

Faster Exo-Earth yield for HabEx and LUVOIR via extreme precision radial velocity prior knowledge

Rhonda Morgan^{a,*}, Dmitry Savransky^{b,c}, Michael Turmon^{b,a},
Bertrand Mennesson^a, Walker Dula^a, Dean Keithly^{b,c},
Eric E. Mamajek^{b,a}, Patrick Newman^{b,d}, Peter Plavchan^{b,d},
Tyler D. Robinson^{b,e}, Gael Roudier^{b,a}, and Chris Stark^f

^aJet Propulsion Laboratory, California Institute of Technology, Pasadena, California, United States

^bCornell University, Sibley School of Mechanical and Aerospace Engineering, Ithaca, New York, United States

^cCornell University, Carl Sagan Institute, Ithaca, New York, United States

^dGeorge Mason University, Fairfax, Virginia, United States

^eNorthern Arizona University, Department of Astronomy and Planetary Science, Flagstaff, Arizona, United States

^fNASA Goddard Space Flight Center, Greenbelt, Maryland, United States

Abstract. The HabEx and LUVOIR mission concepts reported science yields for mission scenarios in which the instruments must search for potentially habitable planets, determine their orbits, and, if worthwhile, invest the integration time for a spectral characterization. We evaluate the impact of prior knowledge of planet existence and orbital parameters on yield for four mission concept architectures: HabEx 4m telescope with hybrid starshade and coronagraph, HabEx 4m telescope with starshade only, HabEx 4m telescope with coronagraph only, and LUVOIR B 8m telescope with coronagraph only. We use perfect prior knowledge to establish an upper bound on yield and use partial prior knowledge from a potential future extreme precision radial velocity (EPRV) instrument with 3 cm/s sensitivity. We detail a modeling framework that performs dynamically responsive observation scheduling with realistic mission constraints. We evaluate exo-Earth yields against three metrics of spectral characterization for the four mission architectures and three levels of prior knowledge (none, partial, and perfect). The EPRV provided prior knowledge increases yields by ~30% and accelerates by a factor of 3 to 6 the time to achieve half of the yield of the mission. Prior knowledge makes all the mission architectures more nimble and powerful, and most especially starshade-based architectures. With prior knowledge, a small telescope with a starshade can achieve comparable yield to a larger telescope with a coronagraph. © 2021 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.JATIS.7.2.021220](https://doi.org/10.1117/1.JATIS.7.2.021220)]

Keywords: exoplanets; starshade; coronagraph; observation scheduling; mission simulation; HabEx; LUVOIR; extreme precision radial velocity.

Paper 20134SS received Sep. 7, 2020; accepted for publication Jun. 2, 2021; published online Jun. 21, 2021.

1 Introduction

As the number of known exoplanets grows, so too grows the desire to directly image habitable zone exoplanets and acquire spectra of their atmospheres. Particularly compelling is the goal to directly image and spectrally characterize an exo-Earth candidate to search for biomarkers. As complex and expensive mission concepts are designed and put forth, this raises the question of the potential of precursor observations. Can precursor observations support these potential future space missions with early indirect detections of exo-Earth candidates, and, if so, how is yield

*Address all correspondence to Rhonda Morgan, rhonda.m.morgan@jpl.nasa.gov

improved? It is critical to identify worthwhile impact sufficiently early to allow time to develop the next generation of ground-based instruments and perform observations over multiple years. Multiple mission concepts were presented to the Astro 2020 Decadal Survey that could be capable of spectrally imaging an Earth- or super-Earth-sized exoplanet. Here, we study the UV to near-infrared mission designs presented in the HabEx and LUVOIR Concept Study Final Reports: HabEx 4m telescope with hybrid starshade and coronagraph, HabEx 4m telescope with starshade only, HabEx 4m telescope with coronagraph only, and LUVOIR B 8m telescope with coronagraph only. This study is not intended as a comprehensive treatise on the design space. These point designs were previously studied in a common comparison^{1,2} that used two different modeling approaches (discussed below), but the same astrophysical inputs and instrument parameters.

The science yield presented in the HabEx and LUVOIR final reports used scenarios in which no prior knowledge of the exoplanets is available. Thus, the mission instrument must perform a blind search to discover the planets, a series of photometric detections to ascertain orbital parameters with sufficient accuracy to determine if the exoplanet is within the habitable zone (HZ), and, if so, perform the lengthy integration required for spectral characterization.

In this study, we extend our previous work to consider the impact of prior knowledge of the existence of the exo-Earth and its ephemeris. We examine two cases of prior knowledge. The first case is an idealized scenario in which the exo-Earth ephemeris and brightness are perfectly known: no photometric detections are required, and the instrument only performs spectral characterization. This is called the omniscient, or perfect, prior case. It defines the yield upper bound that would be obtained under limiting mission constraints, such as mission lifetime, maximum allowed integration time, solar keepout, or starshade fuel, and determines if the available targets can be exhausted in the available observing time. An ample solution space lies between the upper bound of the perfect prior case and the lower bound of the no-prior, blind-search case. As an intermediate between these lower and upper bounds, we examine the case of partial precursor knowledge provided by a potential future EPRV instrument with sensitivity to 3 cm/s, the sensitivity required to detect exo-Earths to a signal-to-noise ratio (SNR) of 10.³⁻⁵ This case has knowledge of the existence and orbital parameters of some exo-Earths, but the sensitivity of EPRV to exo-Earths is dependent on the planet mass and the host star's temperature. This case allows some spectral characterization to begin immediately at the start of the mission and performs a blind search to identify the non-EPRV-positive targets. We examine how great an impact the EPRV-provided partial prior knowledge has on the yield for the LUVOIR B, HabEx 4H hybrid, 4S starshade-only, and 4C coronagraph-only architectures. To evaluate yield as a function of time, we use the Exoplanet Open-Source Imaging Mission Simulator (EXOSIMS),⁶⁻¹⁶ due to its strength in dynamically scheduling starshade observations. We present this EPRV-related yield information in the context of existing results from a study² done by the Exoplanet Exploration Program Office Standard Definitions and Evaluation Team (ExSDET), which computed yields using both EXOSIMS and the Altruistic Yield Optimization (AYO) yield estimation codes.¹⁷⁻²⁰

The paper is organized as follows. In the remainder of the introduction, we discuss the general modeling approach of these two codes, a brief historical context of their development, and the importance of observation scheduling for starshade missions. Section 2 provides more detail on the model components of EXOSIMS, scheduling algorithms, and observing scenarios, and closes with a summary of the astrophysical and instrument parameters used in this study. Section 3 opens with a discussion of yield evaluation metrics and gives yield summaries for mission scenarios, including yield profiles versus elapsed mission time. The last section offers conclusions and discusses the significance of our results.

1.1 Approaches to Exoplanet Imaging Mission Yield Modeling

Broadly speaking, there are presently two basic approaches to yield modeling, using many of the same constituent parts and similar calculations, but arriving at the final yield estimates in fundamentally different ways. The first approach is to evaluate the summed completeness of a set of target stars for a particular planet population and instrument.^{21,22} Completeness is defined as the probability of detecting an exoplanet belonging to some assumed population,

given that one exists about a particular target star, upon a single observation of that star with a particular instrument to a particular Δ magnitude limit. The summed completeness of multiple stars, scaled by the expected overall occurrence rate of planets in the population, produces an expected value for the number of planet detections of a whole observing program. The concept has been extended to account for subsequent observations of the same target stars.²³

The second approach is Monte Carlo Mission Scheduling (MCMS). In this case, we draw a sample from the assumed planet population to create a “simulated universe,” in which synthetic planets are placed in orbits around the target stars in accordance with assumed occurrence rates and planet properties. We then simulate an observing sequence on this simulated universe, using dynamic scheduling rules and an instrument description, and record the numbers of successful planet detections and characterizations. An ensemble of such simulations (keeping the instrument description and mission rules constant while varying the random draw of synthetic planets) provides both the expectation values and full distributions of planet detections and spectral characterizations, along with any other metrics of interest that are tracked by the simulation, such as fuel use.

Full mission simulations make it significantly easier to directly model dynamically evolving factors and constraints throughout the course of the mission, such as the motion of the observatory on its orbit, the observatory keepout regions, and local zodiacal light contributions. They also allow for the simulation of reactive mission rules, such as requiring spectral characterization follow-ups for positive detections. Such elements can be treated broadly in the completeness yield case, but as that approach does not directly model specific observations, nor the execution of a time-dependent observing plan, it is more difficult to capture such details. The full mission simulations may produce lower yields than completeness, not due to any particular issue with the instrument or observatory design, but solely due to enforcement of operating constraints. As such, completeness results (using identical inputs) can be viewed as an upper bound on MCMS yields, describing a situation where a mission is executed in a completely optimal fashion, with little or no scheduling constraints. MCMS results for the same mission description could be more representative of “true” mission performance in cases where scheduling constraints make it impossible to use all available observing time efficiently within the assumed mission duration. An example of this is the case where a large portion of the target list is clustered in right ascension, meaning that those targets can only be observed at a specific time of year, and the mission does not have enough years to observe them all. Given the complex nature of the scheduling problem, it is inherently difficult to determine whether an MCMS represents a true maximum of performance under the stated constraints or whether better performance could be achieved with smarter time allocation. When the two approaches broadly agree, then we should have high confidence in the results, remembering that the results are dependent on our input assumptions.

The completeness yield and full mission simulation approaches agreed to within 50% in the ExSDET final report of the several architectures with common results. We present these previous results of both modeling approaches alongside the new results of this study in Sec. 3.2 to provide continuity and context.

The completeness-based science yield estimation code AYO^{17,18} is written and maintained by Chris Stark and was initially developed for coronagraph-based direct imaging missions and then extended to starshade-based missions.^{19,20,24} For both coronagraph and starshade missions, AYO assumed that optimum scheduling of the observations was possible, including revisits, with the revisit cadence optimized for completeness, not necessarily for orbit determination. While the assumption of target availability at optimal epochs is well-justified for coronagraphs given their large instantaneous field of regard, the scheduling of starshade observations is limited by stronger time-based constraints, such as fuel consumption [in itself, a generalization of the traveling salesman problem (TSP)] and the solar system keepout.

The mission simulation approach with observation scheduling was first attempted by Hunyadi for the Terrestrial Planet Finder concepts with both a coronagraph and external occulter.^{25–27} The observation schedule abided by the solar keepout constraints, prioritized target selection by completeness improvement divided by integration time, and the yield consisted of the cumulative completeness over the scheduled observations. This work was extended to a

scheduler that generated a single starshade path per scenario as an existence proof of a design reference mission (DRM) for several scenarios for the Exo-S probe study.^{28,29} Savransky introduced optimized scheduling for starshades⁶ and calculated optimal, Monte Carlo DRMs for multiple mission concepts, including coronagraphs.^{7–12} This full mission simulation approach was adapted for the Roman Space Telescope (previously known as WFIRST) and became the EXOSIMS framework,^{13–15} an open-source MCMS framework developed by the Space Imaging and Optical Systems Laboratory at Cornell University with contributions from JPL and other sources.¹⁶ Validation of the EXOSIMS code was performed in prior work.³⁰ EXOSIMS continues to improve the fidelity of starshade propulsion modeling,^{31–33} integration time optimization,³⁴ and schedulers for various architectures and observing scenarios.^{1,35} EXOSIMS is described in more detail in Sec. 2.1.

1.2 Scheduling Constraints for Starshade Missions

While completeness determines the expected value of a target, the design of a starshade mission requires evidence that the starshade has a flight path that can traverse the desired targets within the available fuel mass and schedule, while complying with solar keepout constraints. Minimally, this is an existence proof that a starshade flight path exists that obtains the stated yield. At best, the realized schedule is optimized to balance each target's value against a cost function of fuel use, transit duration, integration time, and potential missed targets or revisits. A simulation tool incorporating robust schedule optimization allows for exploration of mission design space and tuning against various priorities. A realized observation schedule, using the realistic solar constraints and slew transits between targets, can show that the DRM is achievable and verifies that the design of the starshade (dry mass, fuel mass, and propulsion) closes.

Figure 1 shows the importance of dynamic scheduling to account for the key observational constraints of solar system keepout and exosystem planet phase. The solar system keepout refers to the presence of bright foreground objects (including the Sun, Earth, Moon, and other bright solar system bodies) within the field of view or, in the case of a system with an external starshade, geometries where light from these bright sources would be reflected or scattered into the telescope aperture. Exosystem planet phase refers to the fact that the orbits of exoplanets (even ones known to exist via previous indirect detections) are not constrained until after multiple imaging observations have been performed.

Figure 1 shows the impact of these observational constraints for a starshade observing case typical for HabEx. The stars used are nearby stars from the ExoCat-1 catalog³⁶ having a range of luminosities; their properties are shown in Table 1 (for more on the catalog, see Sec. 2.5.1). The exo-Earths simulated in the figure are Earth twins at the Earth-equivalent insolation distance (EEID) with a 60-deg inclination, three zodis of exozodiacal brightness, zero eccentricity, and random initial phase. For each day in the simulation, the integration time required to achieve a characterization SNR of 10 was calculated based on the planet phase and separation angle at that day; if the integration time was less than the mission policy of 60 days, then the SNR constraint for spectroscopic characterization was considered satisfied for that day. Most of the integration times to achieve this SNR were a few days and increased exponentially for the last 3 days within the 60 day cutoff. (The intersection of the SNR constraint and solar constraint does not consider if the entire observing window satisfies the solar constraint, only the initial day, so the green bars in the histogram should be considered an upper bound.) The 15 exo-Earths shown are the candidates selected for characterization by the scheduler during a 5-year nominal mission for a particular synthetic universe. In the illustrated cases, the solar system keepout removes over 50% of characterization observing windows, whereas the exosystem planet phase removes as much as 75%. The two constraints operate independently, so the compound observable window is ~30% and varies by year. For several illustrated targets, some years have virtually no observable window. This situation also underscores the importance of orbit determination to allow scheduling of characterization time within high-SNR observing windows and the importance of discovering exo-Earth candidates early in the mission to ensure multiple scheduling opportunities.

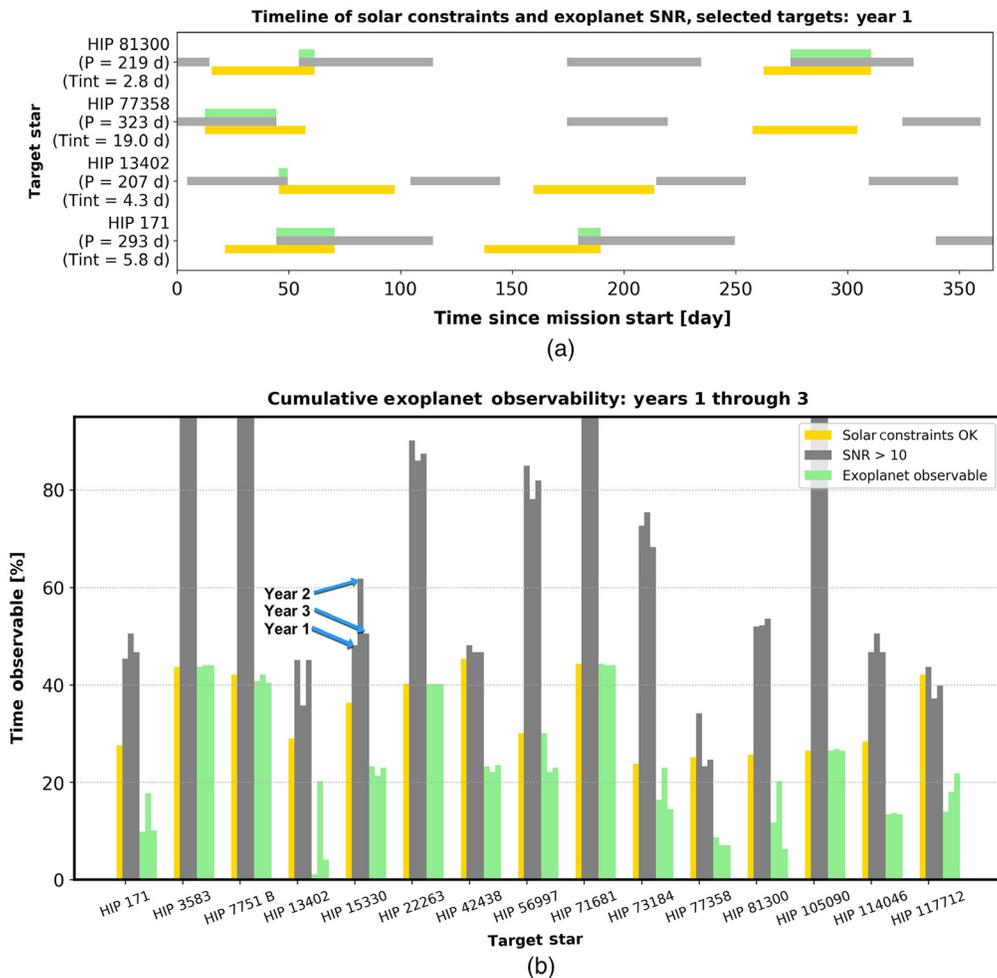


Fig. 1 The independent observing constraints due to solar system keepout (yellow) and exosystem planet phase (gray) and their intersection (green) for a HabEx starshade scenario. (a) Constraints across the first mission year for four targets. Times where solar system observing constraints are satisfied are in yellow, times where an SNR > 10 is attainable are in gray, and the intersection is green. The SNR constraint is satisfied largely depending on the exoplanet's orbital phase (its period, and the median integration time to achieve SNR = 10, are given in days at left). However, the solar constraints have yearly periodicity, resulting in a beating phenomenon for the intersection in green. (b) shows how this expresses itself over 3 years of observations, plotting the cumulative constraint satisfaction (solar, SNR, and combined) for the 15 stars that supported characterization in this simulated universe. Solar constraints, constant for each year, are in yellow, whereas the SNR and intersection constraints vary for each year. For some targets (e.g., HIP 13402, years 1 and 3), the overlap of these essentially unlinked constraints can be nearly zero over an entire year of observations.

