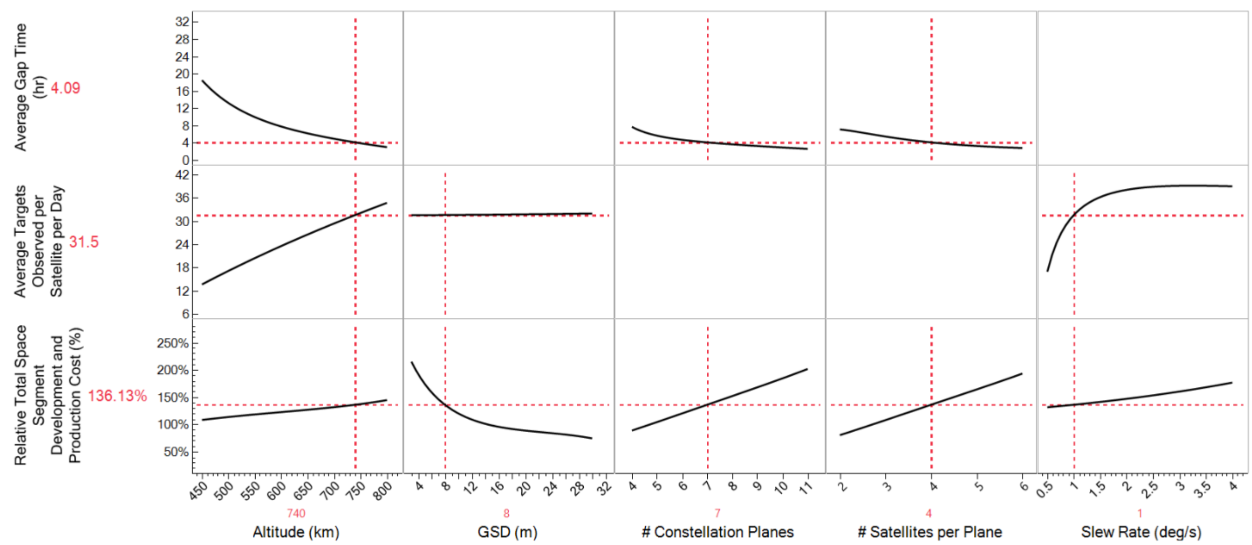

Figure 8: Parametric Surrogate Models

Visualizing the design space trends in the matrix fashion shown in Figure 8 allows for the identification of input sensitivity drivers on outputs being modeled to be isolated. The trends seen in Figure 8 are visualized relative to each other and changes in one input parameter may change or shift the slopes of another's relationship to an output. The first key takeaway ascertained from visualizing the parametric surrogate model is that not all input parameters influence performance metrics, but each input parameter does drive cost. Higher altitudes result in both cost and mission performance improvements (gap time decreasing and targets observed increasing) with faster rates of performance increases compared to cost increases indicating a key tradeoff for balancing cost and performance. The number of planes and number of satellites per plane increase performance at a slower rate than they increase cost, suggesting that these parameters should be increased sparingly. The opposite can be said about slew rate—the mission performance increases from slew rate improvements have slower rates of cost increases between the ranges of 0.5 deg/s and 1.5 deg/s.

4.1 Verification and Validation

Validation and verification are critical exercises necessary for instilling trust within the efficacy of the surrogate models and proving that the trends captured are indeed believable. The surrogate models are assessed in three separate ways, through goodness of fit metrics, verification of linearity in predicted versus actual plots, and independent model verification. The goodness of fit metrics for the surrogate models developed are detailed in Table 3.

