Enhanced Heuristic Algorithm for Optimal Cislunar Space Situational Awareness Architectures

Jacob A. Dahlke and Robert A. Bettinger
Department of Aeronautics and Astronautics, Air Force Institute of Technology
Rachel Oliver
Department of Aeronautics and Astronautics, Air Force Institute of Technology

ABSTRACT

The goal of this study is to utilize heuristic optimization techniques, such as Genetic Algorithms, to examine near-optimal space-based sensor architectures within the cislunar environment for the space situational awareness (SSA) mission. Specifically, this study introduces an adapted heuristic algorithm for optimizing cislunar SSA architectures. The algorithm incorporates a categorical variable for selecting cislunar families' periodic orbits, leading to shorter chromosome lengths and enhanced satellite capacity within a single family at reduced costs. Additionally, the algorithm allows for the removal of hidden genes, facilitating calculations for set-size architectures. Evaluation of the algorithm's efficacy will be evaluated by assessing potential benefits such as reduced runtime and improved solution quality. The optimization problem will seek to analyze the trade-offs between SSA capability, station keeping costs, and number of satellites in a given constellation. These objectives are competing for resources within the optimization problem, so a Pareto front of optimal solutions depicting the cost of favoring a given objective will be provided. Trajectory maintenance requirements in terms of stability, as well as various satellite constellation designs utilizing different periodic orbit designs (e.g., Halo, Lyapunov, distant retrograde orbits) will be investigated to provide a baseline assessment of SSA functionality in the volume of space extending from geosynchronous Earth orbit to the Moon and beyond into translunar space. A "cloud of point" target deck will be employed covering the Earth-Moon corridor and the Earth-Moon L1, L2, L4, and L5 points. Visual magnitude SSA metric will be used to evaluate the detectability of the target deck by the architectures evaluated by the Genetic Algorithm, The research will use the circular-restricted three-body problem (CR3BP) as the foundational dynamical model. The research will comprise three fundamental components: (1) generation of a catalog of three-body periodic orbits in excess of 1,000 candidate trajectories; (2) analysis of SSA models and design criteria; and (3) formulation, simulation, and evaluation of SSA architectures comprising sensors in cislunar periodic orbits utilizing the aforementioned heuristic search algorithm.

1. INTRODUCTION

Optimal design is essential for developing effective architectures under financial constraints, such as those faced by the space situational awareness (SSA) community. The U.S. Space Surveillance Network (SSN) tracks over 47,000 near-Earth space objects. As the focus expands to the cislunar domain, where the required volumetric coverage is significantly greater, optimizing SSA architectures becomes increasingly important, particularly with the anticipated growth in cislunar missions [19, 18]. While cislunar SSA research is still in its infancy compared to near-Earth studies, initial optimization efforts are underway, often leveraging heuristic search methods to manage the complexity of these problems.

Heuristic approaches, particularly those involving sensor placement, have been widely explored in SSA optimization. Recent studies have applied advanced algorithms like genetic algorithms (GAs), deep reinforcement learning (DRL), and Monte Carlo tree searches to optimize SSA architectures, focusing on maximizing coverage while minimizing costs through various metrics [16, 26, 12, 19, 11, 4, 20, 25, 2]. Two notable studies by Fahrner et al. [11] and Visonneau et al. [25] specifically address the challenge of optimizing coverage along the Earth-Moon corridor, a region that is more frequently trafficked and difficult to observe due to Moon's brightness. Fahrner et al. [11] developed a multi-step process to optimize cislunar SSA architecture by maximizing coverage with minimal sensors, combining exhaustive

search, greedy algorithms, and genetic algorithms. Visonneau et al. [25] employed a hidden-genes Non-Dominated Sorting Genetic Algorithm-II (NSGA-II) to allow for a varied number of satellites in the architecture while balancing observability and cost.

Despite the significant progress made by recent studies on cislunar SSA architecture development, such as those by Fahrner et al. [11] and Visonneau et al. [25], gaps remain in developing optimization methodologies that are both scalable and comprehensive for SSA architectures. Expanding the design space to encompass a wide range of constellations, varying numbers of satellites, and diverse sensor characteristics requires a flexible approach that can also accommodate future extensions like sensor tasking and multi-sensor architectures without limiting the potential design space. Current research often addresses individual components, but fails to integrate them into a single, cohesive study. To address these challenges, the present research evaluates a heuristic similar to one used by Visonneau et al. [25], termed the "Non-Categorical" heuristic, and introduces an enhanced "Categorical" heuristic for more scalable and integrated optimization in an expanded design space. Both heuristics are demonstrated and compared in a simplified scenario for cislunar SSA architecture development, with future work demonstrating the "Categorical" heuristic in an expanded design space.

The research is organized as follows: Section 2 provides background information on the Circular Restricted Three-Body Problem (CR3BP) and the SSA models used in this study. It also discusses key concepts in optimization and heuristic analysis, with a specific focus on the hidden-genes NSGA-II. Section 3 reviews the Non-Categorical heuristic from the literature, followed by an explanation of the objective functions, GA settings, and evaluation plan for the two heuristic forms. Section 4 details the enhancement of the Non-Categorical heuristic, leading to the development of a Categorical heuristic that addresses its limitations. The resulting Pareto fronts are presented to evaluate these heuristics, along with a discussion of the two forms. Finally, Section 5 summarizes the chapter and highlights the key takeaways from the development of the new heuristic form.

2. BACKGROUND

2.1 The Circular Restricted Three-Body Problem

The cislunar domain is dominated by the gravitational forces of the Earth and the Moon. To model this environment, the dynamics from the CR3BP are used. The CR3BP describes the motion of a body of assumed negligible mass under the influence of two primary bodies—Earth and Moon in this context—assuming that these primaries are in circular motion around their common barycenter. The CR3BP involves several key assumptions and simplifications: the primaries are assumed to be in circular orbits, units for length, mass, and time are nondimensionalized, and the reference frame is selected to be the synodic frame, which rotates with the primaries.

In the synodic reference frame, the primary bodies remain stationary, with the Earth and Moon located along the x-axis at nondimensional distances of $(\mu,0)$ and $(1-\mu,0)$ from the barycenter, respectively. The z-axis is perpendicular to the plane of motion of the primaries in the direction of their angular velocity, while the y-axis completes the orthonormal frame, as shown in Fig. 1 [3]. The rotating reference frame allows for the emergence of unique dynamical structures, such as stationary equilibrium points known as Lagrange points [3], illustrated in Fig. 2. The non-dimensional equations of motion that describe the motion of a satellite within this framework are given in Eqs. 1-3.

$$\ddot{x} - 2\dot{y} - x = -\frac{(1-\mu)(x+\mu)}{r_1^3} - \frac{\mu(x-1+\mu)}{r_2^3}, \tag{1}$$

$$\ddot{y} + 2\dot{x} - y = -\frac{(1-\mu)y}{r_1^3} - \frac{\mu y}{r_2^3}, \tag{2}$$

$$\ddot{z} = -\frac{(1-\mu)z}{r_1^3} - \frac{\mu z}{r_2^3} \tag{3}$$

where $\mathbf{r}_1$ and $\mathbf{r}_2$ are the nondimensional positions of the third body with respect to each primary, respectively, as follows:

$$\mathbf{r}_1 = (x+\mu)\hat{x} + y\hat{y} + z\hat{z} \tag{4}$$

$$\mathbf{r}_2 = (x-1+\mu)\hat{x} + y\hat{y} + z\hat{z} \tag{5}$$

The CR3BP has no closed-form solution, making it challenging to determine the long-term behavior of trajectories. Periodic orbits, however, are "repeating natural trajectories" [15] where the "same configuration is repeated at regular intervals of time" [21]. Due to this regular, repeating nature, the long-term motion of periodic orbits is well understood, making them extremely valuable for gaining insight into the problem and for mission planning, especially for SSA. The repeating trajectory allows for continual tracking and monitoring of a specific region of space. Within the CR3BP, periodic orbits can be categorized into families or types based on shared characteristics, such as stability and location.