2 Methods

We give an overview of the EXOSIMS components, emphasizing its strategy for observation scheduling. We lay out several mission observing scenarios and then describe the implementation of the prior knowledge cases.

2.1 Overview of EXOSIMS

EXOSIMS¹⁶ is an MCMS simulation code that simulates Monte Carlo universes. For a single universe simulation, synthetic planets are drawn from planet property distributions, then an end-to-end mission is simulated. The mission is dynamically, sequentially scheduled and executed, where the scheduled next targets respond to the results of previous observations (successful, null,

Table 1 Properties of target stars used in Fig. 1.

HIP	Distance (pc)	V	$L_{\text{bol}} (L_{\odot})$	EEID (AU)	EEID (mas)	Δmag
171	12.17	5.80	0.70	0.84	68.65	24.45
3583	15.16	5.80	0.98	0.99	65.44	24.83
7751	7.82	5.68	0.34	0.58	74.55	23.68
13,402	10.35	6.05	0.42	0.65	62.36	23.89
15,330	12.01	5.53	0.80	0.89	74.45	24.60
22,263	13.28	5.49	1.01	1.01	75.75	24.85
42,438	14.36	5.63	1.03	1.01	70.57	24.87
56,997	9.61	5.31	0.65	0.81	83.79	24.37
71,681	1.34	1.35	0.52	0.72	541.30	24.14
73,184	5.86	5.72	0.29	0.54	91.81	23.50
77,358	15.35	6.01	0.87	0.93	60.79	24.69
81,300	9.75	5.76	0.47	0.69	70.60	24.03
105,090	3.95	6.69	0.08	0.28	71.05	22.08
114,046	3.28	7.35	0.04	0.20	62.13	21.39
117,712	10.89	6.38	0.39	0.62	57.29	23.82

or false positive). EXOSIMS consists of a collection of modules, abstracting out various state variables and methods associated with different aspects of the full mission simulation, all with a strictly defined input/output specification. This allows for modules to be extended, or for new modules to be implemented, without requiring any modifications to other portions of the code. EXOSIMS is thus adaptable to entirely new designs for instruments, observatories, or overall mission concepts. Figure 2 shows a schematic representation of how an EXOSIMS Mission Simulation object is constructed out of the individual modules. The module functionality split is described in Appendix A.

Yield is calculated from a generated survey ensemble, which contains the outputs of N survey simulations, each for a given simulated universe, all other mission and instrument parameters being constant. The simulated universe is regenerated for each individual mission simulation by resampling the planet population distributions and synthesizing entirely new planets. The execution of a single mission simulation occurs within a top-level loop, which selects each subsequent target from the pool of “currently” (at the current simulated mission time) available targets, calculates the required integration time, simulates the observation (and followup characterization, if called for by the selected mission rules), and then simulates the outcome, which can be a true positive (detection), false positive (misidentified speckle or background object), true negative (null detection), or false negative (missed detection). The functionality of the default survey execution is extended and augmented in multiple additional survey simulation implementations designed specifically for analyzing HabEx and LUVOIR. Below we detail two of these, which were used extensively to produce the results in this report.

2.2 Observation Scheduling

The scheduler module is the backbone of EXOSIMS. It drives the dynamic observing sequence and contains the logic for autonomous target selection. Multiple schedulers exist for EXOSIMS, each tailored to an architecture and an observing scenario. The weighted linear cost function (WLCF) scheduler is used for architectures having a starshade (HabEx 4S and HabEx 4H).

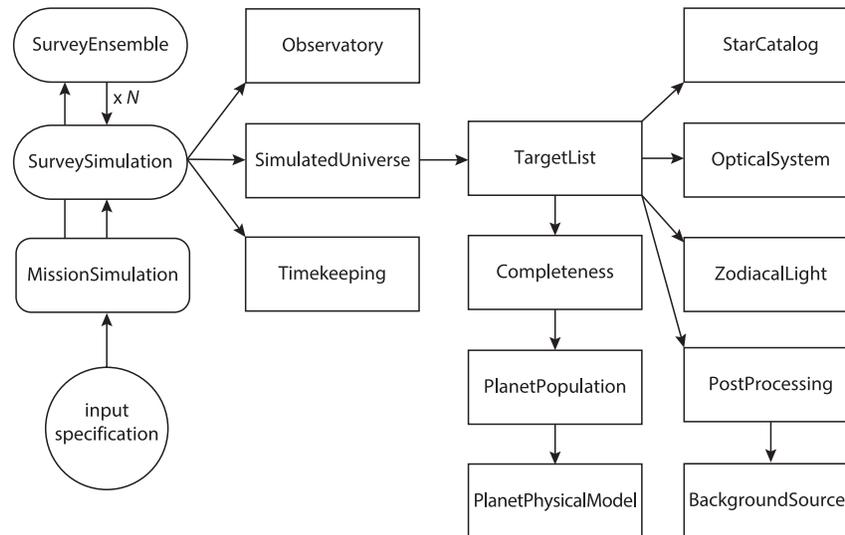


Fig. 2 Schematic depiction of the instantiation path of all EXOSIMS modules. The arrows represent instantiation of an object, and object references to each module are always passed up directly to the top calling module so that a given module has access to any other module connected to it by a direct path of instantiations. For example, TargetList has access to both PostProcessing and BackgroundSources, whereas Observatory does not have access to any other modules. The typical entry point to EXOSIMS is the construction of a MissionSimulation object, which causes instantiation of SurveySimulation, which in turn instantiates all the other modules.

In the case of the HabEx 4H architecture, which uses both a starshade and a coronagraph, a tiered scheduler is further used. For coronagraph-only architectures, a simplified WLCF scheduler is used. Below we introduce the core scheduling approach used, how it is adapted to serve the coronagraph/starshade hybrid case, and how it accommodates starshade-only repeat observations for orbit determination.

2.2.1 Weighted linear cost function scheduler

The WLCF scheduler, implemented in EXOSIMS as “linearJScheduler,” is based on an observation scheduling methodology fully described by Savransky et al.¹¹ A starshade scheduler must find a compromise between high-completeness targets and low-slew targets. The problem of maximizing completeness of a series of observations subject to an overall slew constraint is related to the well-known TSP, but with key differences that make a time-windowed but dynamic approximation more effective.¹¹ A detailed description of the approach and cost function terms are in [Appendix B](#).

Changing the weighting factors on the cost function terms changes mission priorities. Optimal weighting factors can be determined by wrapping the whole simulation in an optimization scheme with an appropriately defined objective function (e.g., the number of Earth-twin spectral characterizations within a fixed set of simulated universes). This approach was used several times to evaluate parameter choices for this scheduler under different operating conditions.

2.2.2 Tiered scheduler

The tiered scheduler was developed for the HabEx hybrid starshade and coronagraph architecture. It is a hierarchical scheduler that schedules the starshade at the top tier and fills in coronagraph observations at the lower tier. Observations are scheduled sequentially to allow for dynamic response to the discoveries and the success or failure of observations. The starshade observations have first priority, using the above WLCF scheduler to find an effective starshade path. During the starshade slews, coronagraph observations are scheduled within the second tier.

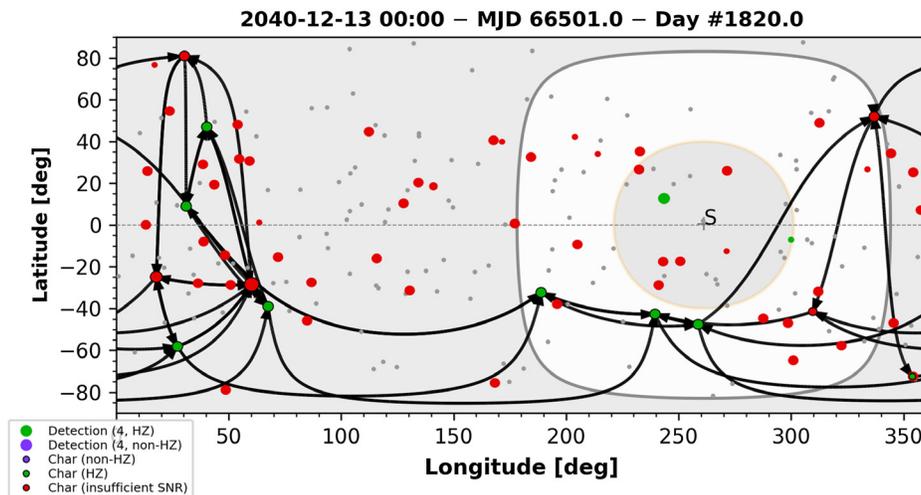


Fig. 3 Map in ecliptic coordinates of a typical hybrid starshade/coronagraph observing sequence. The image shown summarizes the final outcome of the detections and characterizations within an observing sequence shown in full in [Video 1](#). The synthetic planets are observed and considered detected or characterized if the goal SNR is reached. The target is shown in green for rocky planets in the HZ, red for insufficient SNR to detect any planets, or gray for an unobserved star. The size of the circle indicates the number of repeat detections or characterizations, with the case of four detections shown in the legend for scale. Spectral characterizations with the starshade are distinguished by a black edge around the circle marker and are at the tip of a black slew arrow. The white annulus is the region of observability for the starshade, moving across the sky as time advances. The central gray circle with yellow border is the solar keepout zone for both the coronagraph and the starshade. The Sun's position is shown by a gray cross, labeled "S." [Video 1](#), MP4, 1.4 MB [URL: <https://doi.org/10.1117/1.JATIS.7.2.021220.1>].

Also during the slews, general observatory (GO) observation time is allocated at the specified rate (50% for the HabEx baseline mission concept). GO time is allocated whenever "owed" GO time accumulation exceeds 1 day, unless a starshade observation is occurring, in which case the owed time is allocated at the end of the starshade observation. This distributes the GO time evenly throughout the mission. Additional details are in the results section for the HabEx 4H hybrid case (Sec. 2.3.2).

Figure 3, available as a video, shows an observing sequence for one possible universe for the HabEx hybrid starshade and coronagraph scenario, scheduled with the tiered scheduler. (The still-image version shows the final tour executed by the mission.) The region of observability of the starshade given the solar keepout constraints is the white annulus within the figure, and starshade slews are indicated by the black arrows. While the starshade slews to its next target, the telescope rapidly repoints to allow the coronagraph to search for new exoplanets. Revisit observations are indicated by an increased marker size. Once sufficient observations determine the orbit, the target is promoted to the starshade for spectral characterization.

2.2.3 Scheduling for starshade-only orbit determination and spectral characterization

The WLCF scheduler was extended for the starshade-only case to provide orbit determination with multiple detection visits. In the starshade-only observing scenario, the starshade performs a blind search in imaging (detection) mode. Once a planet is detected, it is immediately characterized spectrally. After spectral characterization, the completeness, which is used in the cost function, is set to unity, the highest value allowed for completeness, to increase its priority as a target for follow-up detections for orbit determination. The starshade is required to wait a third of an orbital period of an EEID orbit before it can perform an additional detection attempt. The target is retired from the target pool once the minimum number of detections has succeeded. With this scheduler, a subset of the spectrally characterized planets have a sufficient number of detections for their orbits to be considered determined.

Table 2 Exo-Earths observing scenarios.

Parameter	LUVOIR B	HabEx C	HabEx H	HabEx S
Photometric detections	Coronagraph	Coronagraph	Coronagraph	Starshade
Bandwidth	20%	20%	20%	76%
Spectral characterization	Coronagraph	Coronagraph	Starshade	Starshade
Bandwidth	4 × 20%	4 × 20%	108%	108%
Initial observation	Blind search photometric detection			
Response to detection	Orbit determination via photometric detections			Immediate spectral char.
Final follow-up	Spectral characterization			Orbit determ.
EXOSIMS module	CoroOnly, Sec. 2.3.1	CoroOnly, Sec. 2.3.1	Tiered Scheduler, Sec. 2.3.2	WLCF plus orbit determ.

2.3 Observing Scenarios

Here, we discuss the observing scenarios for the coronagraph-only (LUVOIR B and HabEx 4C), the hybrid starshade and coronagraph (HabEx 4H), and the starshade-only (HabEx 4S) architectures. The scenarios are for a 5-year duration mission with 50% of the mission time allocated to exoplanet observations. The corresponding observing scenarios for the baseline no-prior case are summarized in Table 2. Each observing scenario section discusses the no-prior case and then the perfect-prior case. The partial prior assumptions and implementation for each observing scenario are discussed in Sec. 2.4.

2.3.1 Coronagraph only

The coronagraph-only scenario uses the coronagraph instrument for photometric imaging observations (detections) for the blind search and orbit determination, as well as for spectral characterization. The coronagraph is limited to 20% instantaneous bandwidth, and spectral characterization is accordingly performed sequentially to compose an aggregate spectrum. The maximum allowable integration time of 60 days was applied to the total integration time of the aggregate spectrum. The maximum integration time of 60 days was a mission policy used by both the HabEx and LUVOIR Concept Studies. The field of regard is large; pointing is allowed outside a 40-deg sun angle for both HabEx and LUVOIR. The field of regard for the starshade, for comparison, is limited to an annulus from 40 deg to 83 deg around the sun. The coronagraph benefits from being able to point to the entire anti-sun hemisphere, and the sun constraint does not drive the mission scheduling like it does for the starshade. The importance of observation scheduling for the coronagraph is in the revisit cadence for optimizing the probability of detection and successful orbit determination. EXOSIMS schedules the detection observations at a cadence optimal for orbit determination and promotes the planet only when the minimum number of detections span more than half a period. Lower completeness stars are likely to have more null detections during orbit determination attempts and are likely to be promoted for characterization late in the mission, or the mission may end before sufficient orbit determination epochs have elapsed. The observation scheduling exposes the inefficiency of the coronagraph blind search under realistic mission scheduling constraints.

Blind search. The crafting of the blind search seeks to efficiently discover, revisit, retire, or promote targets. The target-rich environment for the coronagraph means that target pruning should be applied. The first step is to prune the target pool based on potential characterization outcome, not detection outcome. Targets are excluded if the minimum characterization integration time for an Earth twin at quadrature (assuming zero zodiacal brightness and zero exozodiacal brightness) is greater than the mission policy limit of 60 days and if the completeness

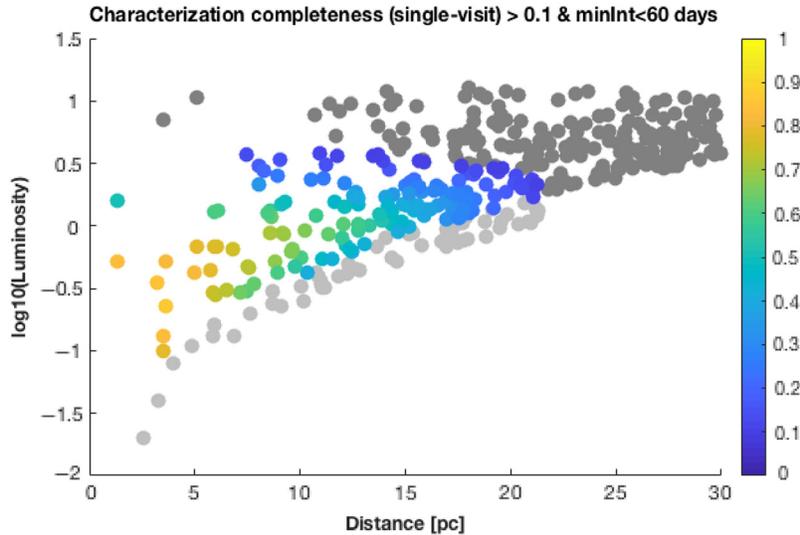


Fig. 4 Target stars for the HabEx 4S architecture are color coded by characterization completeness as function of distance from Sun and stellar bolometric luminosity. Potential targets whose detection integration times are longer than 60 days (dark gray) or that have completeness <10% (light gray) are filtered out.

is below 10%. Figure 4 shows the subset of targets kept after filtering for characterization completeness >0.1 (colored by characterization completeness) and the targets (gray), which have characterization integration time <60 days.

The sequence the targets are observed in is a ranking of completeness/integration time. The integration time used for ranking is the minimum integration time required when assuming the minimum solar zodiacal brightness over the sky and zero exozodiacal brightness. The list is observed in order. An overhead of 5 h per observation is included in the total per target observing time to account for observatory settling, digging the coronagraph dark hole, and focal plane readout. The revisit wait time for each star is a third of the period of an Earth twin at the EEID. As the wait-period is reached for each target, the target is available for another epoch of observing. In this way, the high value targets are observed early and promoted as early as possible; the targets with longest integration times are observed latest in the mission. This ranking, and filtering out targets with completeness <0.1, is the only place where completeness is used. Completeness is not used to calculate yield. Yield is calculated based on successful spectral observations of synthetic planets during the simulated mission.