Table 3: Parametric Surrogate Model Goodness of Fit Metrics

Parametric Surrogate Model	R^2	RMSE	RMSPE
Average Gap Time (hr)	0.90	2.42	23.8%
Average Number of Target Observations per Vehicle per Day	0.97	1.52	6.5%
Relative Total Space Segment Development and Production Cost (%)	0.99	3.50	5.6%

The R^2 values in the goodness of fit metrics being greater than 0.8 suggest that the trends are captured well by the models created. The Root Mean Squared Error (RMSE) represents the average distance a model-predicted value will be from truth. RMSE provides a good general estimate of the expected mean error for the data at-large, but this value can suggest many different things depending on where within the model it is applied. An RMSE value assessed at the lower numeric areas of a model means a higher relative percent error than an RMSE value assessed at the upper numeric end of a model. Because of this, a different goodness of fit metric would be better utilized when assessing stringent customer requirements. Root Mean Squared Percent Error (RMSPE) is an alternative to RMSE that describes the same distance from truth a model-predicted value can be expected to lie except that it is formulated in terms of percent difference instead of the nominal data units. This is made clear when addressing the requirement for average gap time. The average gap time RMSE value is 2.42 hours which could, in the eyes of a customer, represent an expected error bound of about 40% for the 6-hour requirement. Alternatively, the RMSPE value of 23.8% would suggest that the expected error bounds would instead be 1.43 hours.

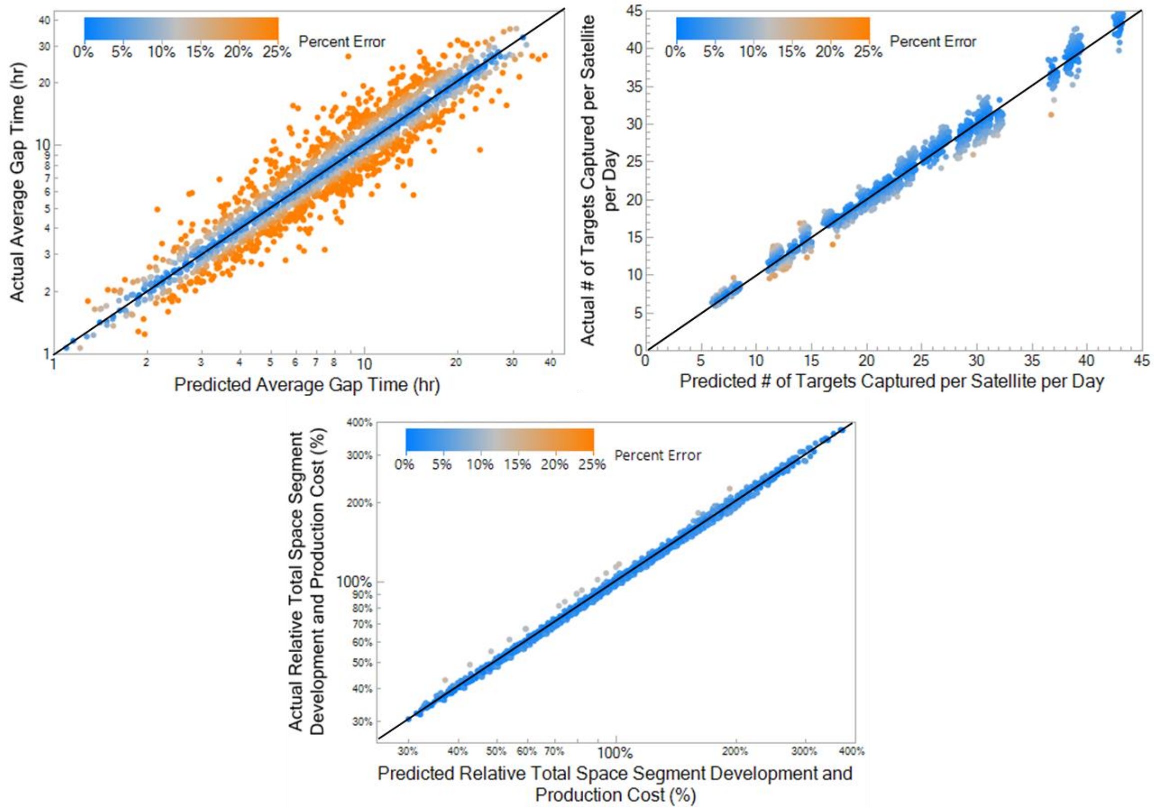

Figure 9: Predicted vs Actual Linearity Assessments

Assessing the linearity of predicted values to actual results provides critical insight into the behavior of the models and demonstrates where the model prediction range variance can be expected. Figure 9 is a composite of predicted vs actual linearity plots for each of the surrogate models created. The top left figure in the composite, the average gap model assessment, suggests that the model predicts well in high performing (low gap time) regions, but experiences the highest amount of variance around the middle of the numeric range. The average targets observed per satellite per day assessment, the top right figure, shows that there is generally low variance in the model, but chunks and holes within the plot suggest that many architectures within certain ranges of performance were not captured by the DoE. The bottom figure, the linearity in the cost model, shows that there is uniform low variance across all ranges.