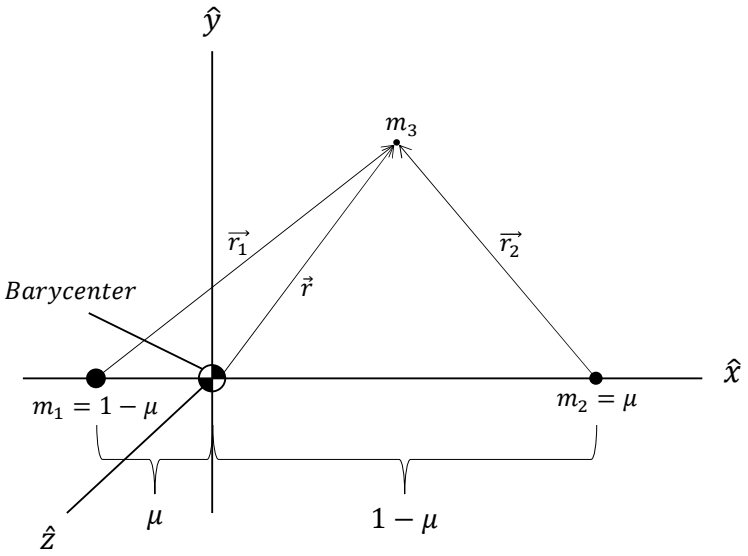

Fig. 1: The CR3BP reference frame

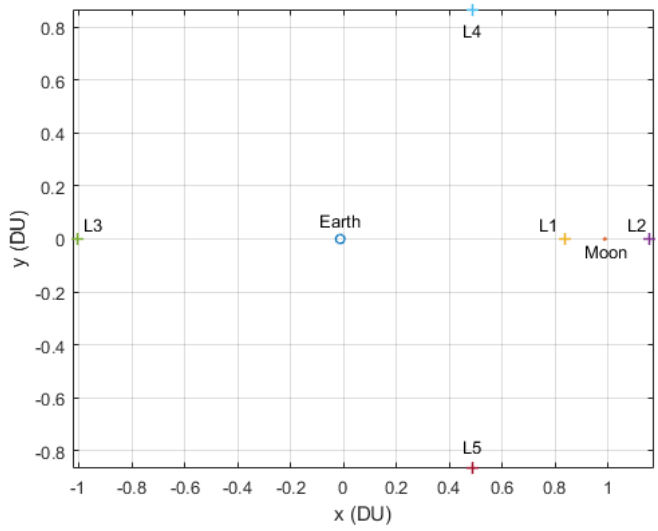

Fig. 2: Equilibrium points in the Earth-Moon system

2.2 Space Situational Awareness Model

Space Situational Awareness (SSA) involves the detection and tracking of objects in space, ensuring the safety and sustainability of space activities. Visual magnitude, a measure of an object's brightness relative to the star Vega, is used to determine the detectability of space objects. It is expressed on a reverse logarithmic scale, where lower values indicate brighter objects. The visual magnitude M_v of a target can be calculated using the methodologies provided by Krag [17] and Coder et al. [5] as:

$$M_v = M_{v,Sun} - 2.5\log_{10}\left[\frac{A_{Tar}}{r_{\text{Tar/Obs}}^2}\left(a_{spec}\rho_{spec}(\phi) + a_{diff}\rho_{diff}(\phi)\right)\right] \tag{6}$$

where:

- $M_{v,Sun}$ is the apparent visual magnitude of the Sun, equal to -26.74 at a distance of approximately 1 astronomical unit (AU) from the Sun.
- $\mathbf{r}_{\text{Tar/Obs}}$ is the vector position of the target with respect to the observer sensor satellite.
- A_{Tar} is the surface area of the illuminated target.
- a_{spec} and a_{diff} are the coefficients of specular and diffuse reflection, respectively.
- ρ_{spec} and ρ_{diff} are the specular and diffuse phase functions, which account for these two components of reflectivity and are functions of the solar phase angle ϕ.

The solar phase angle ϕ determines the direction of incident light on the target from the Sun with respect to the observer's position, is calculated as:

$$\phi = \arccos\left(\frac{\mathbf{r}_{\text{Tar/Obs}} \cdot \mathbf{r}_{\text{Tar/Sun}}}{r_{\text{Tar/Obs}} r_{\text{Tar/Sun}}}\right) \tag{7}$$

where $\mathbf{r}_{\text{Tar/Sun}}$ is the vector position of the target with respect to the Sun.

Krag [17] provides tabulated phase functions for targets with simple geometries. For a spherical target, the phase functions for specular and diffuse reflection are given by:

$$\rho_{spec} = \frac{1}{4\pi} \tag{8}$$

$$\rho_{diff} = \frac{2}{3\pi^2}\left(\sin\phi + (\pi - \phi)\cos\phi\right) \tag{9}$$

A target is considered detectable if its visual magnitude M_v is less than or equal to a designated detection threshold $M_{v,det}$, meaning $M_v \leq M_{v,det}$.

Exclusion regions are also factored into detectability considerations to account for physical limitations in observation. Following Vendl's methodology [24], these exclusion regions are categorized into two types: line of sight and illumination constraints. Line of sight constraints occur when the target is either obstructed by or visually overwhelmed by the brightness of the Earth, Moon, or Sun. Illumination constraints arise when the target is in the shadow of the Earth or Moon.

For space-based observers operating in the cislunar region, exclusion angles must be dynamically calculated because the apparent sizes of the Earth and Moon change with distance. In contrast, the Sun's exclusion angle remains constant due to its vast distance from this region.

To evaluate the performance of an SSA architecture, the observation accessibility metric Γ_{Mv}, developed by Vendl [24], is employed. This metric quantifies the percentage of targets that are visible based on their visual magnitude over a significant portion of the observation period, serving as a proxy for the ability to consistently track objects. The metric is calculated by determining the proportion of observation periods during which each target is detectable and then calculating the percentage of targets that meet a specified visibility threshold. Typically, this threshold is set at 90%, meaning a target must be visible for at least 90% of the observation periods to be considered accessible.

A target is considered accessible if it is detectable for at least 90% of the observation periods, expressed as follows:

$$q_i = \begin{cases} 1 & \text{if } \frac{N_{det,i}}{N_{obs}} \geq 0.9 \\ 0 & \text{if } \frac{N_{det,i}}{N_{obs}} < 0.9 \end{cases} \tag{10}$$

where $N_{det,i}$ is the number of detections for the i-th target, and N_{obs} is the total number of observation periods. The total number of accessible targets $N_{accessible}$ is then the sum of all accessible targets:

$$N_{accessible} = \sum_{i=1}^{N_{targets}} q_i \tag{11}$$

Finally, the observation accessibility metric Γ_{Mv} is expressed as:

$$\Gamma_{Mv} = \frac{N_{accessible}}{N_{targets}} \cdot 100 \tag{12}$$

2.3 Hidden Genes Non-dominated Sorting Genetic Algorithm II

GAs are optimization techniques inspired by the process of natural evolution. GAs work by generating a population of potential solutions, evaluating their performance using a fitness function, and then using crossover (mixing of solutions) and mutation (random changes) to evolve the solutions over several generations, as depicted in Fig. 3. The goal is to find the best, or more accurately, a "near"-optimal solution to the problem at hand. The term "near" optimal is used because GAs are heuristic search techniques, which are stochastic in nature. While these stochastic operations help escape local extrema by expanding the search space, they do not guarantee finding the global optimum. Instead, they offer a strong likelihood of getting close to an optimal solution, hence the term "near" optimal [1, 10]. The process continues through successive generations until a termination criterion is met, which is designed to balance the need to avoid premature termination, which could result in suboptimal solutions, with the need to avoid excessive computation time when further improvements are unlikely [6].

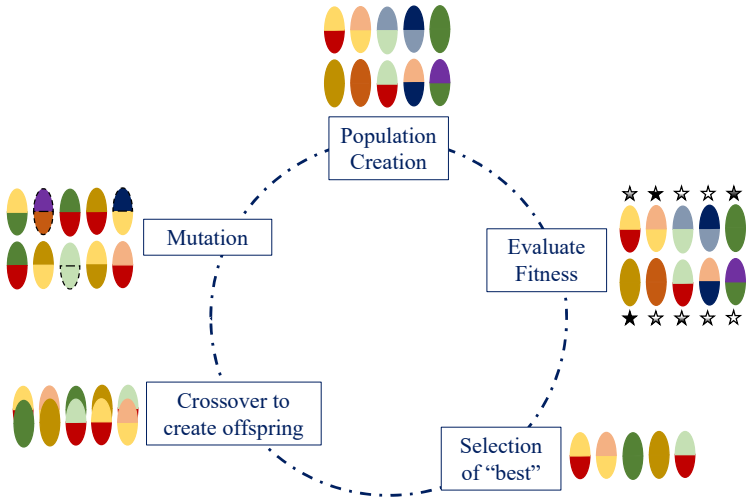

Fig. 3: Genetic algorithm depiction

In some optimization problems, the number of design variables is not fixed and can vary depending on the solution structure—this is known as a Variables-Size Design Space (VSDS). Standard GAs, which typically operate with a fixed-length chromosome, face challenges when the number of variables changes. To address this challenge, the concept of hidden genes is introduced. Hidden genes are part of the chromosome but do not directly influence the fitness evaluation. By marking genes as active or inactive, the GA can maintain a fixed chromosome length while effectively optimizing problems with varying numbers of design variables [7]. With this approach, all chromosomes can have the same length, enabling standard genetic operations without the added complexity of variable length, though all chromosomes must contain the maximum number of possible design variables, even if not all are used in a given solution [13, 25].