The coronagraph blind search raises the question of how many null detections to allow before retiring a target. A null detection could occur when the planet is inside the coronagraph inner working angle (IWA), has a poor SNR, or if a planet does not exist. A high number of visits with null detections, while conservatively thorough, uses an integration time that could be spent on other targets. Following reasoning similar to Brown and Soummer,²³ we note that the probability of a series of N independent missed detections, even though an exo-Earth is present, is $(1 - C)^N$, where C is the single-visit completeness. Turning this around via Bayes' theorem shows that the probability that an exo-Earth is present even though it was missed N times is given as

$$P(\text{Earth}|N \text{ misses}) = \frac{P(N \text{ misses}|\text{Earth})P(\text{Earth})}{P(N \text{ misses})} = \frac{(1 - C)^N \eta}{(1 - C)^N \eta + (1 - \eta)}, \quad (1)$$

from which it is easy to find the number of observations needed to force this below any given threshold (Fig. 5). For example, for moderately difficult targets with $C = 0.5$, five failed observations are needed to force the probability that an exo-Earth is present below 0.01. We choose a policy based on <1% chance of missing an exo-Earth for high-completeness targets, that is, with $C > 0.7$. A sensitivity study showed that $N = 2$ was more efficient and no less effective than $N = 3$. Even for a perfect instrument, there is inherently an inefficiency of $N(1 - \eta)$

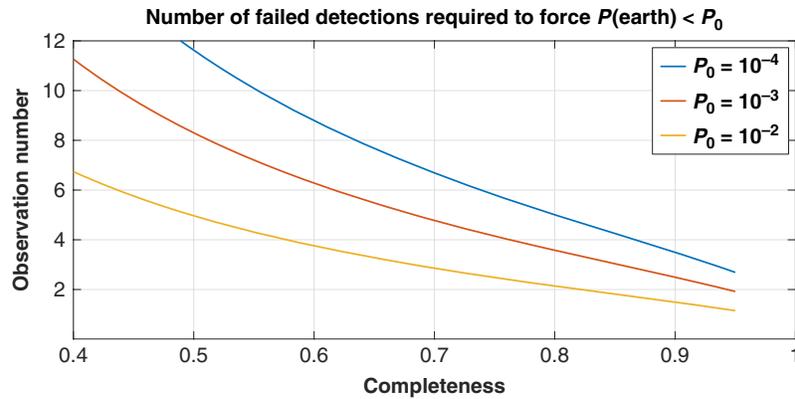


Fig. 5 The number of failed detections required to force the probability of an EEC less than a selected threshold P_0 , as a function of first-visit completeness, taking $\eta = 0.24$.

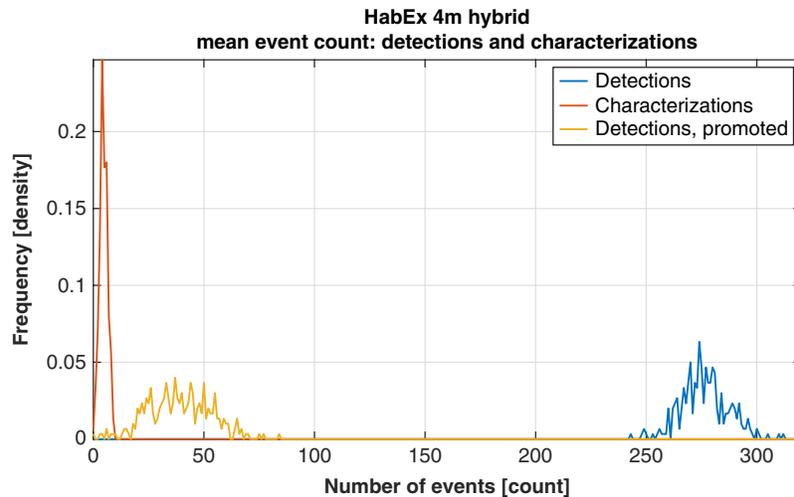


Fig. 6 A histogram over an ensemble of 300 simulations showing the number of characterizations (red), all coronagraph detections (blue, mean is 274), and coronagraph detections contributing to a starshade characterization (yellow, mean is 38.8).

observations. An additional inefficiency arises for planets detected 2 or less times during the orbit determination phase, as these will not be promoted for spectral characterization.

The craftsmanship of the coronagraph blind search involves trading thoroughness for efficiency. A parameter sweep was performed for: the weighting factor of revisits, number of allowed null detections, maximum number of successful detections, and the minimum number of detections required for promotion. The heuristic that produced the highest characterization yield was used for all the coronagraph blind searches. Figure 6 shows the histogram of the total number of detection attempts (blue) and the subset of detections, which contributed to a promoted target (yellow) for the nominal, baseline HabEx case. An average of about 280 detection attempts occurred, of which ~ 50 contributed to targets that were promoted. The detection attempts on promoted targets include nulls and the three successful detections required for orbit determination.

No prior knowledge case. The EXOSIMS coronagraph-only scheduler makes detection observations and revisits until an exo-Earth candidate is discovered and its orbit has been characterized by a minimum of three detections spanning more than half a period, at which point the target is promoted for spectral characterization. Spectral characterization is prioritized over the

coronagraph blind search and happens immediately after the final orbit-determination detection, if sufficient time remains in the window of observability. For the characterization observation, the exoplanet is placed at quadrature, as with the tiered scheduler, to facilitate comparison. Given the large region of observability of the coronagraph, the intersection between solar constraint and favorable planet SNR is likely to occur within a 1-year window, and the justification that the characterization could be scheduled at a favorable orbital phase once the orbit is determined is more robust.

The orbit determination goal is 10% error or less on the estimation of semimajor axis, eccentricity, and inclination. Horning et al.³⁷ and Nielsen²⁴ determined a heuristic minimum of three observations spaced over greater than half the planet's period would provide 10% uncertainty in semi-major axis and eccentricity in the vast majority of cases. (For high-inclination orbits, a fourth observation is required.)

The LUVOIR B nominal observing scenario budgets 6 months total for full spectra characterization and 2 years total for the coronagraph blind search and initial spectral characterization with a single 20% bandwidth spectra for water line detection at 940 nm. The AYO results in the LUVOIR final report use 2 years cumulative integration time to calculate exoplanet yield. The EXOSIMS results presented here use 50% of the 5-year mission time, or 2.5 years cumulative mission time, for the coronagraph blind search and spectral characterization of targets at the water line. The cumulative 2.5 years, instead of 2 years, were used to facilitate comparison with the HabEx architectures and additional metrics, which were evaluated for 2.5 years of cumulative mission time dedicated to direct imaging and spectroscopy of exoplanets.

Perfect prior knowledge case. For the coronagraph-only case, the perfect prior assumes that the existence and orbital parameters are known for all exo-Earths; the blind search is not performed. Instead, the scheduler spectrally characterizes targets in order of characterization completeness divided by total characterization integration time. The schedule limitation in the perfect prior is not windows of observability as in the starshade case, but instead the maximum number of targets that can be spectrally characterized within the integration time policy limit. The perfect prior case shows the upper bound of targets that can be characterized within the limits of the instrument and shows when in the elapsed mission time the targets are exhausted.

2.3.2 Hybrid: coronagraph and starshade

The hybrid case leverages the complementary strengths of the starshade and coronagraph. The strength of the coronagraph is its agility to observe and revisit numerous targets, which is well suited for photometric detections used to discover exo-Earths and revisits for orbit determination. The strength of the starshade is to provide a continuous spectrum over more than 100% bandpass. The starshade offers a shorter total integration time for spectral coverage due to its high throughput and broad bandwidth. For the HabEx hybrid architecture, the starshade was designed to have a similar IWA to the coronagraph (Table 4), so their spatial sensitivity to the habitable zone will be similar. The starshade has a much larger outer working angle (OWA) than a coronagraph, allowing a greater opportunity to simultaneously image multiple planet systems. When paired with an integral field spectrometer (IFS) both a coronagraph and starshade can acquire spectra on multiple planets, though the coronagraph must take a sequence of subspectra while the starshade offers simultaneous full spectra.

No prior knowledge case. The mission begins with photometric detection observations (detections) with the coronagraph. Once a planet is detected, follow up imaging (detection) observations are made until the orbit is determined. If the orbit is within the habitable zone, then the target is promoted to the pool of targets for spectral characterization by the starshade. A maximum integration time of 60 days is allowed for spectral characterizations in accordance with HabEx mission criteria. An overhead of 8 h is included in the total observation time for each starshade repositioning to account for formation flying acquisition and settling. The hybrid scenario uses the EXOSIMS tiered scheduler discussed in Sec. 2.2.2. The coronagraph blind search and orbit determination are the same as described in Sec. 2.3.1. The starshade scheduling is discussed here.

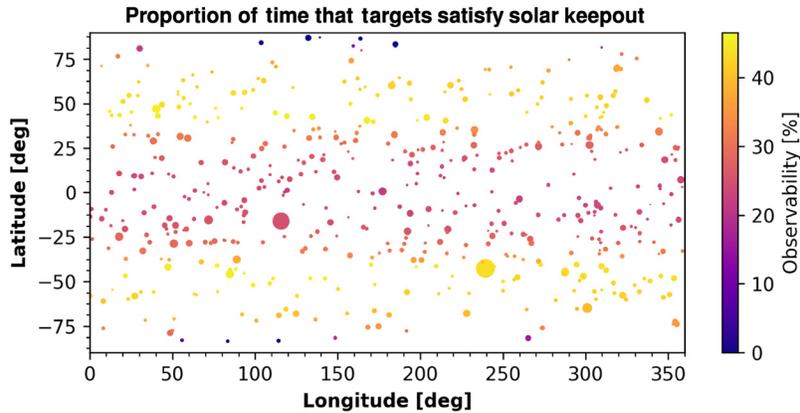


Fig. 7 The starshade sky coverage in ecliptic coordinates, colored by the proportion of time within a 5-year observing interval that the target star satisfied solar keepout constraints. The marker size is proportional to V magnitude.

For starshades, at a fixed time, the solar keepout constraint results in an annular region of observability, shown in white in Fig. 3. This annulus moves smoothly across the ecliptic plane over 1 solar year, and when a given star is within the annulus, it is observable. Thus, the ecliptic latitude of a star determines the temporal windows of observability within a solar year. This can range from two windows of 43 days duration separated by 80 days (for a target on the ecliptic) to one long window of ~ 165 days (for a target near ± 45 deg ecliptic latitude). The temporal windows of observability as a function of sky position are summarized by the sky coverage plots of cumulative observability-fraction (Fig. 7) and longest continuously observable interval (Fig. 8). The first plot shows that low-latitude targets are significantly less-frequently observable than moderate-latitude stars, whereas the second plot shows that some of these small-latitude stars have a bound on integration time considerably less than the nominal 60-day limit. Stars within 7 deg of the ecliptic poles are not observable with the starshade due to its 83-deg anti-sun angle constraint.

Perfect prior knowledge case. To determine an upper limit on the number of planets that can be characterized while abiding by realistic mission constraints (sun angles, slew travel times, and integration time policy limit), we separate the starshade scheduling from the blind search by examining the case of perfect prior knowledge. For the case of perfect prior knowledge, the code is run in an “omniscient mode” in which all of the synthetic exo-Earths are promoted for

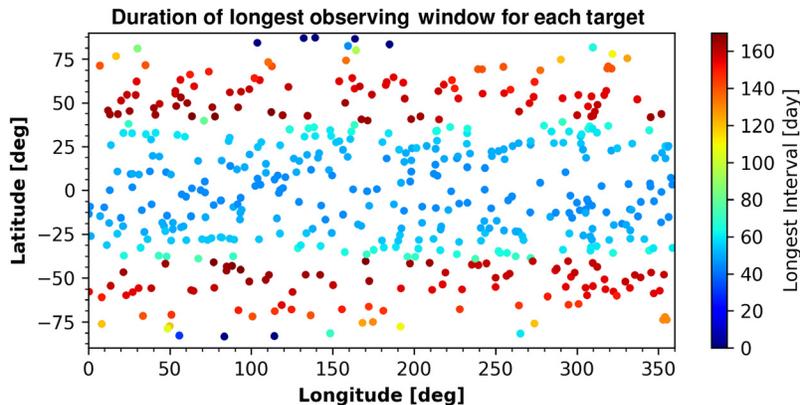


Fig. 8 The starshade sky coverage in ecliptic coordinates, colored by the length of the longest continuous observing interval within a 5-year starshade mission. The starshade solar keepout constraints allow for a maximum observable interval of ~ 160 days.

characterization at the start of the mission; the planet properties and exozodi are used directly to calculate a tailored integration time and only targets that can achieve the target SNR of 10 within 60 days of integration will be characterized. While the ephemerides of the planets are known perfectly in omniscient mode, the implementation of a time-window constrained, TSP scheduler and optimizer was beyond the scope of this study and is a subject for future work. As an alternate, the exo-Earth is placed at quadrature for the characterization observation, with the justification that the characterization could be scheduled at a favorable orbital phase sometime during the remaining mission time. Based on Fig. 1, an intersection of the solar constraints and a favorable planet SNR is likely to occur within a 3-year window. For the 5-year mission duration studied here, which is target-limited, this provides ample schedule margin for the omniscient case. For the no-prior baseline case, targets promoted for characterization in the first 2 years should have sufficient schedule margin for this assumption, but targets promoted later in the mission do risk conflict with observing constraints. In addition, quadrature is not the brightest point for characterizing the planet, so this provides some photometric margin while utilizing the widest separation angle of maximum elongation. Figure 9 shows an omniscient mode starshade observing timeline for a specific DRM realization. For this figure, the mission allocation was allowed to be near 100% to determine if the scenario was target limited. The available targets were depleted half way through the mission, and the spectral characterization took about 20% of the total mission time. Note that the detections (indicated in blue) occur on a separate tier than the starshade scheduling and do not impact the characterization yield for the omniscient case; they were kept as a diagnostic. To facilitate comparison of the three prior-knowledge cases, the exo-Earth was placed at quadrature for characterization for all three prior-knowledge cases, except for the starshade-only observing scenario as discussed in Sec. 2.3.3.

The WLCF weighting coefficients were tuned for each prior knowledge case. For the omniscient case, in which revisits were unnecessary, the cost function variables were effectively reduced to two variables: integration time and slew time. The tuning of the coefficient weights was fairly insensitive because all of the viable targets were observed during the mission. For a shorter mission or more targets, the coefficient weights would have higher sensitivity and impact.

2.3.3 Starshade only

No prior knowledge case. The starshade-only architecture utilizes the starshade for the blind search, the spectral characterization, and the orbit determination, which are executed in that order for a given target star. The observing sequence is different than the coronagraph-only and the hybrid scenario. In the starshade-only observing scenario, the observatory performs a blind search in detection mode, which is imaging over a 450- to 1000-nm bandpass. When an exo-Earth is discovered in imaging mode, a spectral characterization is taken immediately, as long as sufficient time remains in the observable window. Immediate characterization ensures that a spectrum is acquired before the exoplanet can propagate to an unfavorable, and possibly unknown, position at the next epoch of potential starshade observation. After spectral characterization has been conducted, follow-up broad-band imaging detections contribute to orbit determination. For the cases studied, not all targets that are spectrally characterized received all three requisite orbit determination observations, due to the realistic constraints of observing with the starshade. The revisit wait time of a third of a period of an Earth-twin at the EEID, also used in the coronagraph blind search, was enforced, though a more optimal scheme could be developed utilizing orbit information as the orbit fit evolves at each epoch.

It is worth noting that a planet discovered by the starshade in imaging mode at a single epoch may not be Earth-sized nor orbiting in the habitable zone. Astrometric information could indicate if the planet is outside the habitable zone or has an unreasonably high eccentricity, such that a spectrum should not be attempted. This would require the definition of an eccentricity upper limit for planetary orbits that cross the habitable zone. In addition, if a spectrum is started, the integration time could be tailored or terminated when sufficient SNR is achieved to indicate that the further integration time is not desirable. Inclusion of decision-making for this degeneracy in multiplanet systems is part of future work for EXOSIMS. For the modeling in this paper, only exo-Earth populations were used, and this degeneracy was avoided.