5.0 Discussion

Promising design architecture concepts that are likely to meet the example mission requirements with the lowest impact to cost are developed using the parametric surrogate models and are aggregated in Table 5. What is clear from the results developed using the surrogate models is that concept designs likely to meet the mission requirements at lowest cost have between 15 and 20 satellites. These designs take advantage of the higher rate of performance increases relative to cost increases found by increasing altitude and slew rate while minimizing the number of total vehicles. The number of necessary launches for each promising design concept is equal to the number of constellation planes indicating that each design takes full advantage of multiple-manifest capabilities. Concept 2 provides the lowest relative cost solution while meeting requirements primarily by minimizing the number of vehicles and number of launches. A comparison between concepts 1 and 4 suggests that the cost impact from the additional launch required in concept 1 (6 launches, 18 satellites) may be close to the cost impact of the two additional satellites required in concept 4 (5 launches, 20 satellites) with a tradeoff in altitude.

The primary purpose of the parametric surrogate models is to capture trends and sensitivities in the design space, but most importantly, to indicate where promising architecture designs lie relative to performance and cost. The parametric surrogate model allows for the identification of a solution even if that solution was not used to build the surrogate model that predicted it. Because of this, no architecture should be lauded as a solution without a full closed-loop verification confirming that promising design concept does indeed yield a solution within the expected range of performance based on the goodness of fit metrics. Table 5 additionally compiles the actual simulated results of the design concepts identified by the parametric surrogate model.

Table 5: Concept Designs and Model Assessment

Design Parameters	Concept 1	Concept 2	Concept 3	Concept 4
Altitude (km)	780	790	772	745
Number of Constellation Planes	6	4	4	5
Number of Satellites per Plane	3	4	5	4
Total Number of Vehicles	18	16	20	20
Slew Rate (deg/s)	0.80	0.75	0.80	1.00
Ground Sample Distance (m)	8.0	8.0	8.0	8.0
Parametric Surrogate Model Prediction				
Gap Time (hours)	5.01	5.93	5.50	5.52
Average Number of Target Observations per Vehicle per Day	31	31	30	30
Relative Total Space Segment Development and Production Cost (%)	98.3	91.6	105.2	104.8
Simulation Model Actual				
Gap Time (hours)	5.51	5.50	4.57	4.64
Average Number of Target Observations per Vehicle per Day	30	30	30	30
Relative Total Space Segment Development and Production Cost (%)	98.3	88.0	1.01	1.00
Predicted Vs. Actual				
Gap Time Error	9.07%	7.82%	20.35%	18.97%
Target Observation Error	3.33%	3.33%	0.00%	0.00%
Cost Error	0.27%	3.93%	3.99%	4.55%

It is apparent from the results compiled in Table 5 that the parametric surrogate model correctly captures not only the trends present within the design space, but also the values associated with each promising design concept. The errors between predicted values and actual results are all within the expected RMSPE calculated for each model indicating that use of the surrogate models to identify realistic design concepts is valid.