The NSGA-II is a popular algorithm for multi-objective optimization, where multiple conflicting objectives must be optimized simultaneously. Similar to a basic GA, each objective function is evaluated, but in NSGA-II, each chromosome is also sorted into non-domination levels [9]. Optimal points represent a compromise between different

objectives, where any gain in one objective function (e.g., f_1) comes at the cost of another (e.g., f_2), allowing for a subset of "efficient" solutions to be found rather than considering every possible outcome [5]. The subset of feasible solutions, **x**, where this trade-off occurs is called a "Pareto front" (also known as the Pareto Frontier or Pareto Curve), with each point along the curve termed a Pareto optimal solution [5]. Figure 4 illustrates the tradeoff between objective functions along the Pareto front. In this figure, "dominated" solutions are not Pareto efficient, meaning one objective function could be improved without affecting the other objective function [14].

Additionally, the NSGA-II uses a crowding distance metric to maintain diversity in the population and an elite-preserving operator to ensure that the best solutions are retained through successive generations. By incorporating hidden genes, NSGA-II can handle VSDS problems efficiently, allowing for the optimization of complex design spaces with variable numbers of design variables while effectively balancing multiple objectives. The hidden-genes NSGA-II undergoes successive generations of evolution until a termination condition is reached, as in standard GAs, but with the added capability to handle multiple objectives and hidden genes to accommodate a VSDS.

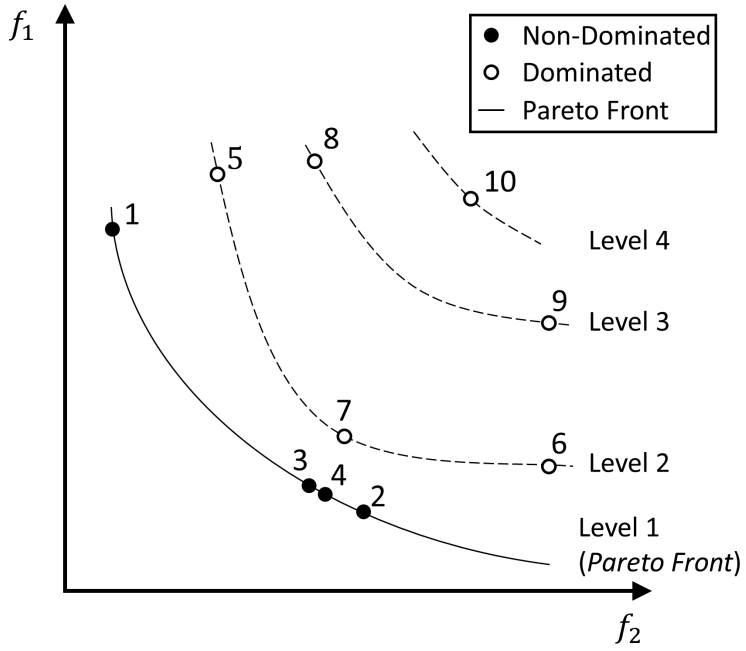

Fig. 4: Pareto front and non-domination levels (adapted from Deb [9])

3. METHODOLOGY

3.1 Non-Categorical Chromosome Form

Utilizing a GA requires the desired system to be representable in the form of a chromosome. A chromosome represents a possible solution to the problem; not necessarily a good or bad solution, but simply a possible one. The GA is then used to identify good solutions and ideally arrive at a "near"-optimal solution. In the present study, one chromosome represents an architecture of satellites that could be used for cislunar SSA. There are numerous ways that an architecture of satellites could be represented in chromosome form, but representing it effectively is part of the challenge in using a GA.

Visonneau et al. [25] developed an NSGA-II to solve a VSDS problem in the cislunar regime, aimed at developing multi-satellite architectures for SSA. The form developed by Visonneau et al. [25] is referred to in the present study as the "Non-Categorical" form of the heuristic. The Non-Categorical form involves creating a chromosome with sets of two genes used to describe a possible observer satellite. These genes, also called design variables, include: 1) selection of an orbit within a family by identifying a unique period, and 2) selection of the phasing of the observer within that orbit. Each of these sets of genes representing an observer also contains a tag indicating whether the observer is "On" (active) or "Off" (inactive). If the observer is "Off" (inactive), then the given satellite is "hidden" and

is not used to calculate performance. Essentially, this is as if the satellite is not in the architecture at all. However, as the genetic algorithm proceeds, the gene may be flipped to a different state through mutations, so it is still tracked to retain evolutionary information. The Non-Categorical form of the heuristic is presented in Figure 5. Each set of two blue squares represents the two genes needed to describe a possible observer sensor. The green box over the set of genes is the tag indicating whether the observer satellite is hidden or not. There are i satellites allowed in each family and n total families. Additionally, there are C chromosomes, with each representing an architecture of satellites that could be used for cislunar SSA. All possible design parameters must be included in each chromosome when using the NSGA-II, so all of the same parameters are incorporated into all C chromosomes in Figure 5.

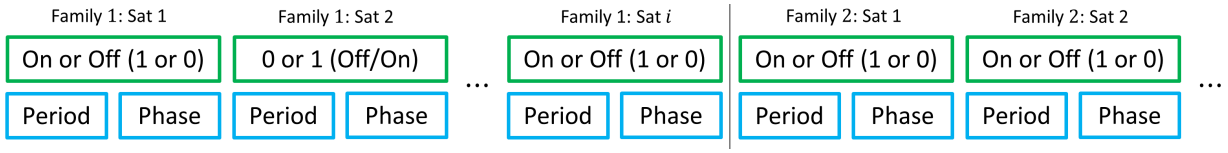

(a) Initial section

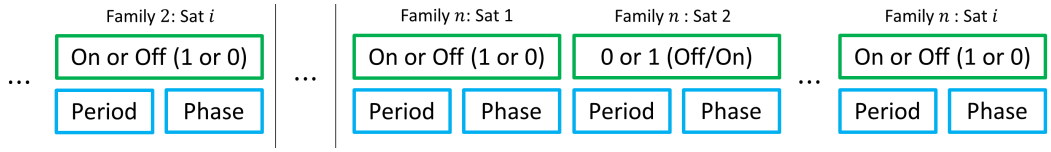

(b) Final section

Fig. 5: Non-Categorical heuristic form with (a) the initial section and (b) the final section of chromosome population

The inherent shortcoming of the Non-Categorical form of the heuristic is its scalability:

1. **Size of Chromosome**

   - Any time a potential new family of orbits wants to be evaluated, all possible observer satellites with their 2 design variable descriptors that can be selected from that family need to be appended to the end of the chromosome. This means that if i possible observers are allowed in that family, then $i \cdot DV$ genes need to be added to the chromosome, where DV is the number of design variables (2 in this instance). Each additional family creates a significantly larger chromosome.
   - If any additional design variables are desired to describe the observer satellite (for instance, some of its sensor characteristics), then an additional gene needs to be added for every possible satellite in every possible family. If there are n possible families, $n \cdot i$ genes need to be added for each additional design variable added to the problem, creating a significantly larger chromosome.

   The maximum length of a chromosome is given by:

   $$\text{Length}_{\text{Max}} = n \cdot i \cdot DV \tag{13}$$

   The maximum length of the chromosome is for each chromosome in the population, and there are C total chromosomes in a population, so the size of recorded information is even larger. The length of the chromosome could become an issue at very large sizes, reducing the performance of the heuristic as all of the satellite information in each chromosome will still need to be tracked even if satellite tags are off in each generation and throughout all of the chromosomes. The rapid increases in chromosome lengths from minor changes and additional characteristics is a concern for the scalability of the problem to a very large design space.

2. **Number of Possible Observer Orbits**
   The number of possible observer orbits is inherently limited in this formulation because the period in a family of orbits is not always unique. This means that the possible observer orbits selected for evaluation must be restricted to a range of orbits within each family that is unique. For example, within a family, no two orbits can have a period of 2 TU, or the heuristic will not know which one to select, so any orbits with repeated periods must be removed.

The Non-Categorical form of the heuristic will be used as the baseline for comparison in the present study. Any improvements in the heuristic form will be measured relative to the performance of the Non-Categorical form. By establishing this baseline, the effectiveness of proposed modifications and enhancements can be quantitatively evaluated. This approach ensures that the assessment of new heuristic forms is grounded in a clear, objective framework. Additionally, using the Non-Categorical form as a benchmark allows us to systematically identify areas where scalability and efficiency can be improved, thereby advancing the overall methodology for designing multi-satellite architectures in the cislunar regime. Through this comparative analysis, the goal is to develop a more robust and efficient heuristic that can better handle the complexities of large-scale design spaces and diverse satellite architectures.