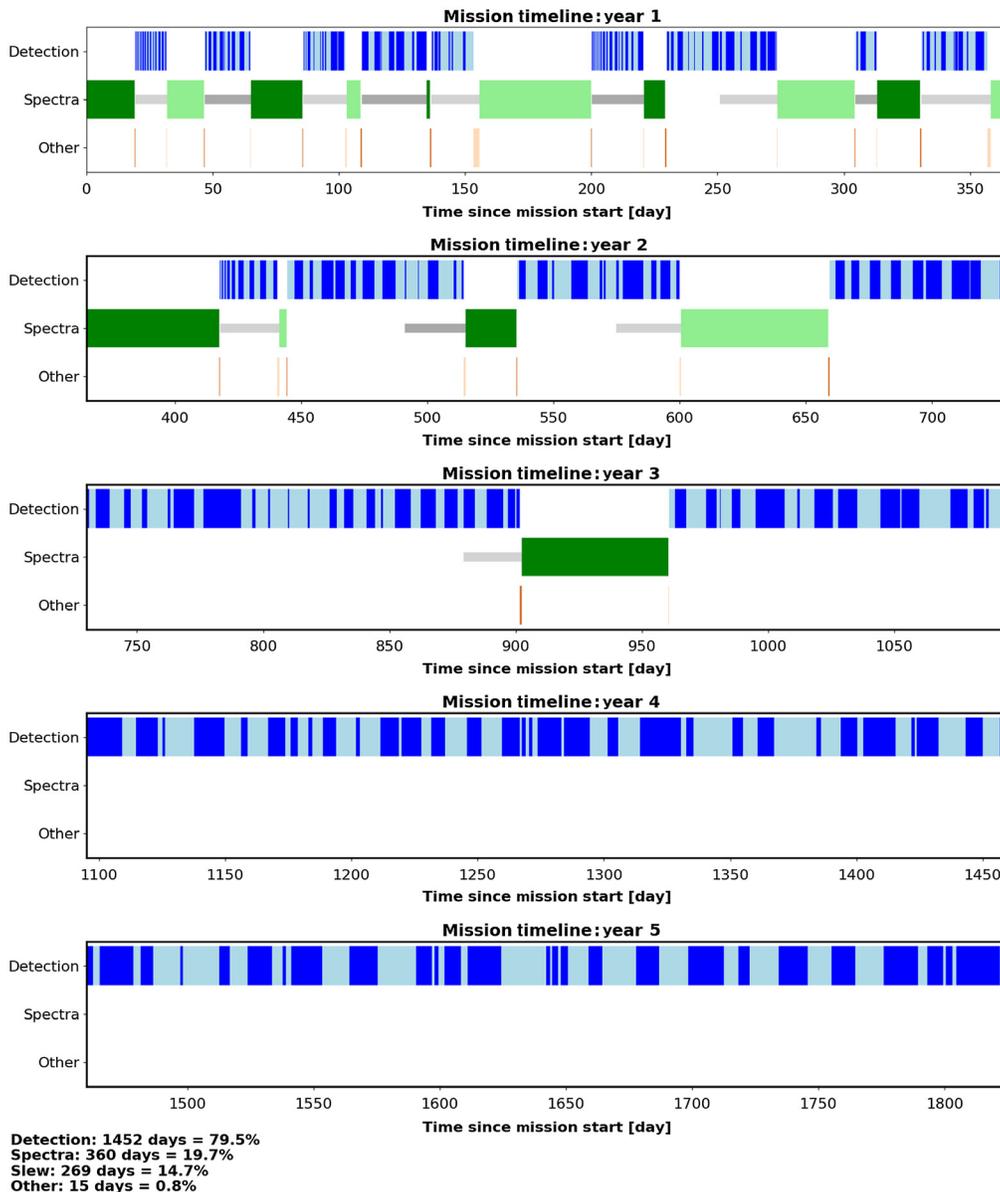


Fig. 9 The observation timeline for an omniscient case for an exemplar synthetic universe in which all characterizable earthlike planets were promoted to the starshade target list at the mission start. Green bars indicate spectral characterization integration times, gray lines indicate starshade slew times. Bars have alternating light and dark colors to distinguish event sequences. Observation scheduling is limited only by keepout constraints and slew times. The characterizable targets were exhausted halfway through the mission. Coronagraph detection observations (blue bars) were allowed to occur during slews to compare the number of targets discovered during the blind search to the number of targets that could be characterized in the omniscient case.

Figure 10 shows the timeline of observations for the starshade-only scenario for an exemplar DRM. The imaging mode detections, shown in blue, include both blind search and follow-up orbit determination observations. The spectral characterizations, indicated in green, occur immediately after a detection (blue). In the starshade-only scenario, the characterization precedes the orbit determination, and so an exoplanet positioned at quadrature is not assumed. The planet phase used for spectral characterization is the phase of the planet at the end of the detection integration time. The mission timeline plot shows that the high-completeness stars are characterized early in the mission, whereas the stars with lower completeness and longer integration times are visited late in the mission. The starshade-only case allows for a generous portion of the mission time for GO time during starshade transit slews: almost 80% in this exemplar DRM.

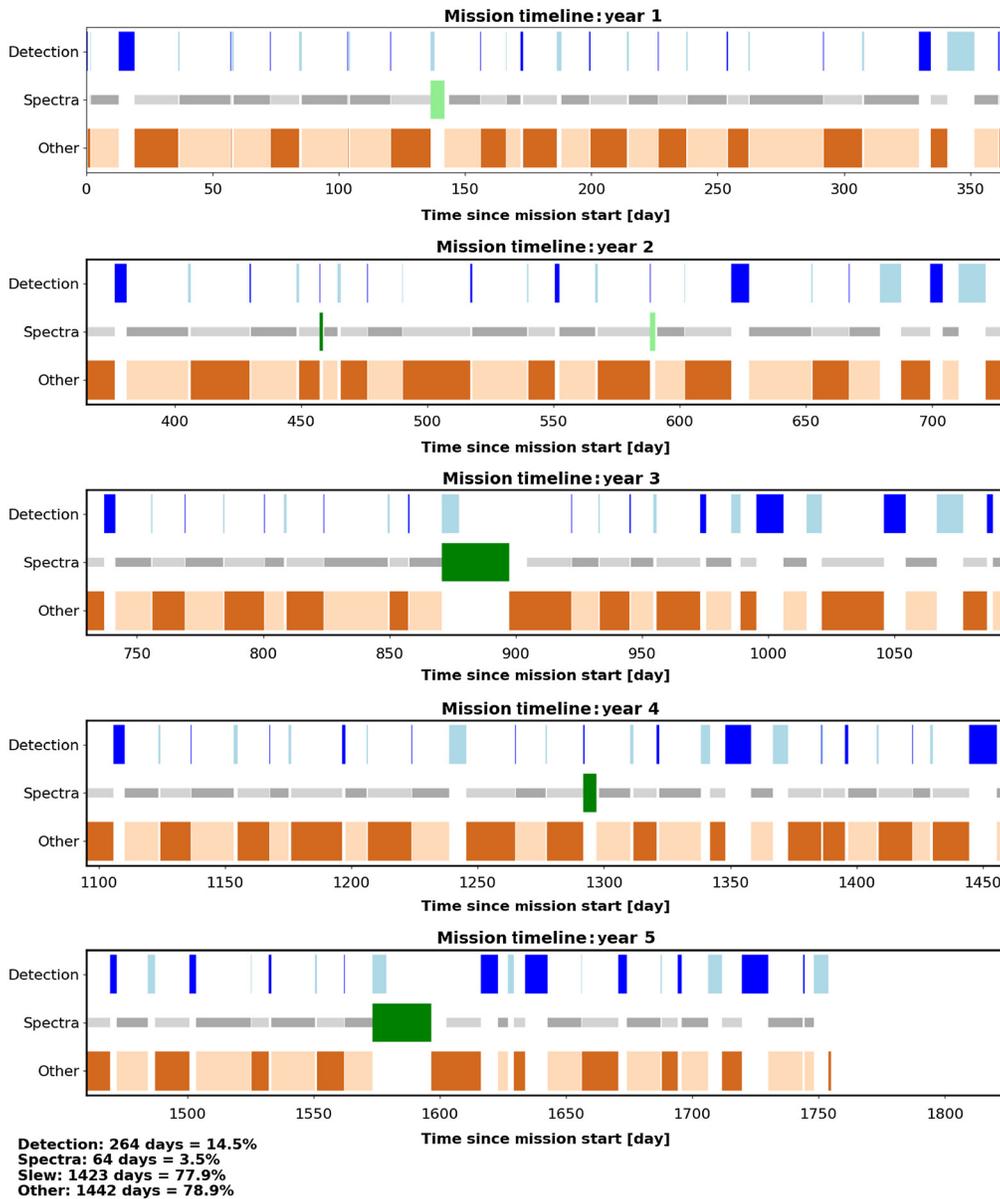


Fig. 10 The starshade-only scenario uses the starshade for blind search, spectral characterization, and follow-up detections for orbit determination. The detection observations with the starshade are shown in blue, and spectral characterizations are shown in green on the “spectra” row. Slews preceding detection observations are shown in gray on the “spectra” row. Telescope time during slews is used for the “other,” nonexoplanet timeline shown in brown, which receives ~80% of the mission time.

The coefficients of the cost function can be tuned to emphasize revisits versus unexplored targets. A coefficient tuning experiment using a Monte Carlo approach determined the highest yield for the starshade-only scenario. The best coefficients were different for the no-prior baseline case than for the EPRV prior case. The coefficient tuning for the starshade-only scenario was more sensitive than for the hybrid scenario.

Perfect prior knowledge case. The case of perfect prior knowledge for the starshade-only architecture is effectively the same as the perfect prior knowledge case for the hybrid starshade and coronagraph architecture since the starshade is used for spectral characterization while the coronagraph is used for the blind search.

2.4 Observations Using Partial Prior Knowledge via EPRV

A feature was implemented in EXOSIMS to flag a synthesized exoplanet if it could be known via precursor knowledge, based on the modeled sensitivity of a potential extreme precision radial velocity (EPRV) observing program. Using this feature, EXOSIMS promotes the known exo-Earths for characterization at the start of the mission. Those exo-Earths are placed at quadrature for characterization, for comparison to the perfect-prior and no-prior cases, on the assumption that the precursor knowledge provides ephemerides of sufficiently small uncertainty to allow scheduling characterization at an ideal epoch. EXOSIMS continues the blind search on all other targets, including targets that have EPRV null detections. A future refinement could employ EPRV null results for portions of the HZ, utilizing a joint ERPV-direct imaging completeness model.

2.4.1 EPRV capability assumptions

The sensitivity limit of the EPRV was found using insertion and recovery simulations by Newman and Plavchan.³ The simulated EPRV system was a ground-based (3 cm/s) RV instrument on a 10-m class telescope surveying ~ 53 HabEx targets in a 5 year, 25% time survey. The simulation assumed that stellar activity was perfectly corrected, the orbital parameters of each planet were estimated correctly, and in particular the right ascension of the ascending node was known perfectly. The simulated recovery results shown in Fig. 11 (blue) demonstrate an EPRV sensitivity from which we drew a heuristic limit.

The heuristic sensitivity limit used by EXOSIMS to determine whether a particular synthetic planet would have been detected around any sample star with effective temperature in the 3000 to 6800 K range and $K > \max [0.03 \text{ ms}^{-1}, \text{period (days)}/10^4 \text{ m/s/day m/s/day}]$ for all periods, where K is related to the planet and target star properties by

$$K = \frac{28.4 \text{ ms}^{-1}}{\sqrt{1 - e^2}} \sin i \left(\frac{m_p}{M_J} \right) \left(\frac{m_S}{M_\odot} \right)^{-2/3} P_{\text{yr}}^{-1/3}, \quad (2)$$

where m_p is the mass of the planet, m_S is the mass of the star, and P_{yr} is the period of the planet in years. Note that the exoplanet orbits synthesized by EXOSIMS had zero eccentricity.

The fraction of exo-Earths detectable with precursor EPRV to exo-Earths present around stars observed by the baseline HabEx 4H hybrid architecture for metric C is 0.29. The fraction for the

Fitting: Period, K, Time of Conjunction; Ecc (max 0.65) (1-9 planet systems)

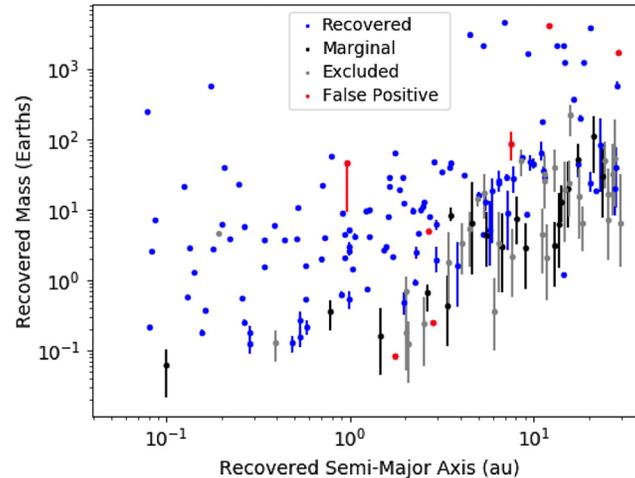


Fig. 11 The sensitivity limit of EPRV is seen by the recovered planets (blue) and the planets with marginal recovery (black). Planet recovery is mostly complete for $K > 1.0 \text{ ms}^{-1}$ at $P 10^4$ days.

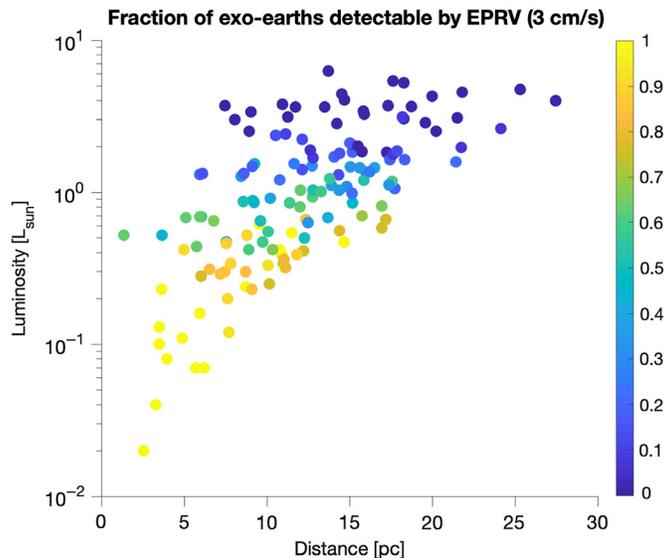


Fig. 12 Fraction per star of exo-Earths detectable by EPRV for the HabEx 4S starshade-only targets.

HabEx 4S starshade-only architecture for metric C is 0.45. The per-star fraction for the HabEx 4S starshade-only case is shown in Fig. 12. The benefit of EPRV decreases with higher luminosity stars. All of the present exo-Earths were detectable by EPRV for M stars and over 80% of the present exo-Earths were detectable by EPRV for K stars.

2.5 Input Parameters

Through a collaborative community-based activity, the ExSDET arrived at a widely accepted set of inputs for yield calculations. The ExSDET comparison of different yield methods shows very similar results for the same input assumptions. Uncertainties in yield are dominated by uncertainty in knowledge of astrophysics inputs.

2.5.1 Astrophysical inputs

The astrophysical assumptions and inputs are summarized here and discussed in greater detail in the ExSDET final report.² The astrophysics input parameters are summarized in Table 3.

The exoplanet occurrence rates were based on the EXOPAG SAG-13 power law model of Kepler data, as modified by Dulz et al.³⁸ for large mass and large semimajor axis planets.

Exo-Earth candidates were assumed to be on circular orbits and to reside within the conservative HZ, spanning 0.95 to 1.67 AU for a solar twin.³⁹ Exo-Earth candidates span radii ranging from $0.8a^{-0.5}$ to $1.4R_{\oplus}$, where a is the semimajor axis. This study focused on the exo-Earth candidate populations, using joint radius and semimajor axis distributions.

The stray light from binary stars in the final image plane was estimated⁴⁰ and included as an astrophysical noise source in exposure time calculations. AYO and EXOSIMS made no artificial cuts to the target list based on binarity. Instead, each code determined whether or not stray light noise made a target unobservable due to excessive SNR.

Zodiacal cloud brightness was estimated as a function of wavelength and ecliptic latitude and longitude by interpolating published tables.⁴¹ EXOSIMS specifically schedules each observation, enabling it to compute the zodiacal brightness based on the target's ecliptic coordinates on the date of the observation.

The exozodiacal light level used in the yield simulations was taken from the recent results of the Large Binocular Telescope Interferometer (LBTI) survey of exozodiacal dust.⁴² The nominal distribution function of the luminosity was used to draw an exozodiacal brightness for each star in each synthetic universe.

Table 3 Adopted astrophysical parameters.

Parameter	AYO	EXOSIMS	Description
η_E	0.24	SAG13 power law	Fraction of sunlike stars w/exo-Earth candidate
R_p		[0.6, 1.4] R_E	Exo-Earth candidate planet radius ^a
a		[0.95, 1.67] AU	Semimajor axis for solar twin
e		0	Eccentricity (circular orbits)
$\cos i$		[-1,1]	Cosine of inclination (uniform distribution)
ω		[0,2 π]	Argument of pericenter (uniform distribution)
M		[0,2 π]	Mean anomaly (uniform distribution)
Φ		Lambertian	Phase function
A_G		0.2	Geometric albedo of rocky planets
A_G		0.5	Geometric albedo of gas planets
z_c	23 mag asec ⁻²	Lindler model ^b	Average V band surface brightness of zodiacal light for coronagraph observations
z_s	22 mag asec ⁻²	Lindler model ^b	Average V band surface brightness of zodiacal light for starshade observations
x		22 mag asec ⁻²	V band surface brightness of 1 zodi of exozodiacal dust ^c
n		LBTI nominal distribution	Distribution of number of zodis for all stars

^aActual lower bound is $R_p > 0.8/\sqrt{a}$ in units of Earth radii R_E where a is the HZ-normalized semimajor axis in AU.⁴⁸

^bLindler zodiacal light model as a function of ecliptic latitude and longitude at observation time.

^cLocal zodi based on ecliptic pointing of telescope. On average, starshade observes into brighter zodiacal light.

^dFor solar twin. Varies with spectral type, as zodi definition fixes optical depth.

EXOSIMS and AYO use different input target list catalogs, but comparisons⁴³ have demonstrated that the minor differences between target lists are inconsequential; target list differences result in yield differences at the few percent level. This is due to the fact that the two target lists, while independently generated, depend heavily on the *Hipparcos* survey, and thus contain largely the same stars out to 30 pc. The overall yield is robust to inaccuracies in individual stellar parameters when large numbers of stars are surveyed.