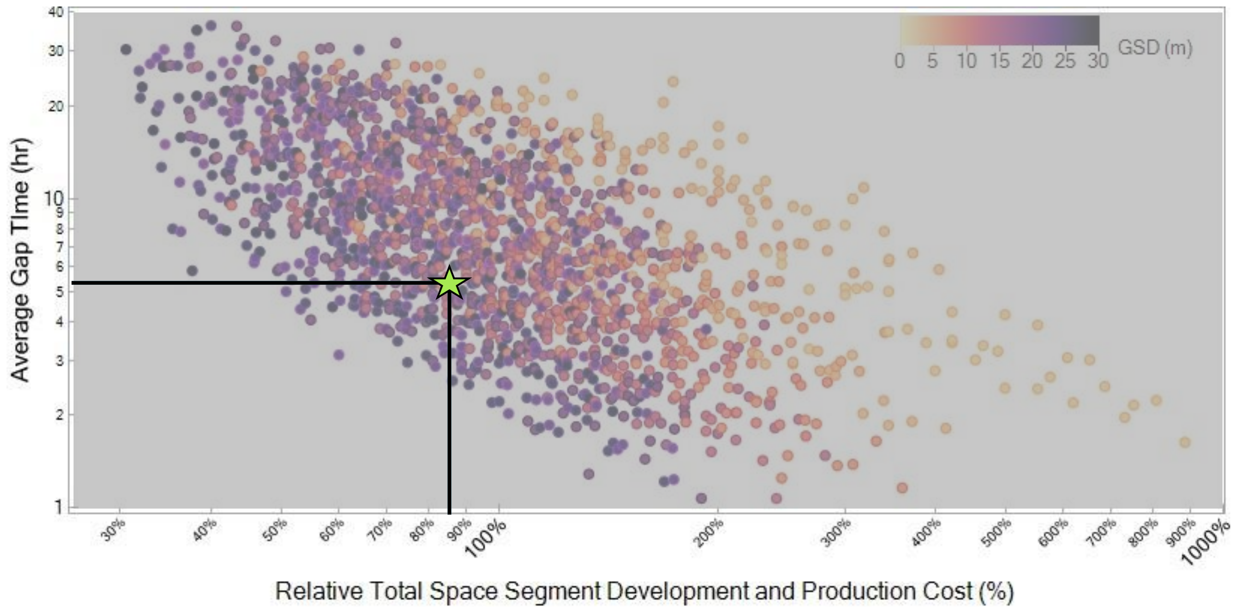

Figure 10: Optimal Design Concept for Cost and Performance

Now that the optimal solution for cost given the mission requirements has been identified and verified, it can be plotted in context with all the data gathered during this process. Figure 10 shows Concept 2 solution on the relative cost vs average gap time colored by GSD visualization previously shown in Figure 5.

6.0 Conclusion

The accelerating complexity within the space mission architecture landscape necessitates the development of next generation tools, practices and procedures that can forecast the downstream impacts to cost and performance from design choices. The data-driven approach presented in this paper hinges on the use of an integrated cost and engineering modeling framework that enables the reasonable development of large and well linked datasets from which higher-level models can be built. Only by leveraging this approach can the distillation of design space sensitives to cost within complex system-of-systems interactions become achievable. The capabilities demonstrated herein provide the ability to begin designs with a meaningful understanding of the entire design space, project where ideal costs should lie given certain requirements, and perform on-the-fly "what-if analyses" to not only mission architects and engineers, but stakeholders and customers.

Though the example scenario formulated in this paper is simple for illustrative purposes, countless more complex and wider ranging requests, trades, and missions may be imagined for which this data-driven approach may prove relevant . From generating cost advantaged design starting points at program/study onset, to tracking downstream cost impacts on choices made, to evaluating and designing requirements to minimize cost, the methodologies showcased in this paper satisfy those needs.

Acknowledgments

We would like to thank BAE Systems for their support in the development of this paper. We also express our thanks and appreciation to Joseph Hangge, Kristopher Sorensen, Brandon Sheridan, Briana LeBlanc, Michael Yantz, and the BAE Space & Mission Systems teams who made significant contributions to our data normalization process and the development of our CERs.

References

1. Barros, P., Kirby, M., and Marvis, D., “Impact of Sampling Technique Selection on the Creation of Response Surface Models,” SAE International, Vol. 113, Section 1: Journal of Aerospace, 2004, pp. 1682-1693.
2. Wautlet, A., Thanavaro, G., and Renkoski, B., “Surrogate-Based SmallSat Architecture Design Leveraging Integrated Parametric Mission Models,” *Proceedings of the AIAA/USU Conference on Small Satellites*, Mission at Scale, SSC23-P4-29. https://digitalcommons.usu.edu/smallsat/2023/all2023/208/
3. International Cost Estimating & Analysis Association (ICEAA), “Cost Estimating Body of Knowledge (CEBok)”. https://wikidev.iceaaonline.com/