3.2 Objective Functions

This analysis utilizes a multi-objective approach to compare three distinct objectives, resulting in the generation of a Pareto front. While the objectives are not weighted, the primary emphasis is placed on the architecture's effectiveness in performing SSA. The metric used to assess SSA performance in this analysis is the Observation Accessibility Metric, Γ_{Mv}.

The second objective function seeks to minimize the total number of satellites within the architecture. This objective is based on the premise that a larger number of satellites will incur higher costs, making this metric a proxy for minimizing the overall cost of the architecture.

The third objective function focuses on the stability of periodic orbits. Stability in this context refers to whether a perturbed object, such as a satellite, will return to or deviate from its orbit, as well as the speed at which this process occurs. Higher stability indices correspond to faster departures from the orbit when perturbed. Given that the CR3BP is an approximation that does not account for all real-world perturbations, orbit stability is a reasonable proxy for estimating the station-keeping efforts required to maintain a spacecraft in its orbit. For example, Davis et al. [8] found that for Halo and Butterfly orbits, station-keeping costs "generally scale with the stability index" [8]. Therefore, minimizing the stability index aims to reduce station-keeping costs.

In the Pareto front plots, stability values are represented by a colorbar. While most architectures exhibit nominal stability values, the colorbar is capped at displaying stability values up to 5,000 to prevent excessively high values from distorting the color distribution in the plot. In this scheme, yellow indicates stability values that may exceed 5,000, whereas orange and lower colors provide meaningful insights into the stability of the architecture. This approach ensures that very high stability values do not skew the color representation, allowing for a more accurate analysis of the stability across typical architectures.

3.3 Genetic Algorithm Configuration and Settings

MATLAB's `gamultiobj` function was used to carry out the multi-objective optimization, utilizing an NSGA-II [22]. The default crossover and mutation operators were used in this analysis, as detailed in [23], with a termination criteria of 50 generations. The `gamultiobj` process is described here for study-specific clarifications and completeness, as detailed in [22]:

1. An initial population is created using the designated chromosome form with each gene randomly selected. For the Non-Categorical form, a random orbital period and phase are chosen for each observer satellite in the initial chromosome. All design variables in the chromosomes are assigned random values between their lower and upper bounds, as specified in Table 4 for all observer satellites in each chromosome.
2. Each chromosome's performance is evaluated using the objective functions detailed in Section 3.2.
3. Based on each chromosome's performance, they are selected for reproduction using tournament selection of four chromosomes at a time.
4. Children are created through the reproduction process, which includes crossover and mutation:
   (a) **Crossover Operator**: In this study, 80% of the offspring population is created via crossover, while the remaining 20% are retained from the parent population based on their elite performance. The crossover operator used is the default in MATLAB, which calculates each element of the child as a weighted average of the corresponding elements of the parents, with the weight being a random number between 0 and 1, instead of a typical single crossover point. This ensures that the offspring has a mix of characteristics from both parent solutions, determined by the random weights [23].

$$\mathbf{c} = \mathbf{p}_1 + w \cdot (\mathbf{p}_2 - \mathbf{p}_1) \tag{14}$$

where $\mathbf{c}$ represents the child chromosome, $\mathbf{p}_1$ represents the first parent chromosome, $\mathbf{p}_2$ represents the second parent chromosome, and w represents the weight.

(b) **Mutation Operator**: Genes are initially randomly selected for mutation and mutated by a random amount in a random direction with a random step length [23]. After this initial mutation, future mutations are adapted based on the success of the new population. If the last mutation resulted in a better solution, then the function might use a similar direction or step length. Conversely, if the last mutation was unsuccessful, the function might choose a different direction or adjust the step length.

(c) **Hidden-Genes Tags Operator**: Tags indicating which genes should be hidden can undergo their own mutation and crossover. In this study, these tags were treated as normal design variables, undergoing the same crossover and mutation processes as the other design variables.

5. Each offspring's performance is evaluated using the objective functions.

6. An extended population is created by combining all parents and newly created offspring.

7. The extended population is reduced back to the original population size by retaining 35% of the chromosomes on the Pareto front, if possible, with the remaining population being selected from other fronts. Crowding distance is used as a tie-breaker for selecting chromosomes on the same front.

8. The process is repeated using this newly created population as the initial population for the next generation until the termination criteria of 50 generations is met, with the settings provided in Table 1:

Table 1: Heuristic and chromosome properties

Parameter	Value	Notes
Population size	100	*The number of chromosomes per generation*
Maximum number of generations	50	—
Maximum number of satellites allowed in architecture	15	*By default, this sets $i = 3$ in the Non-Categorical form, i.e., the maximum number of observers in each family*

3.4 Evaluation Plan and Scenario Setup

A new form of the Non-Categorical Heuristic, called the Categorical form, will be introduced to address its shortcomings. Section 4 will demonstrate the changes, improvements, and inherent benefits of the new form. To evaluate the Categorical form's efficacy in cislunar SSA architecture development, both the Non-Categorical and Categorical forms of the heuristic will undergo the same testing for comparison to determine if there are other possible advantages the Categorical form has beyond just the inherent benefits it has from an enhanced form. Identical scenarios will be used for both heuristics. The results will be compared by examining the Pareto fronts to assess overall performance in each objective and by evaluating the total run-time of the heuristics to determine efficiency.

The analysis will replicate the scenario used by Visonneau et al. [25]. The region of interest for SSA is the cone of shame, which extends from Earth to L_2 and includes targets that are difficult to observe from Earth due to the brightness of the Moon. This region will be modeled as a frustum extending from Earth to geosynchronous Earth orbit (GEO) with a half-angle of 20 degrees. A total of 1,112 targets will be equally spaced within this frustum, as depicted in Figure 6a. The possible orbit families for placing observer satellites include L_1 Lyapunov, L_2 Lyapunov, Distant Retrograde Orbits, L_1 Southern Halo, and L_2 Northern Halo, shown in Figure 6b. The overlay of the target deck on these orbit families is presented in Figure 6c.

To assess performance across a range of scenarios, two visual magnitude detection thresholds (16 and 20) will be used with a termination criteria of limiting the maximum number of generations to 50. The settings used to determine detectability through visual magnitude are presented in Table 2. Additional scenario settings and paramters are provided