EXOSIMS uses the ExoCat-1 catalog,³⁶ as stored in the MissionStars table of the NASA Exoplanet Archive hosted by the NASA Exoplanet Science Institute.⁴⁴ Missing photometric information for targets is optionally synthesized by interpolating over Eric Mamajek's Mean Dwarf Stellar Color and Effective Temperature Sequence.⁴⁵ ExoCat-1 is further modified in the case when the binary leakage model is being used with updated information from the Washington Double-Star catalog, maintained at the U.S. Naval Observatory.

2.5.2 Instrument parameters

We calculate yields for four different mission concepts: HabEx baseline 4m Hybrid (4H) architecture, HabEx 4m starshade-only (4S) architecture, HabEx 4m coronagraph-only (4C) architecture, and LUVOIR B (coronagraph-only) architecture. LUVOIR B uses an 8m primary (max distance from tip-to-tip) that provides a 6.7-m-diameter clear circular aperture and equivalent collecting area for coronagraphic observations. The instrument parameters for the HabEx and LUVOIR B architectures are summarized in Table 4. The contrast curves of the HabEx and LUVOIR coronagraphs and the HabEx starshade are in their respective final reports.

The simulated missions were 5 years in duration with 50% of the mission allocated to exoplanet observing. The 50% allocation to exoplanets was kept throughout the mission.

Table 4 Instrument parameters.

Parameter	LUVOIR B	HabEx 4C	HabEx 4S
Primary diameter (m)	8.0	4.0	4.0
Obscuration factor	0.14	0	0
Integration time limit	60 days	60 days	60 days
Starlight suppression performance			
Raw contrast floor	1×10^{-10}	1×10^{-10}	1×10^{-10}
Raw contrast stability	1×10^{-11}	2×10^{-11}	2×10^{-11}
Postprocessing factor	0.25	0.29	0.29
Systematic noise floor	26.5	26.5	26.5
Core throughput	0.46	0.52	0.7
Photometric aperture	$0.8 \lambda/D$	$0.7 \lambda/D$	$0.7 \lambda/D$
IWA _{0.5}	$3.9 \lambda/D$	$2.4 \lambda/D$	$1.15 \lambda/D^a$
IWA _{0.1}	$1.5 \lambda/D$	$1.5 \lambda/D$	$0.9 \lambda/D^a$
OWA	$60 \lambda/D$	$26 \lambda/D$	$40 \lambda/D$
Imaging channel 1 ^b			
Noncoronagraph throughput	0.17	0.28	0.45
Bandwidth	20%	20%	100%
Imaging channel 2 ^c			
Noncoronagraph throughput	0.39	0.42	—
Bandwidth	20%	20%	—
Spectral channel			
Noncoronagraph throughput	0.39	0.42	0.39
Bandwidth	20%	20%	108%
Detectors			
Quantum efficiency	0.9	0.9	0.9
Photon counting efficiency	0.75	0.75	0.75
Dark current (e/s)	3×10^{-5}	3×10^{-5}	3×10^{-5}
Read noise (e/pix)	0	0	0
Clock-induced charge (e/s)	1.3×10^{-5}	1.3×10^{-5}	1.3×10^{-5}
Starshade			
Starshade thrust (mN)	—	—	1040
Starshade slew I_{sp} (s)	—	—	3000
Starshade stationkeeping I_{sp} (s)	—	—	308
Starshade wet mass (kg)	—	—	11,180
Starshade dry mass (kg)	—	—	4550
Starshade separation (km)	—	—	76,600

^afor $\lambda = 975$ nm.^bBlue imaging band for HabEx, UV imaging band for LUVOIR B.^cRed imaging band for HabEx, Visible imaging band for LUVOIRB.

3 Results

3.1 Metrics

The desired end product of the direct imaging of exo-Earth candidates is not only their discovery and not only determination of an orbit in the habitable zone, but also a full spectra to reveal potential biomarkers. A starshade and coronagraph arrive at a full spectra in different ways and are suited to different metrics. Three spectral characterization metrics were evaluated. The first metric was used by LUVOIR to report yield of exo-Earth candidates, i.e., targets with interesting enough spectra in the 20% bandwidth to merit a full spectra, such as a water line. The second metric is a full spectra tailored for a coronagraph, as suggested in the LUVOIR final report. The third metric is a full spectra as seen by a starshade and was used in the HabEx final report.

Figure 13 shows the simulated spectra of an Earth-Sun twin system at 7.5 pc with five exozodis and added noise, as observed by each of the three metrics. Noise simulations were performed with a publicly available high-contrast imaging noise model.^{46,47} This tool includes noise treatments for host stellar light leakage, exozodiacal and zodiacal light, dark current, read noise, and clock-induced charge noise, and relevant model parameters were taken from

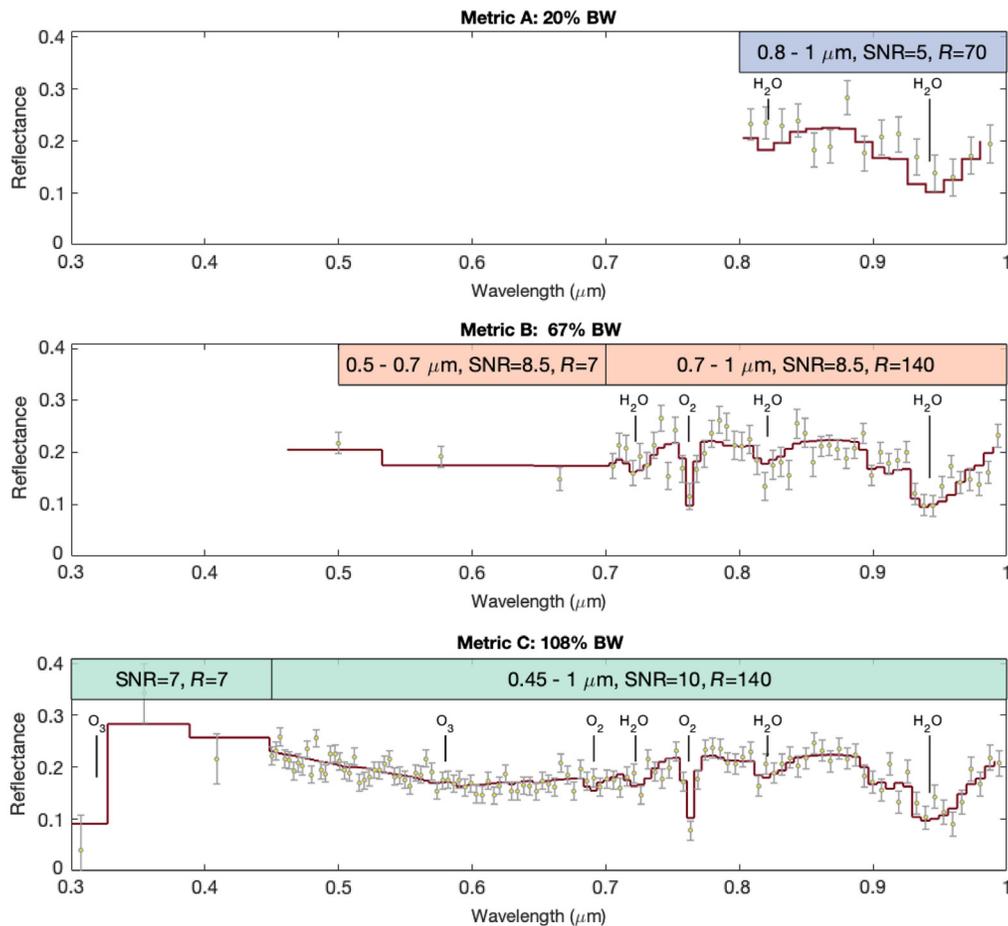


Fig. 13 Synthetic spectra with added noise for an Earth-Sun twin at 7.5 pc with a 4m HabEx. The magenta line is the Earth's spectrum at the goal spectral resolution. (a) A single coronagraph channel of 20% bandwidth for $R = 70$, SNR = 5 used 107 h of integration time. (b) The coronagraph 20% bandpasses have some overlap and use 392 h of cumulative integration time. HabEx 4m coronagraph, 500 to 700 nm: $R = 7$, SNR = 8.5; 700 to 1000 nm: $R = 140$, SNR = 8.5. Cumulative integration time is 392 h. (c) The starshade spectra for $R = 140$, SNR = 10 used 390 h of integration time.

the HabEx final report. The features visible in these simulated spectra are from H₂O, O₂, and O₃. The resulting spectra provide various levels of fidelity of the final science product.

Characterization metric A facilitates a quick search for the deep water line at 940 nm and the shallower water line at 820 nm. It is used by the LUVOIR mission concept to vet candidates for follow-up full spectra observation and is used as their official yield metric. The 20% bandwidth is achievable by a coronagraph in one observation. The spectral resolution of $R = 70$ and $\text{SNR} = 5$ is sufficient for detection of the waterline while keeping the integration time minimal. The simulated spectra were achieved by the HabEx coronagraph in 107 h of integration time.

Characterization metric B is a full spectra tailored to minimum required SNR and spectral resolution to allow for more favorable integration times by coronagraphs. From 500 to 700 nm, the spectral resolution of 7 and SNR of 8.5 are sufficient to detect atmospheric scattering indicated by a rise in the blue end. From 700 to 1000 nm, the spectral resolution of 140 and SNR of 8.5 are considered minimal for adequate detection of species of interest.⁴⁸ The simulated earth spectra show three water lines, a deep oxygen line, and potential Raleigh scattering toward the blue end. The cumulative integration time is 392 h for a sequence of four subspectra at 20% bandwidth with the HabEx 4C coronagraph.

Characterization metric C is the full spectra attainable by the HabEx starshade. The baseline HabEx architecture uses the starshade to achieve a continuous spectra from 300 to 1000 nm in a single integration time to $\text{SNR} = 10$ at 650 nm. The starshade IFS produces a spectra from 450 to 1000 nm with spectral resolution $R = 140$. The starshade UV camera uses a grism to achieve $R = 7$ from 300 to 450 nm. An SNR of 5 is sufficient to detect Raleigh scattering in the UV, though a higher SNR is likely to be achieved during the integration time that is set for the $R = 140$, $\text{SNR} = 10$ region. The simulated earth spectra show three water lines, a deep oxygen line, a shallow oxygen line, a deep ozone feature in the UV, and Raleigh scattering through the UV. The simulated spectra for metric C were achieved by the HabEx starshade in 390 h of integration time.

3.2 Results Summary Table

Table 5 compares the average number of exo-Earths spectrally characterized by four different mission concepts. The targets spectrally characterized in the hybrid and coronagraph-only scenarios also received orbit determination; about half to two-thirds of the targets spectrally characterized in the starshade-only scenarios received orbit determination. The potential of each mission architecture is compared for the three different levels of spectral characterization, corresponding to the three figures of merit defined in the previous section. The fidelity of spectral characterization improves from metric A to B to C.

Table 5 Average number of characterized exo-Earths for four mission architectures, three levels of prior knowledge, and three metrics. Entries are for 50% of a 5-year mission.

Metric type	A			B			C		
Metric	20% BW, SNR = 5			4 × 20% BW, SNR = 8.5			100% BW, SNR = 10		
Prior	None	EPRV	Omni	None	EPRV	Omni	None	EPRV	Omni
HabEx 4H	9	12	32	8	10	16	7 ^a	9	13
HabEx 4S	7	10	32	5	7	16	4	6	13
HabEx 4C	7	8	13	2	3	4	2 ^b	2 ^b	3 ^b
LUVOIR B	15 ^c	17	33	8	9	12	6	7	9

^aAYO also evaluated HabEx 4H under no prior information, obtaining a yield estimate of 8.

^bUV (300 to 450 nm) is not supported.

^cAYO evaluated LUVOIR B for 40% of a 5-year mission and no prior information, obtaining a yield estimate of 28.

For each spectral characterization depth (i.e., each “metric”), the table indicates the number of planets characterized by the four mission concepts, assuming no prior knowledge of exo-Earths around the target stars (prior = none), partial knowledge of which stars have exo-Earths around them and what their orbits are (prior = EPRV), and perfect knowledge of exo-Earth presence and orbit for all targets (prior = Omni). Results from the AYO and EXOSIMS yield estimation codes are indicated. AYO results are only given for metric A for LUVOIR and metric C for HabEx, in the no prior knowledge case, to provide context and consistency with the HabEx and LUVOIR final reports and ExSDET common comparison.

For the LUVOIR B metric A results, EXOSIMS has a lower yield than AYO due to the inefficiencies of the coronagraph blind search. The coronagraph initiates some search observations too late in the mission to allow for sufficient epochs to perform orbit determination. Front-weighting the exo-planet mission portion early in the mission could provide more epochs for orbit determination and increase the number of stars promoted for characterization. In addition, the integration times used in the coronagraph blind search could be optimized, such as the way AYO optimizes integration time across all targets for a given mission time, to increase yield. The omniscient prior case demonstrates that EXOSIMS can achieve an upper bound similar to the AYO yield.

For the HabEx 4H baseline mission, the EXOSIMS results are slightly lower than the AYO results due to the same inefficiencies of the scheduled blind search. The evaluation of the omniscient upper bound proved to be target limited, so the slight discrepancy in yield is not likely due to the realistic observation scheduling of the starshade.

We discuss hereafter only the EXOSIMS results.

3.3 Yield as a Function of Elapsed Mission Time

The cumulative yield as a function of mission time for all of these cases is shown in Figs. 14 and 15. Figure 14 is normalized by the omniscient prior yield at end of mission. This normalization highlights the shape of the cumulative yield curves, particularly in the first 2 years of the mission when the EPRV prior knowledge achieves spectral characterizations more quickly in the mission. This is seen most dramatically in HabEx 4H and 4S with metrics B and C and significantly for coronagraphs for metric A, both HabEx 4C and LUVOIR B.

EPRV prior knowledge impacts the starshade architectures more strongly than the coronagraph-only architectures in terms of early mission and end-of-mission yield. The stars with high EPRV sensitivity consistently coincide with the stars with high starshade-only yields. The EPRV prior knowledge allows these stars to be bypassed and the blind search to reach to more distant, slightly lower completeness targets. Use of the EPRV null results, particularly for the high-sensitivity and high-completeness targets, would have the strongest impact on the starshade-only scenario, by saving three to six starshade slews for those targets without exo-Earths (a proportion of $1 - \eta_E$), a possible savings of ~ 30 to 60 slews.

Target exhaustion is indicated when the cumulative yield flattens out, which is seen for the omniscient case for most of the architectures and metrics. The exception is the HabEx hybrid and starshade for metric A, which is on the verge of target exhaustion; in most DRMs these two architectures were observing none or only one target in the final year of the mission.

The non-normalized cumulative yields are shown in Fig. 15 on the same scale to allow for comparison of relative yield across architectures and metrics. The relative yield increases and acceleration is most dramatic for the HabEx hybrid architecture across all three metrics. The EPRV prior cases are accelerated by three to six times that of the no prior cases. It is important to remember that the results from this simulation allow for a relative comparison of yield. Yield is nearly linear with occurrence rate,¹⁸ and any future refinements in exo-Earth candidate occurrence rates will affect all architectures similarly.

From these yield timelines in Fig. 15, we extract the time to achieve the 50% of the final average yield for each of the cases and present them in Fig. 16. The no-prior cases reached half of their mission yield earlier than halfway through the 5-year mission. This is understandable given the diminishing returns of the blind search as the remaining targets become more challenging. All of the EPRV cases reach 50% of their final yield within the first year of the mission. The EPRV prior knowledge enables a significant acceleration of yield for all architectures.

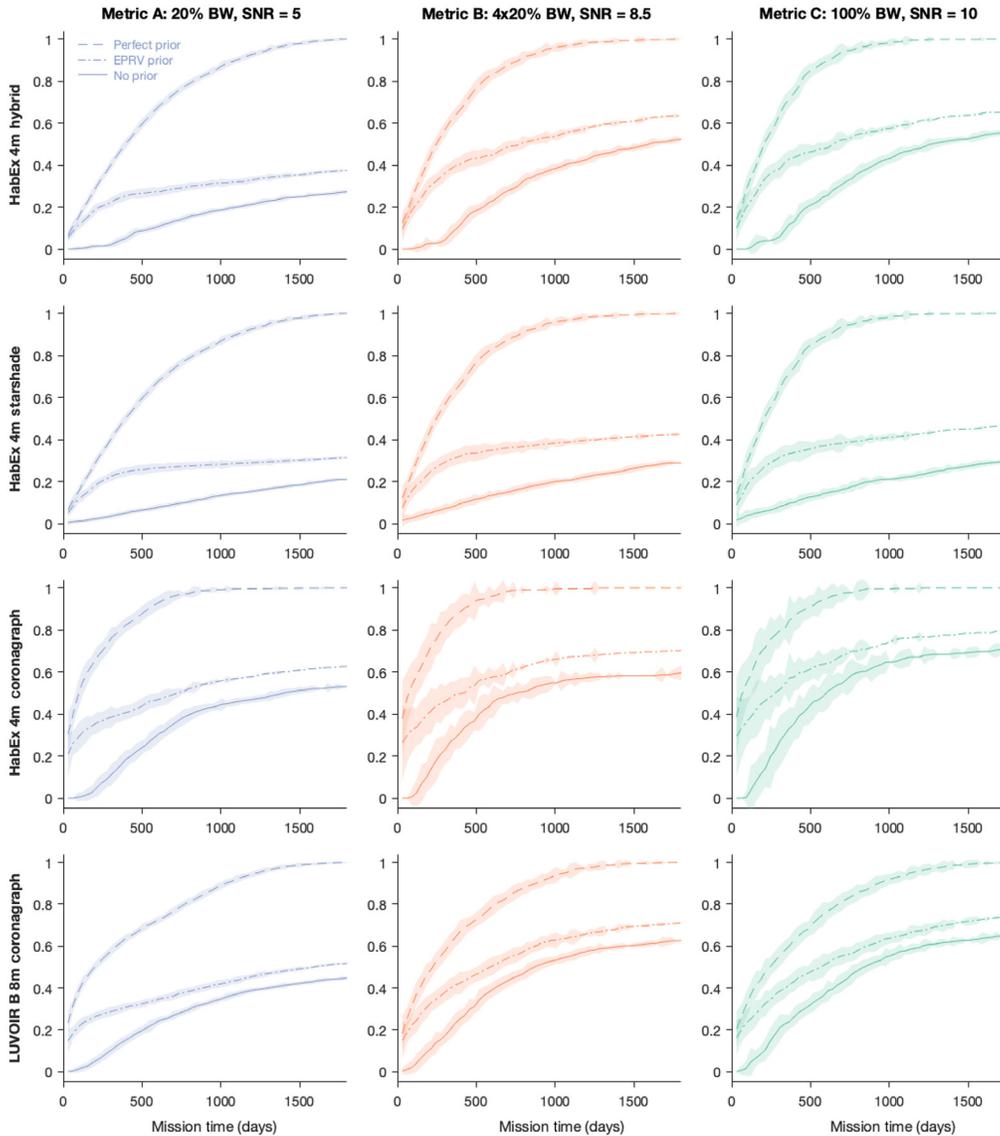


Fig. 14 The cumulative yield over the 5-year mission is plotted for the four mission architectures (rows), the three yield metrics (columns), and three levels of prior knowledge (the top line is the perfect prior, the middle line the EPRV prior, and the bottom line is no prior). The shaded regions indicate the 1σ standard deviation. Each case is normalized by the total yield of its perfect prior.

We can also examine the acceleration of yield by benchmarking the time to achieve the final yield of the no prior case (Fig. 17) by the other prior knowledge cases. The cluster of square markers shows the end of mission average yields for the no-prior cases. The color of the marker corresponds to the yield metric, and the mission architectures are not indicated. The EPRV cases achieved the same yield as the no prior case in an average of 750 days with a standard deviation of 750 days. With the help of EPRV, the science yield that takes 5 years with blind search alone could be accomplished in half that time and for the starshade-only architecture could be accomplished in a single year.

EPRV prior knowledge increased yield speed by promoting some targets for characterization at the mission start. For these targets, the integration time of at least three detection observations was saved and allowed the blind search to examine targets of slightly longer integration times, thereby slightly increasing overall yield. The blind search initiated detection observations for the long integration time targets ~ 200 to 300 days earlier than in the no-prior case; this allowed more orbit determination epochs to fit in the remaining mission time and contributed to increased yield.

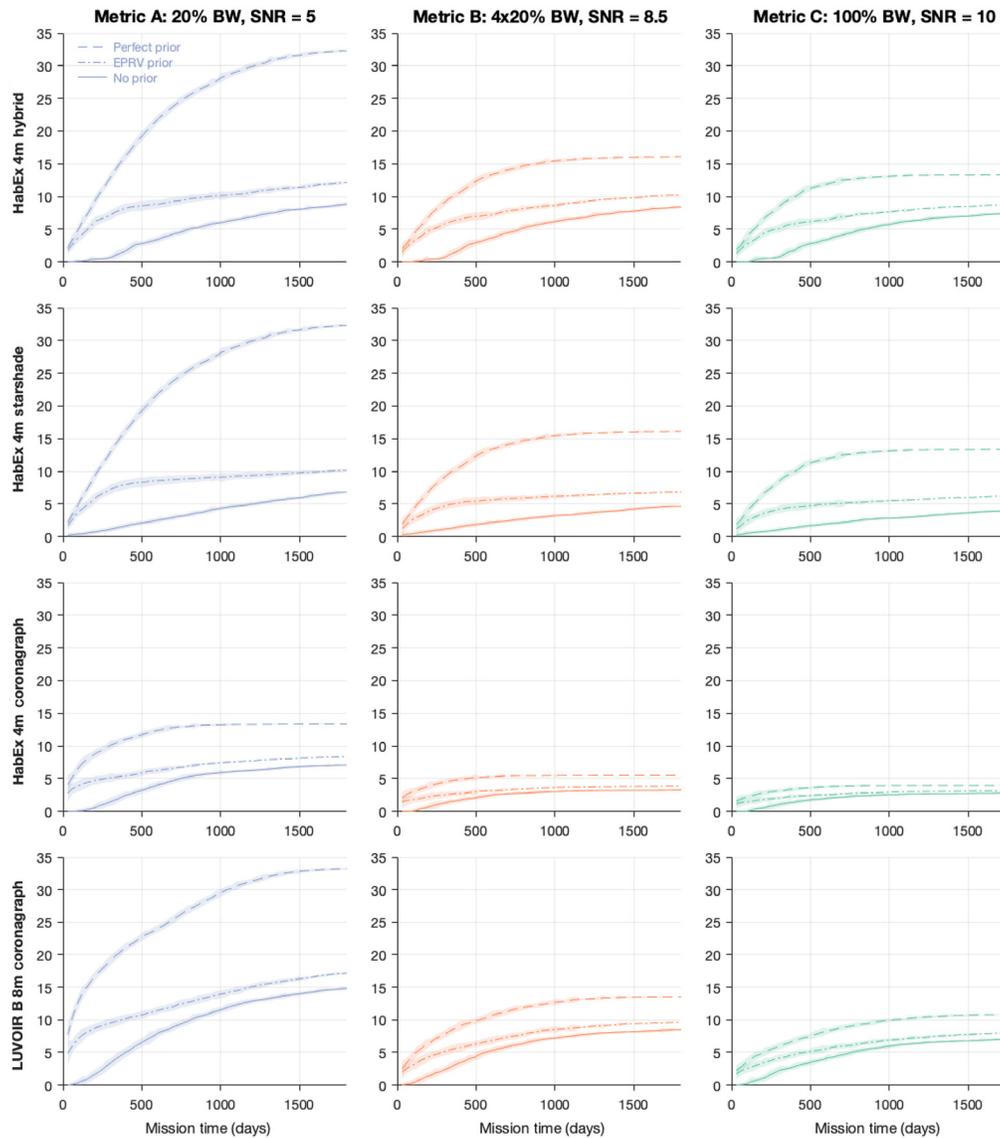


Fig. 15 The cumulative yield over the 5-year mission time is plotted for the four mission architectures (rows), the three yield metrics (columns), and three levels of prior knowledge using the same scale (the top line is the perfect prior, the middle line the EPRV prior, and the bottom line is no prior). The shaded regions indicate the 1σ standard deviation.

3.4 Interpretation of Comparisons and Trends

3.4.1 Metric A

Metric A corresponds to an $R = 70$, $\text{SNR} = 5$ spectrum over a single 20% band centered at 940 nm, searching for water vapor.

Without any prior knowledge of which target stars host exo-Earths, nor what their orbits are, the starshade-only case provides the lowest yield, as expected. That is because the starshade needs to travel to all targets to search for earth-like planets, immediately take a spectrum of any exo-Earths candidates detected, and then revisit those systems multiple times to estimate the orbit. This makes for a very inefficient observing scenario, and if there is no starshade refueling in space, the yield is severely limited by the number of available slews (starshade transits). Missions having a coronagraph benefit from its pointing agility and are able to detect and search for water around significantly more exo-Earths. The yield goes from 7 to 15 when increasing the

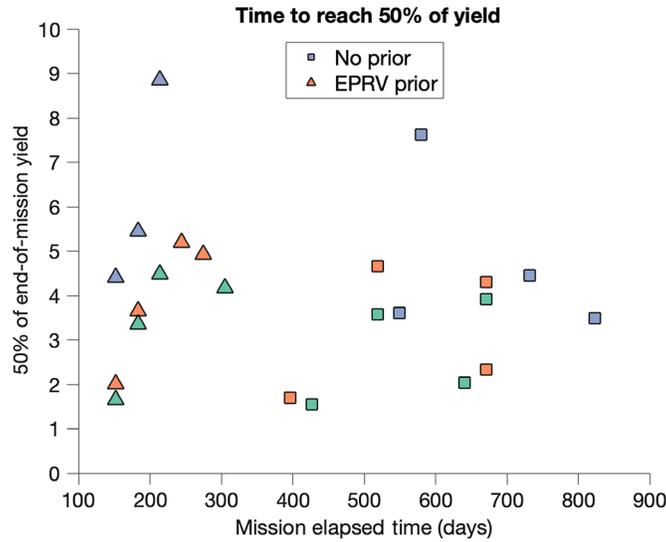


Fig. 16 The mission time at which the 50% value of the yield is achieved, and the 50% value of the yield are plotted without reference to mission architecture. Metric A is blue, metric B is orange, and metric C is green. The marker shape indicates the level of prior knowledge.

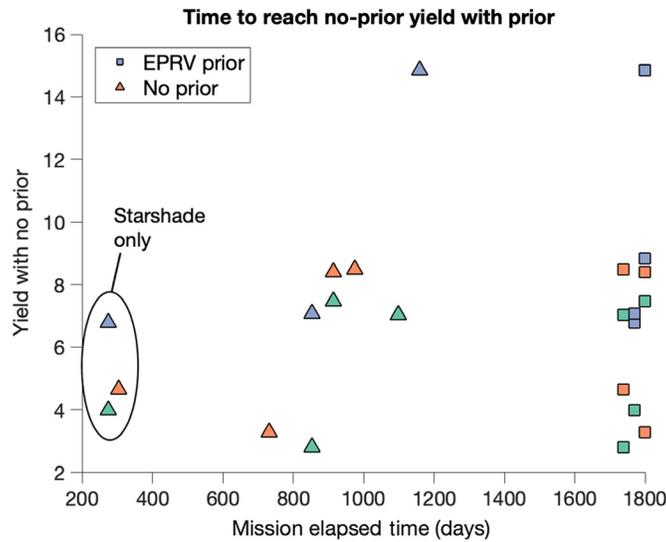


Fig. 17 The mission time to reach the final yield of the no-prior case is plotted for the EPRV prior case and no-prior case. The yields for a particular mission architecture are flat horizontally and the time to achieve that yield is indicated by a square marker for the no prior case and a triangle marker for the EPRV case. Metric A is blue, metric B is orange, and metric C is green.

effective aperture from 4m (HabEx 4C) to 6.7m (LUVOIR B), roughly in line with the D^2 dependence of yield versus diameter already predicted for coronagraph-only missions.¹⁸ The HabEx hybrid (4H) architecture, which uses a coronagraph for detection and a starshade for spectral characterization benefits from the sensitivity gain of the starshade instrument, can detect and characterize a few more systems than HabEx 4C in the same allocated total mission time.

In the case where some prior knowledge is available, e.g., from the particular set of precursor EPRV measurements described above, the yield increases for all missions, but more strongly for missions using a starshade (4H and 4S). In particular, the HabEx 4H yield goes up to 12 (30% increase), compared to 17 (13% increase) for LUVOIR B.

Finally, assuming perfect prior knowledge of exo-Earths and their orbits around the targets, the starshade-based missions HabEx 4H and 4S can search for water around 32 exo-Earths, similar to LUVOIR B's coronagraph on a larger telescope. When every starshade slew counts, the advantages of the starshade (smaller IWA in λ/D units and higher throughput) make up for the reduced aperture of HabEx versus LUVOIR B.

3.4.2 Metrics B and C

Metrics B and C offer a much more detailed and broader spectral characterization than metric A (see Fig. 13). Metric C in particular covers the wavelength range of 0.3 to 1 μm and offers high quality spectral characterization of exo-Earths at $R = 140$ and $\text{SNR} = 10$. The results obtained for metrics B and C are fairly similar in the qualitative trends that they show.

Without any prior knowledge of which target stars host exo-Earths, nor what their orbits are, the architectures providing the lowest yield are HabEx 4C and 4S. HabEx 4C's yield is low because of the exposure time required to cover the whole spectral range using several 20% bandwidth observations, stitching 20% bandwidth spectra obtained one at a time. The HabEx 4S yield remains low as well because of the inefficiency of exo-Earth blind searches with starshade-only missions. The mission architectures that provide the highest yield are now HabEx 4H and LUVOIR B, which both provide about the same number of exo-Earths with deep spectral characterization. Compared to the numbers obtained for metric A, HabEx 4H yield remains close, whereas the coronagraph-only LUVOIR B architecture sees a more significant reduction in yield, due again to the number of sequential observations required to cover the broader spectral range of metrics B and C.

Similar to what was observed with metric A, having either some (EPRV) or full (OMNI) prior knowledge of exo-Earths benefits starshade missions the most. For metrics B and C, with prior knowledge at the assumed EPRV level or higher, HabEx 4H provides the largest number of exo-Earths with fine spectral characterization. As for the coronagraph-only missions, the total yield benefits slightly from EPRV prior exo-Earth detections, and the timeline for spectral characterizations is accelerated.

These results highlight that the optimum architecture to be used for fine exo-Earths spectral characterization strongly depends on the level of prior knowledge:

- If no exo-Earths are known, our results suggest that the highest yield mission will have to include a coronagraph. It could either be a coronagraph-only mission on a larger telescope or a hybrid coronagraph and starshade mission on a smaller telescope.
- If a large enough fraction ($>30\%$) of target stars have precise enough RV (or astrometric) measurements to detect (or reject) the presence of exo-Earths around them, then our results suggest that the highest yield mission must include a starshade. In particular, a hybrid starshade and coronagraph mission provides the same yield as a coronagraph-only mission on a significantly larger telescope, and even higher yield if prior knowledge is available on most systems. Finally, in the asymptotic case where prior knowledge is available on all targets, a starshade-only mission does just as well as a hybrid.

4 Conclusions

We discussed the importance of observation scheduling for simulating science yields for starshade-based mission architectures and detailed the MCMS EXOSIMS that employs dynamic scheduling. We presented observing scenarios for starshade-only (HabEx 4S), coronagraph-only (HabEx 4C and LUVOIR B), and hybrid coronagraph and starshade (HabEx 4H) mission architectures and the modeling approach for blind search, orbit determination, and spectral characterization. We utilized three spectral characterization metrics: metric A (20% BW, $R = 70$, $\text{SNR} = 5$) highlights the strength of coronagraph architectures to perform a blind search and vet targets for a water line feature; metric B (500 to 700 nm: $R = 7$, $\text{SNR} = 8.5$; 700 to 1000 nm: $R = 140$, $\text{SNR} = 8.5$) is a broad spectrum tailored to minimize integration time for coronagraphs that are limited to a sequence of 20% bandwidth subspectra; metric C (108% BW, $R = 140$, $\text{SNR} = 10$) provides the deepest, broadest spectra, which is the strength of starshades.

We evaluated the yield of exo-Earths for the four mission architectures against the three metrics for three levels of prior knowledge.

We established an optimistic upper bound with the omniscient prior case that evaluates the ability of the mission architecture to characterize exo-Earths with realistic mission constraints and scheduling. We evaluated a lower bound, in which no prior information on exo-Earths is known, and the mission architecture must perform a blind search, determine orbits, and spectrally characterize targets. Yields for the lower bound case could be mildly improved with more efficient search implementations and finer tuning of optimization parameters. EPRV prior knowledge increases yield by up to 50% and accelerates by three to six times the mission time to achieve half of the yield. With the help of EPRV, the science yield that takes 5 years with blind search alone could be accomplished in half that time.

For metric A, 20% BW with $R = 70$ and $\text{SNR} = 5$, EPRV prior knowledge increases yield for all missions, and most strongly for starshade architectures (4H and 4S). In particular, the HabEx 4H yield goes up to 12 (30% increase), compared to 17 (13% increase) for LUVOIR B. With perfect prior knowledge and only scheduling limitations, the starshade-based missions HabEx 4H and 4S can search for water around 32 exo-Earths, similar to 33 exo-Earths LUVOIR B's coronagraph. When every starshade slew counts, the advantages of the starshade (smaller IWA in λ/D units and higher throughput) make up for the reduced aperture of HabEx as compared to LUVOIR B.

For metrics B and C, which provide more detailed and broader spectra than metric A, the lowest yields are for the no-prior knowledge cases of HabEx 4S and 4C. The mission architectures that provide the highest yield are HabEx 4H and LUVOIR B, which provide ~ 8 characterized exo-Earths. Prior knowledge benefits starshade missions the most. For metrics B and C, with prior knowledge at the assumed EPRV level or higher, HabEx 4H provides the largest number of exo-Earths with fine spectral characterization.

Independent of architecture, prior knowledge increases total yields and accelerates the time to achieve half the yield. The prior knowledge enables high impact science results to occur earlier in the mission. Achieving characterization results earlier in the mission is more resilient against detector degradation due to radiation exposure and other risks of mission failure. Earlier characterization results could lead to changes in observing priorities, such as affording more time for GO science, or providing sufficient epochs in the remaining mission time to perform follow up characterizations, particularly critical for starshade scheduling. Follow-up characterizations could include finer spectral characterization at different illumination phases, with higher SNRs, and/or over different wavelength ranges than the nominal 0.3 to 1 μm , i.e., with extension to the near UV (down to 200 nm) and to the near-infrared for the most favorable targets. The access to different illumination phases in total intensity or in polarized light is key to the detection of surface oceans (e.g., HabEx Final report Sec. 3.1.4). The access to longer infrared wavelengths, looking for CO_2 and methane is key to ruling out the most likely false-positive scenarios, i.e., nonbiological O_2 or CH_4 build-up processes in the atmosphere of HZ rocky planets (e.g., HabEx report Sec. 3.1.3.2 and LUVOIR report Sec. 3.3.3). Precursor spectral targets on mission day 1 ensure robust scheduling opportunities at optimal viewing epochs, particularly for starshade arrival. Prior knowledge makes all the mission architectures more nimble and powerful.