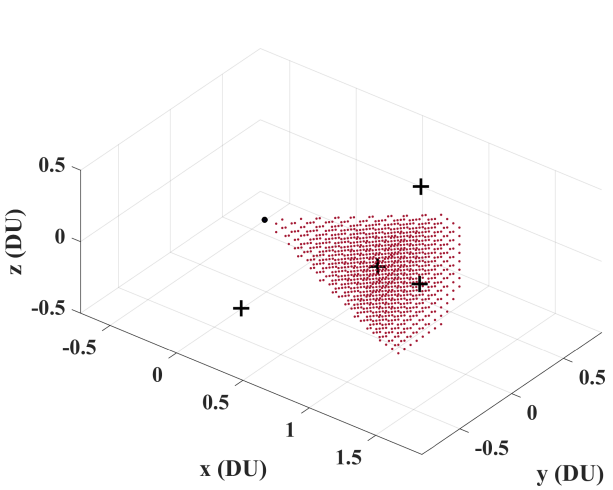

(a) Target deck: Cone of shame

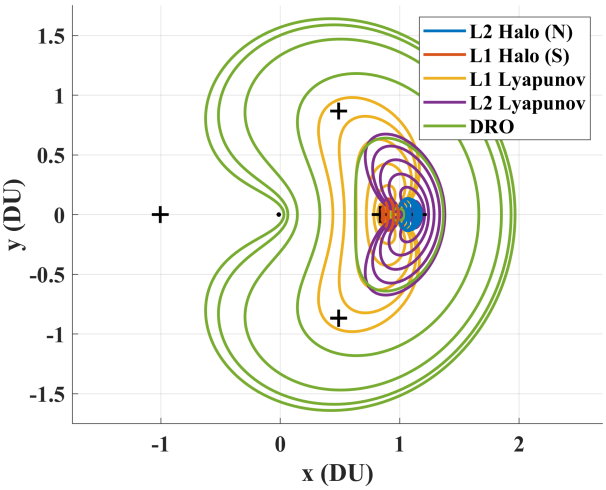

(b) Observer orbits families

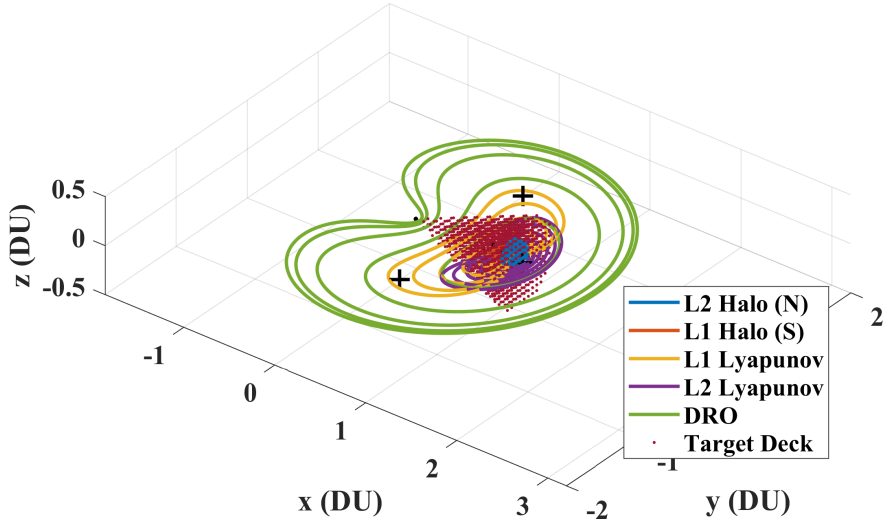

(c) Target deck overlay on observer orbits families

Fig. 6: Observer families and target deck used to evaluate Categorical form of heuristic against Non-Categorical form

in Table 3. The analysis will focus on performance metrics for SSA, the minimization of satellite count or cost, and the stability index, as described in Section 3.2. This comprehensive evaluation will provide insights into the effectiveness and efficiency of the enhanced Categorical heuristic form.

Table 2: Simulation parameters for visual magnitude evaluation

Variable	Parameter	Value
$M_{v,Sun}$	Visual Magnitude of the Sun at 1 AU	-26.74
Target Shape	(*For Phase Function*)	Spherical
d_{Tar}	Diameter of Target	1 meter
a_{spec}	Coefficient of Specular Reflection	0
a_{diff}	Coefficient of Diffuse Reflection	0.2
Sun_{excl}	Exclusion Angle of the Sun	10 degrees

Table 3: Scenario settings and parameters

Setting/Parameter	Value	Notes
Orbit Families	L_1 Lyapunov, L_2 Lyapunov, Distant Retrograde, L_1 Southern Halo, L_2 Northern Halo	—
Region of Interest (Target Deck)	Cone of Shame	*Frustum with 20-degree angle extending from GEO to L_2*
$M_{v,\text{det}}$	Multiple: 16 and 20	—
Start Date	March 25, 2024, 07:12:36.633	*Date when Earth-Moon plane is aligned with the ecliptic plane of the Sun-Earth [25]*
Simulation Duration	1825 days (5 years)	—
Time Steps	2 days	—

4. RESULTS AND ANALYSIS

4.1 Categorical Chromosome Form

To address the inherent shortcomings of the Non-Categorical form of the heuristic, a new form called the Categorical form has been developed. The primary issue with the Non-Categorical form is the rapid increase in chromosome length, which negatively affects scalability. To mitigate this, a categorical variable representing the family is introduced as a design variable. In this context, a categorical variable is assigned a distinct value to represent each member of a group, allowing each family to be selected by the design variable.

Additionally, the design variable "period" is replaced with a new variable called the Fractional Orbital Index (FOI). This change resolves a significant challenge in the crossover process: different families of orbits may have varying ranges of periods, leading to inconsistencies during crossover. Specifically, a period selected from one family might not correspond to a valid orbit in another family, causing problems in the solution representation.

To overcome this issue, the FOI is introduced. The FOI is determined by indexing all the orbits within a family and then converting each index to a fractional value between 0 and 1. This conversion is performed using the following equation:

$$\text{Fractional} = \frac{\text{Current Index} - \text{Min Index}}{\text{Total Range}} = \frac{\text{Current Index} - \text{Min Index}}{\text{Max Index} - \text{Min Index}} \tag{15}$$

To pull out a specific orbit from a family using the FOI, the reverse conversion is done:

$$\text{Current Index} = \text{Fractional} \cdot (\text{Max Index} - \text{Min Index}) + \text{Min Index} \tag{16}$$

With the categorical variable for family and FOI for period, the heuristic is illustrated in Figure 7. The three descriptors of the observer, shown in the blue rectangles, are the categorical variable for family, the FOI for selecting the orbit within the designated family, and the phase. The tags indicating whether an observer is active or inactive are shown in the green rectangle above the three descriptors. The variable i represents the number of satellites included in the architecture.

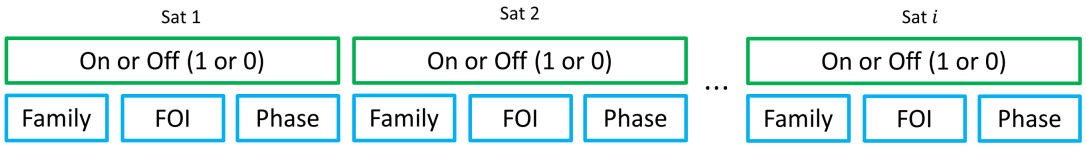

Fig. 7: Categorical heuristic form

The introduction of the categorical variable reduces the length of the chromosome by eliminating the multiplication by the n families, at the cost of increasing the number of design variables by one. This results in a maximum chromosome length of:

$$\text{Length}_{\text{Max}} = i \cdot DV \tag{17}$$

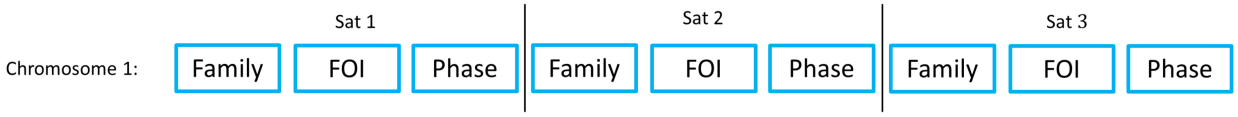

Fig. 8: Categorical heuristic form without hidden-genes tags

In this formulation, *DV* will always be one larger than the *DV* from the Non-Categorical form given in Equation 13. Regardless, for any substantial number of families, design variables, or satellites allowed in each family, the Categorical form will be significantly shorter. Since *DV* cannot be smaller than two in the Non-Categorical form and not smaller than three in the Categorical form, the Non-Categorical form is only shorter when the number of families, n, is equal to one. Otherwise, the Non-Categorical form will be larger by default.

Additionally, this Categorical form adds flexibility to the allowable number of satellites in each family within a chromosome. In the Non-Categorical form, a defined possible number of satellites per family must be selected, which is highly restrictive. For instance, if it is unknown whether an architecture would be best served with 15 satellites in a single family or spread across 5 families, then the Categorical form can simply set $i = 15$ and the heuristic can place satellites in any families, resulting in a $\text{Length}_{\text{Max}} = 45$. For the Non-Categorical form, if $i = 15$, $\text{Length}_{\text{Max}} = 150$, which is significantly higher than the Categorical form. Alternatively, if it tried to match the length of the Categorical form's chromosome, it could only have $i = 5$, restricting to 5 satellites per family. If optimal performance requires all 15 satellites in one family, then the Non-Categorical form would not be able to determine that solution.

The Categorical form is also more adaptable. For example, if it is desired to no longer remain in a VSDS by removing the hidden-genes component of the NSGA-II, then the number of satellites in an architecture can be fixed. This can be useful if the number of observer satellites for SSA is already known or a certain number is desired. Removing hidden-genes in the Categorical form is straightforward because the hidden-genes are only used to vary the number of satellites. By selecting the desired architecture size and creating a chromosome of that size without hidden-genes tags, flexibility is maintained. For instance, a 3-satellite architecture would appear as in Figure 8. In contrast, removing hidden-genes in the Non-Categorical form is not feasible because both the number of spacecraft and the families they are placed in are restricted.

By indexing all the orbits and converting them to fractionals, the following benefits are achieved:

- Families no longer need shared ranges in their descriptors. With FOI, all orbits have a value between 0 and 1, allowing for crossover between any family.
- Fractionals allow the number of orbits in each family to vary, providing greater flexibility and ease when adding new families to the solution space.
- Families with non-unique orbital parameters, such as periods, can still be utilized, removing the restriction of selecting unique ranges of orbits within a family.
- Large segments of orbits can be skipped within each family by re-indexing and ignoring the skipped orbits if desired.