5 Appendix A: EXOSIMS Modules

The EXOSIMS module functionality is split as follows:

- Optical system: a description of the observatory optics, all science instruments, starlight suppression systems (coronagraphs and occulter), and methods for computing integration times for different targets to varying levels of contrast.
- Star catalog: an input stellar catalog with astrometric and photometric information on potential target stars.
- Planet population: a statistical description of the planet population including factorized or joint density functions for all planetary physical and orbital parameters.

- Observatory: a description of the observatory spacecraft, its orbit, and methods for computing target look vectors, keepout regions, and fuel use for starshade slewing and stationkeeping.
- Planet physical model: methods for converting between planet mass and radius, and descriptions of planet phase functions.
- Time keeping: methods for tracking mission time throughout the simulation, including time specifically devoted to exoplanet science.
- Zodiacal light: descriptions of local and exo-zodiacal light.
- Background sources: models for false positives generated from background sources.
- Postprocessing: models for the effects of postprocessing.
- Target list: methods for filtering the full star catalog to a list of targets for a specific instrument and mission.
- Simulated universe: methods to sample the planet population and assign simulated planets to stars from the target list.
- Survey simulation: methods for simulating the full mission, including logic for automated scheduling of targets.
- Survey ensemble: methods for generating a Monte Carlo ensemble of survey simulations.

The completeness, zodiacal light, and optical system modules are described in additional detail in Ref. 15.

6 Appendix B: Weighted Linear Cost Function Scheduler

This appendix describes the WLCF scheduler in EXOSIMS. The WLCF scheduler is related to the well-known TSP but with key differences that make a time-windowed but dynamic approximation more effective.¹¹ This time-windowed approximation starts from a directed, cyclic graph containing the target stars as nodes and slews as edges. A merit function is defined to represent the cost of transitioning from observing the current target (i) to observing the next target (j):

$$A_{ij} = \sum_k a_k c_k(i, j), \quad (3)$$

where each $c_k(i, j)$ is an independent cost/reward term (depending on the current and next targets), and the $\{a_k\}$ are weighting factors. We can form a square adjacency matrix $[A]_{ij} = A_{ij}$ that expresses the cost of any possible single slew. A path through the graph encodes any subsequent observing sequence, and the sum of the A_{ij} values on the edges of the path is the cost of these observations. While a single fixed adjacency matrix cannot encode the dynamic constraints and system evolution described above, it represents a valid approximation for a fixed number of observations forward from a given epoch. We can thus approximate a small number of future observations using the current adjacency matrix and determine the minimum cost path through the graph for that fixed number of steps. The root node of this graph is then selected as the next target, the observation is simulated, and the full adjacency matrix and observing plan are recalculated after each observation.

The EXOSIMS WLCF scheduler implementation uses a cost function with four terms. The WLCF uses a cost penalty term for slewing the starshade (function of the angle between the current target and the next), an observation completeness reward term, and two more regularization terms that increase overlapped mission time, balancing weight to previously unobserved stars (to discover exoplanets), and stars observed less frequently (to complete revisits for orbit determination). The terms are described in more detail:

1. For an external occulter, the first term encodes the cost of slewing the occulter between the two targets, either via a direct calculation of fuel use³¹ or via the angle between the telescope look vectors to the two targets^{2,11} (used in this study). For an internal coronagraph,

it is assumed that all telescope repointings have approximately equal cost, and this term is set to zero.

2. The second term represents the reward metric of the observation, encoded as 1 minus the target completeness, with negation serving to turn the reward into a cost.
3. The third and fourth terms bias toward previously unobserved stars and stars observed less frequently. The weights on these terms increase with mission elapsed time, making it less costly to make new observations (versus revisits) toward the end of the mission execution, and more urgent to revisit targets in need of orbit determination.
4. When including prior knowledge targets, two terms are added. The fifth and sixth terms are analogous to the third and fourth terms but for the targets with prior knowledge. This allows for prioritization of the targets with EPRV detected exoplanets against the blind-search targets in terms of initial visits versus revisits.

7 Appendix C: Target List

The union of the target lists for the four architectures, evaluated for metric B, is shown in Table 6. Metric B is for $4 \times 20\%$ bandwidth spectra at an SNR of 8.5 and provides a common ground for comparison of full spectra characterization yield by coronagraphs and starshades. The target list is sorted by descending characterization yield of the LUVOIR B coronagraph-only architecture. The characterization yield is the average number of unique exo-Earths characterized per star. The associated integration time is the mean integration in days spent on that star for spectral characterization.

Table 6 Target list for four architectures (LUVOIR B, HabEx 4H hybrid, HabEx 4S starshade-only, and HabEx 4C coronagraph-only) for metric B ($4 \times 20\%$ bandwidth at SNR = 8.5). Char is the number of unique exo-Earths characterized for that star, averaged over the Monte Carlo DRM ensemble. T_{int} is the associated spectral integration time in days. EEID is in milliarc-seconds.

Target	Vmag	d (pc)	L (L_{\odot})	EEID	LUVOIR B		HabEx 4H		HabEx 4S		HabEx 4C	
					T_{int}	Char	T_{int}	Char	T_{int}	Char	T_{int}	Char
HIP 84478	6.33	6.0	0.16	66.7	6.2	0.33	2.9	0.25	2.6	0.18	5.0	0.11
HIP 49908	6.60	4.9	0.11	67.1	4.6	0.32	0.9	0.16	1.8	0.24	3.5	0.14
HIP 108870	4.69	3.6	0.23	131.4	1.0	0.28	0.8	0.27	1.1	0.27	7.9	0.17
HIP 104217	5.95	3.5	0.10	89.1	1.6	0.26	0.7	0.22	2.3	0.28	3.6	0.24
HIP 54035	7.49	2.5	0.02	61.7	0.3	0.23	0.2	0.16	0.6	0.06	—	—
HIP 84405 B	4.33	6.0	0.69	139.7	8.5	0.22	2.3	0.11	0.9	0.03	—	—
HIP 64924	4.74	8.6	0.87	109.3	3.8	0.21	0.8	0.08	1.5	0.09	2.2	0.04
HIP 3765	5.74	7.5	0.30	73.3	6.3	0.20	1.8	0.21	1.4	0.14	2.6	0.05
HIP 12114	5.82	7.2	0.29	74.5	4.6	0.20	1.9	0.18	2.9	0.24	4.0	0.06
HIP 72659 A	4.67	6.8	0.65	119.3	3.4	0.20	0.6	0.06	0.9	0.04	4.3	0.06
HIP 114622	5.57	6.5	0.31	84.5	2.4	0.19	1.2	0.13	1.3	0.14	7.7	0.17
HIP 13402	6.05	10.4	0.42	62.4	5.6	0.18	2.1	0.11	—	—	—	—
HIP 8102	3.49	3.7	0.52	197.4	0.4	0.17	0.9	0.17	1.0	0.25	4.0	0.29

Table 6 (*Continued*).

Target	Vmag	d (pc)	L (L_{\odot})	EEID	LUVOIR B		HabEx 4H		HabEx 4S		HabEx 4C	
					T_{int}	Char	T_{int}	Char	T_{int}	Char	T_{int}	Char
HIP 56452	5.96	9.7	0.62	82.1	6.7	0.17	2.3	0.13	1.5	0.06	0.8	0.01
HIP 68184	6.49	10.1	0.33	57.1	8.8	0.17	2.0	0.09	0.7	0.04	—	—
HIP 105090	6.69	4.0	0.08	71.1	0.8	0.17	0.4	0.14	0.2	0.04	2.2	0.11
HIP 114046	7.35	3.3	0.04	62.1	0.6	0.17	0.2	0.08	0.4	0.10	2.2	0.15
HIP 1599	4.23	8.6	1.33	134.5	3.1	0.15	1.6	0.16	0.8	0.04	4.5	0.08
HIP 23693	4.71	11.7	1.54	106.6	4.9	0.15	1.2	0.06	0.8	0.05	1.6	0.03
HIP 88601 A	4.03	5.1	0.68	161.7	1.4	0.14	1.5	0.18	0.3	0.04	3.8	0.17
HIP 99240	3.53	6.1	1.33	188.9	1.8	0.14	0.9	0.13	0.6	0.07	3.1	0.17
HIP 104214	5.20	3.5	0.13	104.7	0.5	0.14	1.8	0.29	0.6	0.23	2.9	0.22
HIP 113283	6.48	7.6	0.20	58.6	2.8	0.14	2.7	0.16	1.5	0.09	0.7	0.01
HIP 7981	5.24	7.5	0.47	90.9	4.4	0.13	1.5	0.20	1.8	0.18	4.6	0.11
HIP 15510	4.26	6.0	0.69	137.6	1.3	0.13	1.9	0.22	0.2	0.11	0.8	0.10
HIP 19849	4.43	5.0	0.42	130.8	2.0	0.13	1.8	0.23	1.5	0.15	1.5	0.12
HIP 57443	4.89	9.2	0.86	100.6	3.3	0.13	1.7	0.15	1.1	0.12	2.8	0.06
HIP 96100	4.67	5.8	0.44	114.8	2.0	0.13	1.3	0.19	1.6	0.16	2.0	0.14
HIP 106440	8.67	5.0	0.03	35.2	3.9	0.13	—	—	—	—	—	—
HIP 5336	5.17	7.6	0.46	89.5	3.4	0.12	1.9	0.10	0.5	0.03	—	—
HIP 45343	7.64	6.1	0.08	46.9	3.3	0.12	0.7	0.05	—	—	1.5	0.04
HIP 61317	4.24	8.4	1.27	133.4	2.3	0.12	1.5	0.14	0.7	0.09	4.7	0.12
HIP 120005	7.70	6.1	0.08	46.5	2.4	0.12	1.3	0.04	0.1	0.01	—	—
HIP 3821 A	3.45	5.9	1.31	192.7	1.4	0.11	2.0	0.16	0.1	0.03	2.3	0.12
HIP 7751 B	5.68	7.8	0.34	74.6	2.4	0.11	1.1	0.12	0.2	0.04	0.8	0.01
HIP 10138	6.12	10.8	0.42	59.9	4.6	0.11	1.4	0.08	2.4	0.10	—	—
HIP 15457	4.84	9.1	0.88	102.5	2.8	0.11	2.5	0.15	1.1	0.11	3.3	0.06
HIP 32984	6.58	8.7	0.24	55.8	3.9	0.11	2.7	0.12	2.1	0.08	—	—
HIP 56997	5.31	9.6	0.65	83.8	3.1	0.11	2.0	0.16	1.3	0.10	1.6	0.03
HIP 72848	6.00	11.5	0.54	64.1	4.8	0.11	1.7	0.12	3.3	0.12	—	—
HIP 85295	7.49	7.7	0.12	45.5	3.8	0.11	1.6	0.07	1.3	0.04	—	—
HIP 99461	5.32	6.0	0.28	88.1	1.8	0.11	0.4	0.14	0.2	0.01	3.2	0.11
HIP 99701	7.97	6.2	0.07	41.8	3.5	0.11	—	—	—	—	—	—
HIP 23311	6.23	8.7	0.30	62.5	2.4	0.10	2.3	0.12	1.8	0.12	0.8	0.01
HIP 1475	8.13	3.6	0.02	43.2	1.1	0.09	0.9	0.04	—	—	—	—
HIP 29295	8.15	5.6	0.06	41.7	2.1	0.09	0.6	0.04	—	—	—	—

Table 6 (*Continued*).

Target	Vmag	d (pc)	L (L_{\odot})	EEID	LUVOIR B		HabEx 4H		HabEx 4S		HabEx 4C	
					T_{int}	Char	T_{int}	Char	T_{int}	Char	T_{int}	Char
HIP 40693	5.94	12.5	0.63	63.5	5.1	0.09	3.4	0.14	1.7	0.06	—	—
HIP 67155	8.43	5.4	0.04	36.8	3.1	0.09	—	—	—	—	—	—
HIP 8362	5.63	10.1	0.55	73.5	2.5	0.08	1.7	0.13	0.9	0.10	2.7	0.05
HIP 29271	5.08	10.2	0.91	93.7	3.0	0.08	2.3	0.12	1.8	0.14	—	—
HIP 86400	6.53	11.0	0.34	52.9	3.8	0.08	2.0	0.08	1.0	0.04	—	—
HIP 88972	6.38	11.0	0.36	54.7	3.0	0.08	3.8	0.14	2.2	0.08	—	—
HIP 99825	5.72	8.9	0.42	73.0	2.1	0.08	1.5	0.11	2.3	0.15	—	—
HIP 25878	7.97	5.7	0.07	45.5	2.1	0.07	0.2	0.02	—	—	—	—
HIP 41926	6.37	12.2	0.41	52.8	3.0	0.07	1.6	0.06	2.5	0.08	—	—
HIP 47080	5.39	11.4	0.85	80.9	3.3	0.07	1.9	0.09	3.6	0.15	—	—
HIP 81300	5.76	9.8	0.47	70.6	2.5	0.07	2.2	0.19	2.0	0.15	0.7	0.01
HIP 105858	4.22	9.3	1.54	134.2	1.3	0.07	0.4	0.04	0.5	0.04	1.2	0.03
HIP 439	8.56	4.3	0.02	34.8	1.6	0.06	—	—	—	—	—	—
HIP 15371	5.24	12.0	1.03	84.5	2.8	0.06	1.6	0.07	2.7	0.06	1.6	0.02
HIP 26779	6.20	12.3	0.50	57.7	2.9	0.06	2.5	0.07	—	—	—	—
HIP 42808	6.58	11.1	0.32	50.7	3.1	0.06	1.2	0.05	1.1	0.04	—	—
HIP 80337	5.37	12.8	1.03	79.6	2.7	0.06	1.2	0.05	1.1	0.03	—	—
HIP 84720 A	5.52	8.8	0.52	81.7	2.0	0.06	0.4	0.05	0.8	0.04	2.3	0.03
HIP 86162	9.15	4.5	0.02	33.0	2.4	0.06	—	—	—	—	—	—
HIP 12777	4.10	11.1	2.42	139.8	0.9	0.05	0.3	0.02	0.2	0.01	0.7	0.01
HIP 43587	5.96	12.3	0.66	65.8	2.5	0.05	2.7	0.11	0.4	0.02	—	—
HIP 58576	5.53	12.8	0.93	75.5	3.0	0.05	0.7	0.02	—	—	—	—
HIP 85235	6.44	12.8	0.44	51.6	2.6	0.05	—	—	—	—	—	—
HIP 113576	7.87	8.2	0.10	38.0	3.2	0.05	—	—	—	—	—	—
HIP 14632	4.05	10.5	2.36	145.6	0.6	0.04	0.4	0.07	1.3	0.03	1.9	0.03
HIP 15330	5.53	12.0	0.80	74.5	2.1	0.04	2.0	0.15	2.4	0.06	0.8	0.01
HIP 27072	3.59	8.9	2.51	177.6	1.1	0.04	0.5	0.09	1.0	0.03	1.2	0.04
HIP 51459	4.82	12.8	1.68	101.3	1.7	0.04	1.0	0.05	—	—	—	—
HIP 57939	6.42	9.1	0.23	52.9	1.0	0.04	2.3	0.10	0.4	0.02	—	—
HIP 69972	6.66	11.8	0.39	52.7	2.1	0.04	3.0	0.09	3.7	0.10	—	—
HIP 79672	5.50	13.9	1.11	75.8	1.7	0.04	1.3	0.05	1.5	0.04	—	—
HIP 98767	5.73	15.9	1.20	69.2	2.9	0.04	2.4	0.07	0.9	0.03	—	—
HIP 117712	6.38	10.9	0.39	57.3	2.3	0.04	1.4	0.07	0.1	0.01	—	—

Table 6 (*Continued*).