Ultimately, the Categorical form of the heuristic provides significant inherent benefits over the Non-Categorical form, including:

- Shortening the length of the chromosomes, especially in larger design spaces with more families, design variables, or allowed satellites in the architecture.
- Allowing satellites to be placed in any family or all in one without defining a maximum number per family.
- Easily adapting to a non-hidden genes form if the desired number of satellites in the architecture is fixed or known.

The Categorical form of the heuristic, with its introduction of a categorical variable for the family and the FOI, offers a streamlined and flexible approach to chromosome representation. By addressing the limitations of the Non-Categorical form, it reduces chromosome length and enhances adaptability in multi-satellite architecture design. These

improvements facilitate the handling of larger design spaces and more complex satellite configurations. In the next section, practical application of the Categorical form will be explored, comparing its performance against the Non-Categorical heuristic to highlight the advantages and potential areas for further refinement.

4.2 Comparison of Non-Categorical and Categorical Chromosome Forms

To compare the Non-Categorical and Categorical forms of the chromosome, both heuristics were run using two visual magnitude detection thresholds ($M_{v,Det}$) of 16 and 20. These scenarios allowed for multiple analyses to ensure that the performance of each form was consistent. To initialize the heuristics, an initial population was created using the respective chromosome forms—either Non-Categorical or Categorical. Each design variable, or gene, was assigned a random value between its lower and upper bounds, as specified in Table 4 for all observer satellites in the chromosomes.

Table 4: Design variables and their bounds

Design Variable	Lower Bound	Upper Bound
Family (Categorical)	0	5
FOI	0	1
Phase (radians)	0	2π

4.2.1 Run Times

Tables 5 show the total run times. In both chromosome forms, the run times were fairly similar across the visual magnitude detection thresholds. The Non-Categorical form performed better for $M_{v,Det}$ of 20, whereas the Categorical form performed slightly better at $M_{v,Det}$ of 16.

Due to run time being influenced by a variety of factors, such as other computer processes occurring at the same time, it is prone to a degree of fluctuation. Without definitive speed improvements, it is difficult to conclude any significant difference based on run time alone. Neither form had definitively better run times, so there is no distinguishable difference in run times between the two chromosome forms.

Table 5: Run times for Non-Categorical and Categorical forms with 50 generations termination

$M_{v,Det}$	Non-Categorical	Categorical
16	5 hours, 59 minutes	5 hours, 56 minutes
20	3 hours, 18 minutes	4 hours, 7 minutes

4.2.2 Pareto Fronts

By limiting the heuristics to terminate after 50 generations, a direct comparison of the heuristic performance can be made by observing how effectively each form analyzes the scenarios. Fig. 9 and 10 depict the results for the visual magnitude detection thresholds of 16 and 20, respectively, with the Non-Categorical Form shown in the upper subfigure (a) and the Categorical Form in the lower subfigure (b) of each figure.

In Fig. 9, the Categorical Form produces better performing architectures. By comparing fixed "Total Number of Observer Satellites" and examining the performance of the other two objectives, the superiority of the Categorical form becomes even more evident. Consider the highest Γ_{Mv} values at 14 total observer satellites at $M_{v,Det}$ of 16. The Categorical Form has a Γ_{Mv} of 3.42% and a Total Stability Index of 1,333.67, compared to the Non-Categorical form's Γ_{Mv} of 2.70% and a Total Stability Index of 4,402.98. The architecture is illustrated in Figure 11 to highlight these differences. The Categorical Form focused almost exclusively on the L_1 Southern Halo Family, keeping the orbits packed near the Moon and L_1, while the Non-Categorical form chose significantly more expansive orbits, spreading its selection among various possible observer families.

A major disadvantage of the Non-Categorical chromosome form is evident—despite having 14 satellites, only 3 satellites are allowed in each family due to the chromosome's structure. The Categorical form does not have this restriction and thus places 11 of its 14 satellites in the priority L_1 Southern Halo orbit family. For the Non-Categorical form to

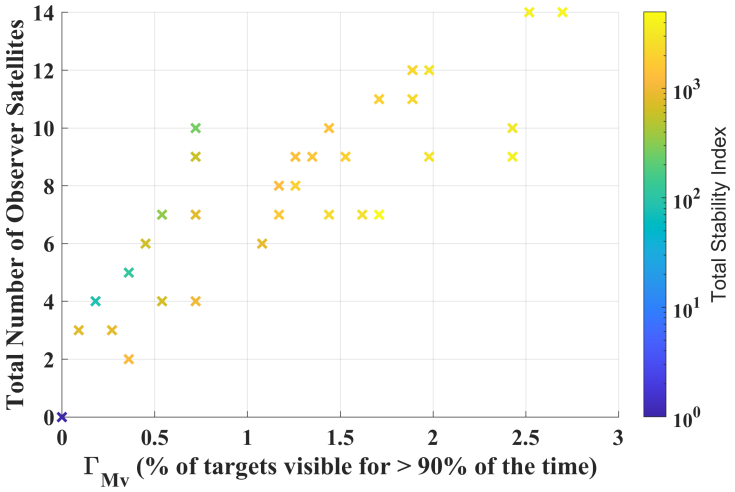

(a) Non-Categorical form at $M_{v,Det}$ of 16

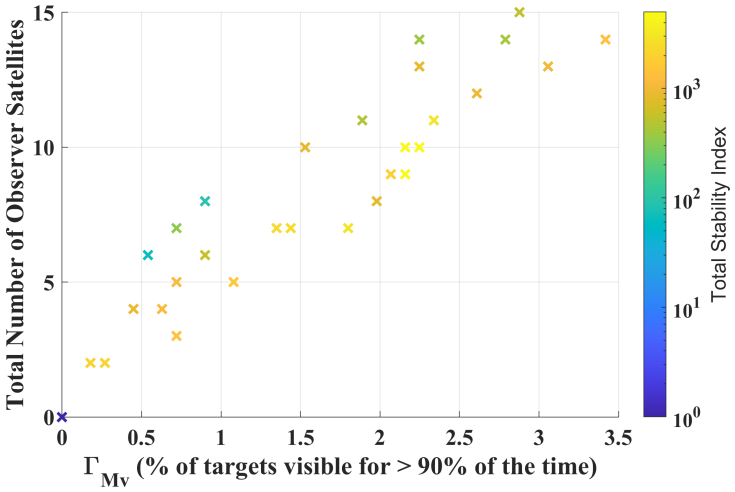

(b) Categorical form at $M_{v,Det}$ of 16

Fig. 9: Comparison of Non-Categorical and Categorical forms at $M_{v,Det}$ of 16

place 11 satellites in the same family, a large expansion of the chromosome would be needed, totaling a length of 110 genes, whereas the Categorical variable achieves this with a length of 45 (and theoretically could do so with only 33). This demonstrates the inherent downsides of the Non-Categorical form's structure and the improvements gained by the Categorical Form, both in structure and actual application.

Both forms exhibit relatively poor performance at an $M_{v,Det}$ of 16. However, the Categorical form demonstrates better SSA observation performance and lower stability for the same number of satellites compared to the Non-Categorical form. Furthermore, the Categorical form was able to explore more of the design space, identifying a 15-satellite architecture on the Pareto front. In contrast, the Non-Categorical form did not discover any 15-satellite architectures.

At $M_{v,Det}$ of 20, similar performance between the heuristic forms is SSA performance occurs. At this visual magnitude detection threshold, the region of interest (i.e., the cone of shame) is almost entirely visible, meaning both chromosome forms can achieve high Γ_{Mv}. Compared to the other $M_{v,Det}$, the stability values are lower across the Pareto fronts, likely because Γ_{Mv} is no longer the limiting factor due to being near 100%, so now optimizing stability and the total number of observer satellites is the priority. The Categorical form outperforms the Non-Categorical form when examining these other objectives. For example, consider the highest Γ_{Mv}-performer of the 2-satellite architectures. In the Categorical form, it has a Γ_{Mv} of 93.98% and a total stability index of 134.26. The Non-Categorical form has a Γ_{Mv} of 92.99%

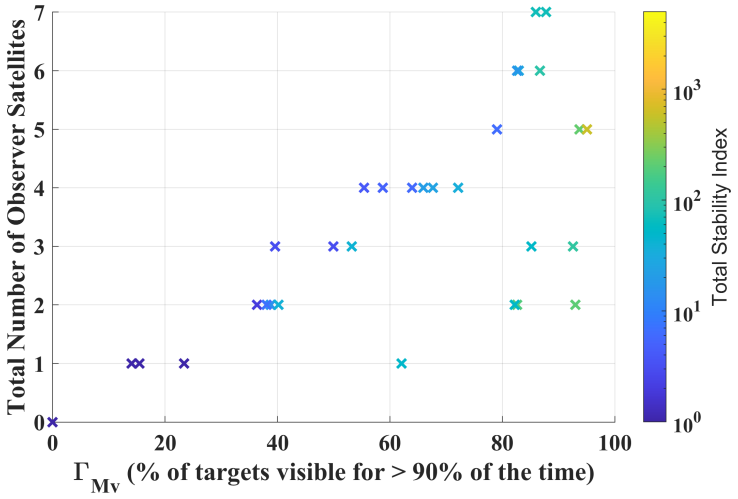

(a) Non-Categorical form at $M_{v,Det}$ of 20

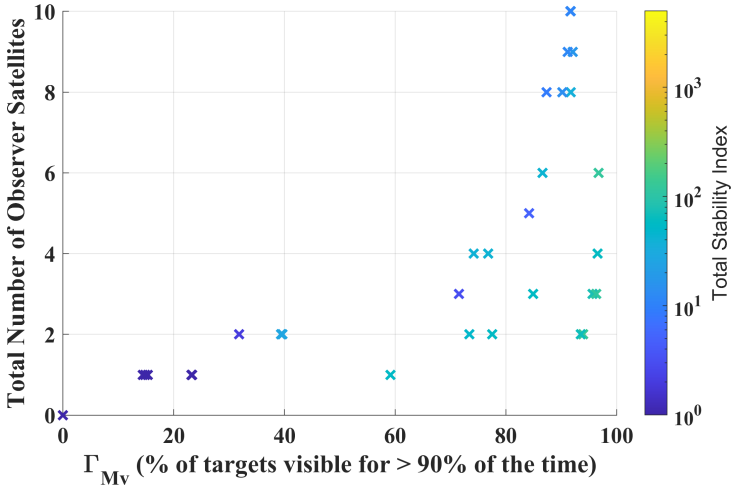

(b) Categorical form at $M_{v,Det}$ of 20

Fig. 10: Comparison of Non-Categorical and Categorical forms at $M_{v,Det}$ of 20

and a total stability index of 210.19, a lower performer in both metrics. The Categorical Form also has another high-performing 2-satellite architecture with a Γ_{Mv} of 93.53% and a stability of 68.29, which further outperforms the Non-Categorical form. This 2-satellite architecture from the Categorical Form and the highest performing 2-satellite from the Non-Categorical form are shown in Figure 12.

In Figure 12, it is evident that both forms of the heuristic found a similar high-performing architecture. Both selected an L_1 Lyapunov orbit of similar period and placed the observer satellite in similar phasing. The key difference between the two architectures lies in the second orbit. The Non-Categorical form selected a planar L_2 Lyapunov orbit, while the Categorical Form selected a spatial L_2 Northern Halo Orbit. Although these two orbits do not have exactly the same planar extent, they are very similar. The fact that they have such similar Γ_{Mv} performance and both architectures share a similar shape suggests that the spatial aspect of the L_2 Northern Halo Orbit may not play a significant role in its selection. Instead, its incredibly low stability is likely the reason for its choice. The L_2 Northern Halo Orbit provides similar, if not slightly better, Γ_{Mv} performance compared to the L_2 Lyapunov, but with a much lower stability index, significantly improving one of the objectives. This indicates that the L_2 Northern Halo Orbit may have been selected more for its planar shape and stability characteristics than its spatial extent. This architecture again demonstrates the Categorical form outperforming the Non-Categorical form, even when both are making minor changes to narrow in on the best architecture, i.e., the heuristics are only swapping out a single orbit to find better performance.

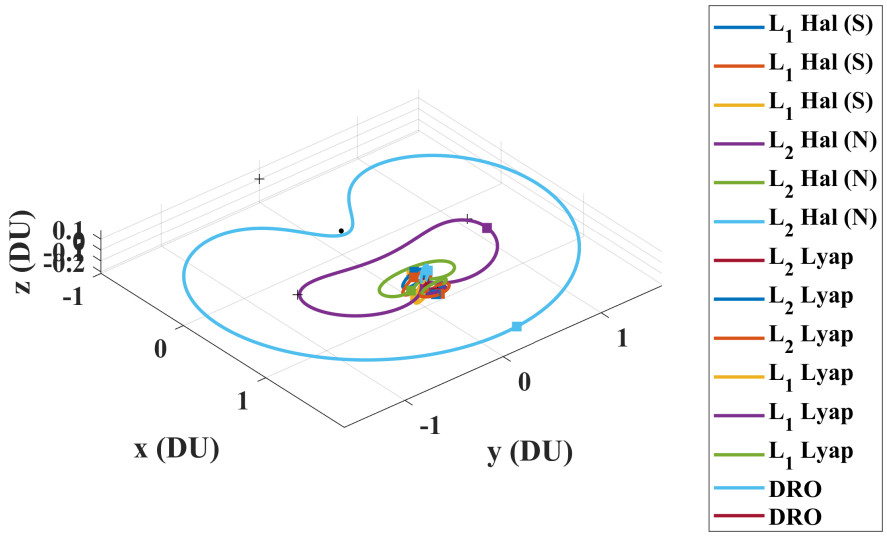

(a) Non-Categorical form

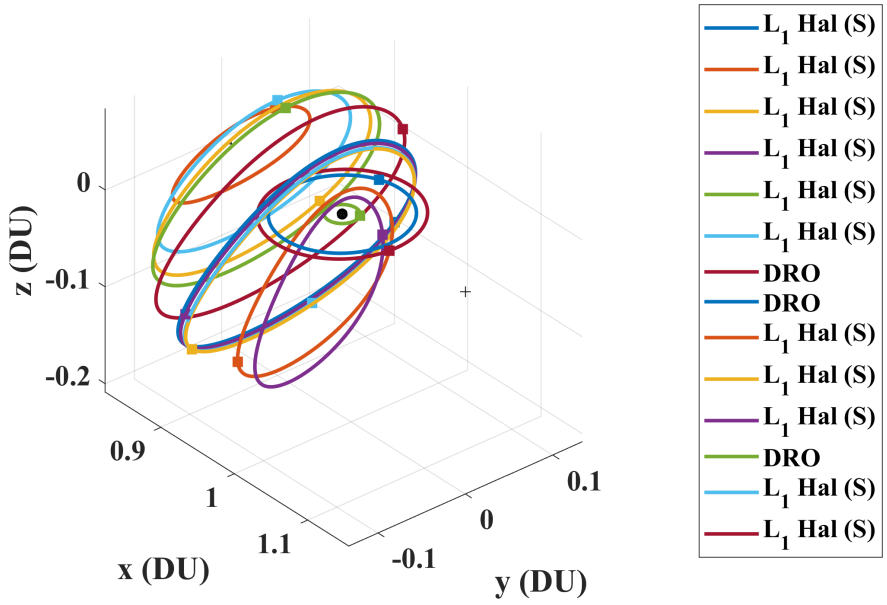

(b) Categorical form

Fig. 11: 14-satellite architectures at $M_{v,Det}$ of 16

4.2.3 Discussion

The Non-Categorical form generally underperforms compared to the Categorical form across both visual magnitude detection thresholds. This is largely due to the inherent restriction of having only three satellites per family in the Non-Categorical form. Consequently, the Non-Categorical form must distribute satellites across multiple families in larger architectures, while the Categorical form, being unrestricted, can find higher-performing architectures using more satellites in individual families. In theory, this limitation should only become noticeable in architectures with four or more satellites since smaller architectures can place all three satellites in one family. However, even in the examined 2-satellite architecture in Figure 11, the Non-Categorical form underperformed compared to the Categorical

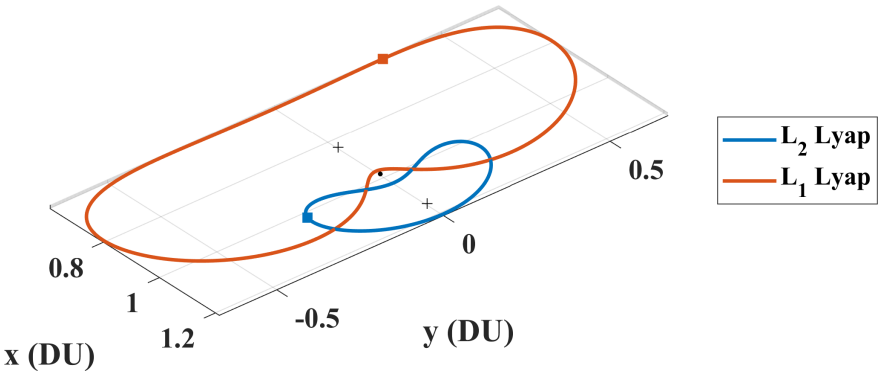

(a) Non-Categorical form

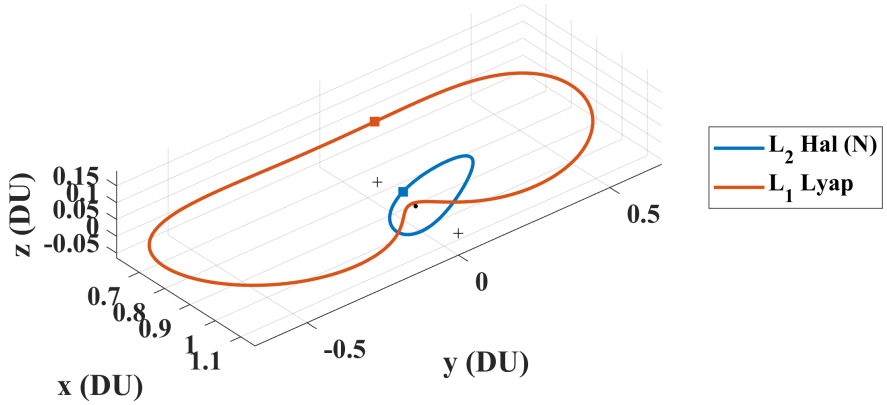

(b) Categorical form

Fig. 12: 2-satellite architectures at $M_{v,Det}$ of 20

form. This is likely due to the limited number of generations. The Categorical form can explore the design space quickly due to its lack of restrictions, allowing it to focus on exploitation in high-performing regions faster. When the number of generations is limited, this rapid exploration and exploitation matter.

The sporadic overperformance of the Non-Categorical form indicates that it can find good architectures, but only in situations where it is not inherently disadvantaged by its form, i.e., being restricted to only three satellites per family. Thus, while the Non-Categorical form is still a viable heuristic, the Categorical form does not have the practical drawbacks that prevent it from finding good solutions consistently.

The Categorical form produces significantly better Pareto fronts and individual architecture solutions than the Non-Categorical form with negligible differences in run time. The Categorical form is a better-performing form of the hidden-genes NSGA-II compared to the Non-Categorical form, with no trade-off in performance or other metrics. Therefore, the Categorical form is the clear choice for developing SSA architectures.

5. CONCLUSION

The core contribution of this research is the development of the Categorical form of the heuristic and its evaluation against the previous Non-Categorical form. The Categorical form not only produced better results than the Non-Categorical form at indistinguishable run times, making it the definitive winner in performance, but it also offers several inherent benefits due to its new chromosome structure. These benefits include:

- Removal of restrictions on orbit families and the orbits contained within, such as requiring unique periods, with the addition of the FOI. This allows for significantly greater flexibility in the placement of observer satellites in various orbits.
- Simple incorporation of new design variables.
- Shortening the length of chromosomes, especially in larger design spaces with more families, design variables, or allowed satellites in the architecture.
- Allowing satellites to be placed in any family or all in one, without defining a maximum number per family.
- Easy adaptation to a non-hidden genes form if the desired number of satellites in the architecture is fixed or known.

By demonstrating the Categorical form as a complete enhancement over the Non-Categorical form, and definitively outperforming it without any of its downsides, a scalable heuristic has been created, allowing for a full-scale analysis of cislunar SSA architecture development. Additional design variables and large observer orbit decks can now be utilized in extensive and diverse scenarios to demonstrate the scalability and extensiveness of the Categorical form of the heuristic. Future work includes demonstrating the Categorical form of the heuristic in an expanded design space to demonstrate it scalability, flexibility, and to apply it to a comprehensive cislunar SSA architecture development scenarios

REFERENCES

[1] Jasbir Arora. *Introduction to Optimum Design*. Elsevier, 2004.

[2] Gregory P Badura, Matthew Gilmartin, Yuri Shimane, Stef Crum, Lois Visonneau, Christopher R Valenta, Michael Steffens, Selcuk Cimtalay, Francis Humphrey, Mariel Borowitz, et al. Optimizing Distributed Space-Based Networks for Cislunar Space Domain Awareness in the Context of Operational Cost Metrics. In *Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 70, 2023.

[3] Kenza K Boudad. Disposal Dynamics from the Vicinity of Near Rectilinear Halo Orbits in the Earth-Moon-Sun System. Master's thesis, Purdue University, 2018.

[4] T Henry Clareson, Matthew C Fox, Dominic Amato, and Hang Woon Lee. Optimization Framework for Space-Based Multi-Sensor Systems in Cislunar Space Domain Awareness.

[5] Ryan D Coder and Marcus J Holzinger. Multi-objective Design of Optical Systems for Space Situational Awareness. *Acta Astronautica*, 128:669–684, 2016.

[6] Albert Corominas. On Deciding When to Stop Metaheuristics: Properties, Rules, and Termination Conditions. *Operations Research Perspectives*, 10:100283, 2023.

[7] Shadi A Darani. *System Architecture Optimization Using Hidden Genes Genetic Algorithms with Applications in Space Trajectory Optimization*. PhD thesis, Michigan Technological University, 2018.

[8] Diane C Davis, Sean M Phillips, Kathleen C Howell, Srianish Vutukuri, and Brian P McCarthy. Stationkeeping and Transfer Trajectory Design for Spacecraft in Cislunar Space. In *AAS/AIAA Astrodynamics Specialist Conference*, volume 8. Springer Nature London, 2017.

[9] Kalyanmoy Deb. *Multi-Objective Optimization Using Evolutionary Algorithms*. John Wiley and Sona, LTD, West Sussex, England, first edition, 2001.

[10] H.A. Eiselt and C.L. Sandblom. *Integer Programming and Network Models*. Springer-Verlag Berlin Heidelberg, Heidelberg, Germany, first edition, 2000.

[11] Naomi Owens Fahrner, Jeremy Correa, and Joshua Wysack. Capacity-based Cislunar Space Domain Awareness Architecture Optimization. Advanced Maui Optical and Space Surveillance Technologies Conference, 2022.

[12] Samuel Fedeler, Marcus Holzinger, and William Whitacre. Sensor Tasking in the Cislunar Regime Using Monte Carlo Tree Search. *Advances in Space Research*, 70(3):792–811, 2022.

[13] Ahmed Gad and Ossama Abdelkhalik. Hidden Genes Genetic Algorithm for Multi-Gravity-Assist Trajectories Optimization. *Journal of Spacecraft and Rockets*, 48(4):629–641, 2011.

[14] Nyoman Gunantara. A Review of Multi-objective Optimization: Methods and Its Applications. *Cogent Engineering*, 2018.

[15] M J Holzinger, C C Chow, and P Garretson. A primer on cislunar space. 2021.

[16] Michael Klonowski, Marcus J Holzinger, and Naomi Owens Fahrner. Optimal Cislunar Architecture Design Using Monte Carlo Tree Search Methods. Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS), 2022.
[17] William E Krag. Visible Magnitude of Typical Satellites in Synchronous Orbits. Technical report, Massachusetts inst of tech lexington Lincoln Lab, 1974.
[18] J. Olson, S. Butow, E. Felt, T. Cooley, J. Mozer, and Peter Garretson. State of the Space Industrial Base 2021: Infrastructure and Services for Economic Growth and National Security. *State of the Space Industrial Base 2021 Workshop*, Nov 2021.
[19] Peng Mun Siew and Richard Linares. Optimal Tasking of Ground-Based Sensors for Space Situational Awareness Using Deep Reinforcement Learning. *Sensors*, 22(20):7847, 2022.
[20] Clint Spesard and Robert A Bettinger. Preliminary Lunar Surface SSA Architecture Optimization for the Observability of Cislunar and Lunar Resident Space Objects. In *AIAA SCITECH 2024 Forum*, page 1674, 2024.
[21] Victor Szebehely. *Theory of Orbits*. Academic Press INC, New York, NY, 1967.
[22] The MathWorks Inc. Gamultiobj Algorithm Documentation, 2024.
[23] The MathWorks Inc. Genetic Algorithm Options Documentation, 2024.
[24] Jacob K Vendl and Marcus J Holzinger. Cislunar Periodic Orbit Analysis for Persistent Space Object Detection Capability. *Journal of Spacecraft and Rockets*, 58(4):1174–1185, 2021.
[25] Lois Visonneau, Yuri Shimane, and Koki Ho. Optimizing Multi-Spacecraft Cislunar Space Domain Awareness Systems via Hidden-Genes Genetic Algorithm. *Preprint arXiv:2302.09732*, 2023.
[26] R Wright, L Tafur, N Owens Fahrner, and J Wysack. Monitoring and Tracking Accessible Invariant Manifolds in the Cislunar Regime. In *Proceedings of the Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, page 204, 2023.