Target	Vmag	d (pc)	L (L_{\odot})	EEID	LUVOIR B		HabEx 4H		HabEx 4S		HabEx 4C	
					T_{int}	Char	T_{int}	Char	T_{int}	Char	T_{int}	Char
HIP 3583	5.80	15.2	0.98	65.4	2.7	0.03	1.5	0.04	—	—	—	—
HIP 10798	6.33	12.7	0.45	53.2	1.9	0.03	1.2	0.03	1.5	0.05	—	—
HIP 22449	3.17	8.1	3.00	214.6	0.6	0.03	0.1	0.01	0.6	0.04	0.8	0.03
HIP 24813	4.69	12.6	1.89	108.8	1.4	0.03	0.6	0.03	0.2	0.01	—	—
HIP 38908	5.58	16.2	1.36	72.0	1.5	0.03	0.6	0.03	0.4	0.01	—	—
HIP 54211	8.82	4.9	0.02	31.1	1.5	0.03	—	—	—	—	—	—
HIP 58345	6.99	10.2	0.25	48.8	1.6	0.03	—	—	—	—	—	—
HIP 64394	4.24	9.1	1.48	133.1	1.0	0.03	1.3	0.15	0.7	0.06	1.9	0.04
HIP 64797 A	6.56	11.1	0.32	51.3	1.3	0.03	0.7	0.03	—	—	—	—
HIP 75181	5.65	14.8	1.09	70.4	2.0	0.03	1.0	0.04	0.6	0.02	—	—
HIP 80686	4.90	12.1	1.41	98.1	1.4	0.03	1.0	0.06	1.3	0.05	—	—
HIP 84862	5.38	14.3	1.30	79.6	1.6	0.03	1.4	0.05	—	—	—	—
HIP 42438	5.63	14.4	1.03	70.6	1.0	0.02	2.1	0.08	0.9	0.01	—	—
HIP 53721	5.03	14.1	1.71	93.1	1.2	0.02	0.6	0.04	0.9	0.03	—	—
HIP 58910	7.78	5.5	0.03	30.1	1.2	0.02	—	—	—	—	—	—
HIP 77257	4.41	12.1	2.22	122.9	1.0	0.02	0.7	0.05	1.2	0.04	0.4	0.01
HIP 95447	5.16	15.2	1.84	89.3	0.9	0.02	1.0	0.04	0.4	0.01	—	—
HIP 107649	5.57	16.0	1.35	72.6	1.4	0.02	0.3	0.02	0.4	0.02	—	—
HIP 113296	8.64	6.8	0.06	34.4	1.2	0.02	—	—	—	—	—	—
HIP 3909	5.16	15.8	1.85	86.3	0.8	0.01	0.4	0.02	0.3	0.01	—	—
HIP 7978	5.52	17.4	1.64	73.5	0.7	0.01	0.1	0.01	0.7	0.01	—	—
HIP 27435	5.97	15.2	0.85	60.6	0.7	0.01	1.6	0.04	0.4	0.02	—	—
HIP 34065	5.56	16.5	1.45	72.9	0.8	0.01	—	—	—	—	—	—
HIP 47592	4.92	15.0	2.11	96.7	0.4	0.01	—	—	—	—	—	—
HIP 49081	5.37	15.1	1.47	80.6	0.7	0.01	1.2	0.04	0.2	0.01	—	—
HIP 78072	3.85	11.3	3.13	157.3	0.1	0.01	—	—	1.3	0.04	—	—
HIP 85523	9.41	4.5	0.02	28.7	0.5	0.01	—	—	—	—	—	—
HIP 86796	5.12	15.5	2.01	91.4	0.6	0.01	0.8	0.02	0.6	0.03	—	—
HIP 97944	6.21	14.1	0.80	63.5	0.5	0.01	0.3	0.01	—	—	—	—
HIP 110109	5.36	13.8	1.23	80.5	0.8	0.01	—	—	—	—	—	—
HIP 113357	5.45	15.6	1.45	77.1	0.8	0.01	0.5	0.03	—	—	—	—
HIP 544	6.06	13.7	0.68	60.4	—	—	0.3	0.01	1.0	0.03	—	—
HIP 2021	2.82	7.5	3.70	258.	—	—	—	—	0.2	0.01	—	—

Table 6 (*Continued*).

Target	Vmag	d (pc)	L (L_{\odot})	EEID	LUVOIR B		HabEx 4H		HabEx 4S		HabEx 4C	
					T_{int}	Char	T_{int}	Char	T_{int}	Char	T_{int}	Char
HIP 5862	4.96	15.1	2.07	95.1	—	—	1.2	0.04	0.3	0.01	—	—
HIP 7513	4.09	13.5	3.64	141.3	—	—	—	—	0.7	0.03	—	—
HIP 10644	4.86	10.8	1.21	102.1	—	—	1.2	0.08	1.7	0.07	1.7	0.02
HIP 12653	5.40	17.2	1.78	77.8	—	—	—	—	0.4	0.01	—	—
HIP 12843	4.47	14.2	2.83	118.2	—	—	—	—	1.0	0.02	—	—
HIP 16852	4.29	14.0	3.29	130.0	—	—	—	—	0.6	0.02	—	—
HIP 17378	3.52	9.0	3.37	203.2	—	—	—	—	0.3	0.02	—	—
HIP 22263	5.49	13.3	1.01	75.8	—	—	3.6	0.14	2.3	0.08	—	—
HIP 26394	5.65	18.3	1.64	69.9	—	—	0.5	0.02	—	—	—	—
HIP 32480	5.25	16.7	1.95	83.5	—	—	0.3	0.01	—	—	—	—
HIP 33277	5.75	17.2	1.32	66.6	—	—	1.4	0.05	—	—	—	—
HIP 35136	5.54	16.9	1.54	73.5	—	—	—	—	0.4	0.01	—	—
HIP 37279	0.40	3.5	7.12	760.8	—	—	0.6	0.13	—	—	—	—
HIP 43726	5.99	17.4	1.09	60.1	—	—	0.4	0.01	—	—	—	—
HIP 45333	5.19	19.6	2.87	86.6	—	—	—	—	0.6	0.01	—	—
HIP 48113	5.07	18.4	2.78	90.7	—	—	—	—	0.2	0.01	—	—
HIP 48331	7.65	11.2	0.17	37.3	—	—	—	—	—	—	1.2	0.01
HIP 50954	3.98	16.2	5.62	146.2	—	—	—	—	1.0	0.01	—	—
HIP 57757	3.59	10.9	3.79	178.2	—	—	—	—	1.0	0.03	—	—
HIP 62207	5.95	17.4	1.12	60.9	—	—	1.1	0.03	1.3	0.03	—	—
HIP 70497	4.04	14.5	4.42	144.7	—	—	—	—	0.3	0.01	—	—
HIP 71681	1.35	1.3	0.52	541.3	—	—	0.5	0.18	0.1	0.06	1.2	0.24
HIP 71683	-0.01	1.3	1.61	948.4	—	—	1.6	0.19	0.1	0.02	—	—
HIP 77358	6.01	15.4	0.87	60.8	—	—	1.2	0.03	—	—	—	—
HIP 77760	4.60	15.9	3.26	113.6	—	—	—	—	0.9	0.02	—	—
HIP 79248	6.61	17.8	0.70	47.6	—	—	—	—	0.5	0.01	—	—
HIP 83389	6.76	18.6	0.65	43.5	—	—	—	—	—	—	0.4	0.01
HIP 89042	5.47	17.6	1.78	75.7	—	—	2.1	0.05	1.5	0.03	—	—
HIP 96441	4.49	18.3	4.51	115.8	—	—	—	—	1.0	0.02	—	—
HIP 98959	6.07	17.7	1.06	58.0	—	—	0.2	0.01	—	—	—	—
HIP 101997	6.36	14.4	0.56	52.0	—	—	0.3	0.01	—	—	—	—
HIP 109176	3.77	11.7	3.63	162.5	—	—	—	—	1.0	0.01	—	—
HIP 114924	5.58	20.5	2.15	71.5	—	—	—	—	0.5	0.01	—	—

Acknowledgments

We deeply appreciate those that contributed to the ExSDET effort: Shannon Dulz (Notre Dame), Andrew Horning (Cornell University), Garreth Ruane (JPL), Dan Sirbu (Ames Research Center), Gabriel Soto (Cornell University), and Margaret Turnbull (SETI Institute). The ExSDET members are Rhonda Morgan (lead), Dmitry Savransky, Chris Stark (STScI), Avi Mandell (NASA Goddard), Ruslan Belikov (NASA Ames), John Krist (JPL), and Eric Nielsen (Stanford University). We appreciate the EPRV Working Group for their feedback on early results. Monte Carlo simulations in this study were run with Gnu parallel.⁴⁹ Graphics in Figs. 14 and 15 were produced in part with third-party code.^{50,51} This information is predecisional and is provided for planning and discussion purposes only. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). ©2021. All rights reserved.

References

1. R. Morgan et al., “Standard exoplanet yield evaluation for the LUVOIR and HabEx concept studies,” *Proc. SPIE* **11117**, 1111701 (2019).
2. R. Morgan et al., “The standard definitions and evaluation team final report: a common comparison of exoplanet yield,” 2019, https://exoplanets.nasa.gov/system/internal_resources/details/original/1434_Standards_Team_Final_Report_20191007.pdf.
3. P. Plavchan et al., “EarthFinder probe mission concept study: characterizing nearby stellar exoplanet systems with earth-mass analogs for future direct imaging,” arXiv:2006.13428v1 (2020).
4. S. Gaudi et al., “Extreme precision radial velocity initiative plan,” 2020, https://exoplanets.nasa.gov/internal_resources/1556/.
5. J. Crass, S. Gaudi, and S. Leifer, “Extreme precision radial velocity working group final report,” 2021, <https://exoplanets.nasa.gov/exep/NNExplore/EPRV/>.
6. D. Savransky and N. J. Kasdin, “Design reference mission construction for planet finders,” *Proc. SPIE* **7010**, 70101T (2008).
7. D. Savransky and N. J. Kasdin, “Dynamic filtering for the analysis of astrometric and radial velocity data sets for the detection of exoplanets,” in *AIAA Guidance, Navigation, and Control*, 6083 (2009).
8. D. Savransky and N. Kasdin, “Automated design reference mission generation for theia,” *Bull. Am. Astron. Soc.* **41**, 363 (2009).
9. D. Savransky, “Designing the next generation of extra-solar planet observatories,” in *IPMJ Astrostats Meeting* (2009).
10. D. Savransky, N. Kasdin, and D. Spergel, “Results from the automated design reference mission constructor for exoplanet imagers,” *Proc. SPIE* **7440**, 744009 (2009).
11. D. Savransky, N. J. Kasdin, and E. Cady, “Analyzing the designs of planet finding missions,” *Publ. Astron. Soc. Pac.* **122**, 401–419 (2010).
12. D. Savransky et al., “Occulting ozone observatory science overview,” *Proc. SPIE* **7731**, 77312H (2010).
13. D. Savransky, C. Delacroix, and D. Garrett, “EXOSIMS: Exoplanet Open-Source Imaging Mission Simulator,” Astrophysics Source Code Library (2017).
14. D. Savransky and D. Garrett, “WFIRST-AFTA coronagraph science yield modeling with EXOSIMS,” *J. Astron. Telesc. Instrum. Syst.* **2**(1), 011006 (2015).
15. C. Delacroix et al., “Science yield modeling with the Exoplanet Open-Source Imaging Mission Simulator (EXOSIMS),” *Proc. SPIE* **9911**, 991119 (2016).
16. D. Savransky and EXOSIMS team, “Exoplanet Open-Source Imaging Mission Simulator,” 2020, <https://github.com/dsavransky/EXOSIMS>.
17. C. C. Stark et al., “Maximizing the exoearth candidate yield from a future direct imaging mission,” *Astrophys. J.* **795**(2), 122 (2014).
18. C. C. Stark et al., “Lower limits on aperture size for an exoearth detecting coronagraphic mission,” *Astrophys. J.* **808**(2), 149 (2015).

19. C. C. Stark et al., “Maximized exoearth candidate yields for starshades,” *J. Astron. Telesc. Instrum. Syst.* **2**(4), 041204 (2016).
20. C. C. Stark et al., “A direct comparison of exoearth yields for starshades and coronagraphs,” *Proc. SPIE* **9904**, 99041U (2016).
21. R. A. Brown, “Obscurational completeness,” *Astrophys. J.* **607**, 1003–1013 (2004).
22. R. A. Brown, “Single-visit photometric and obscurational completeness,” *Astrophys. J.* **624**, 1010–1024 (2005).
23. R. Brown and R. Soummer, “New completeness methods for estimating exoplanet discoveries by direct detection,” *Astrophys. J.* **715**, 122 (2010).
24. B. S. Gaudi et al., “The Habitable Exoplanet Observatory (HabEx) mission concept study final report,” arXiv:2001.06683 (2020).
25. S. Hunyadi, S. Shaklan, and R. Brown, “Single visit completeness optimization,” tech. rep., Terrestrial Planet Finder – Coronagraph STDT #5, 2005, <http://hdl.handle.net/2014/38535>.
26. S. L. Hunyadi, A. S. Lo, and S. B. Shaklan, “The dark side of TPF: detecting and characterizing extra-solar earthlike planets with one or two external occulters,” *Proc. SPIE* **6693**, 669303 (2007).
27. S. Hunyadi, S. Shaklan, and R. Brown, “The lighter side of TPF-C: evaluating the scientific gain from a smaller mission concept,” *Proc. SPIE* **6693**, 66930Q (2007).
28. R. Trabert et al., “Design reference missions for the exoplanet starshade (exo-s) probe-class study,” *Proc. SPIE* **9605**, 96050Y (2015).
29. S. Seager et al., “Exo-s: starshade probe-class exoplanet direct imaging mission concept final report,” Tech. Rep., available at exep.jpl.nasa.gov/stdt (2015).
30. R. Morgan et al., “ExEP yield modeling tool and validation test results,” *Proc. SPIE* **10400**, 104001K (2017).
31. G. J. Soto et al., “Optimal starshade observation scheduling,” *Proc. SPIE* **10698**, 106984M (2018).
32. G. Soto et al., “Parameterizing the search space of Starshade fuel costs for optimal observation schedules,” *J. Guid. Control Dyn.* **42**(12), 2671–2676 (2019).
33. G. Soto, D. Savransky, and R. Morgan, “Analytical model for starshade formation flying with applications to exoplanet direct imaging observation scheduling,” *J. Astron. Telesc. Instrum. Syst.* **7**(2), 021209 (2021).
34. D. R. Keithly et al., “Optimal scheduling of exoplanet direct imaging single-visit observations of a blind search survey,” *J. Astron. Telesc. Instrum. Syst.* **6**(2), 027001 (2020).
35. C. Spohn and D. Savransky, “How orbital fit uncertainties impact dynamic scheduling,” *Proc. SPIE* **11449**, 114492K (2020).
36. M. C. Turnbull, “Exocat-1: the nearby stellar systems catalog for exoplanet imaging missions,” arXiv:1510.01731 (2015).
37. A. Horning, R. Morgan, and E. Nielsen, “Minimum number of observations for exoplanet orbit determination,” *Proc. SPIE* **11117**, 111171C (2019).
38. S. Dulz et al., “Joint radial velocity and direct imaging planet yield calculations: I. Self-consistent planet populations,” *Astrophys. J.* **893**(2), 122 (2020).
39. R. K. Kopparapu et al., “Habitable zones around main-sequence stars: new estimates,” *Astrophys. J.* **765**, 131 (2013).
40. D. Sirbu et al., “Demonstration of multi-star wavefront control for WFIRST, Habex, and LUVOIR,” *Proc. SPIE* **11117**, 1111719 (2019).
41. C. Leinert et al., “The 1997 reference of diffuse night sky brightness,” *Astron. Astrophys. Suppl. Ser.* **127**(1), 1–99 (1998).
42. Ertal2018Report, “Report on hosts survey return,” 2018, http://nexsci.caltech.edu/missions/LBTI/report_hosts_dec2018_v1.2.pdf.
43. C. C. Stark et al., “Exoearth yield landscape for future direct imaging space telescopes,” *J. Astron. Telesc. Instrum. Syst.* **5**(2), 024009 (2019).
44. NASA Exoplanet Archive, “Mission Stars,” <https://exoplanetarchive.ipac.caltech.edu/docs/data.html> (2018).
45. E. Mamajek, “A modern mean dwarf stellar color and effective temperature sequence,” 2019, http://www.pas.rochester.edu/emamajek/EEM_dwarf_UBVIJHK_colors_Teff.txt.

46. T. D. Robinson, K. R. Stapelfeldt, and M. S. Marley, “Characterizing Rocky and Gaseous Exoplanets with 2 m class space-based coronagraphs,” *Pub. Astr. Soc. Pacific* **128**, 025003 (2016).
47. T. D. Robinson, “Hci-noise,” 2016, https://github.com/hablabx/hci_noise.
48. The LUVOIR Team, “The LUVOIR mission concept study final report,” arXiv:1912.06219 (2019).
49. O. Tange, “Gnu parallel 20200522 (‘kraftwerk’),” 2020, <https://doi.org/10.5281/zenodo.3841377>.
50. Charles, “cbrewer: colorbrewer schemes for matlab,” 2020, <https://www.mathworks.com/matlabcentral/fileexchange/34087-cbrewer-colorbrewer-schemes-for-matlab>.
51. K. Kearney, “boundedline.m,” 2020, <https://github.com/kakearney/boundedline-pkg>.

Rhonda Morgan is a senior optical engineer at the Jet Propulsion Laboratory, California Institute of Technology. She received her BS degree in electrical engineering from Caltech and her MS and PhD degrees in optical sciences from the University of Arizona. Her interests include exoplanets, high contrast imaging, and advanced technologies for space telescopes. She is the HabEx technologist and lead of the Exoplanet Program Office yield Standard Definitions and Evaluation Team.

Dmitry Savransky is an assistant professor in the Sibley School of Mechanical and Aerospace Engineering at Cornell University and PI of the Space Imaging and Optical Systems Laboratory. He received his PhD from Princeton University. As a postdoctoral researcher at Lawrence Livermore National Laboratory, he worked on integration, test, and commissioning of the Gemini Planet Imager. His research focuses on ground and space-based exoplanet imaging survey optimization, autonomous optical system control, and advanced image processing.

Michael Turmon is a principal engineer at the Jet Propulsion Laboratory, California Institute of Technology. His research interests are in uncertainty quantification for science measurements and model-driven statistical inversion for science data.

Bertrand Mennesson is a principal scientist at the Caltech Jet Propulsion Laboratory. His expertise is in the design, assembly and scientific exploitation of infrared interferometers and optical coronagraphs used for high contrast, high resolution astronomical imaging. His main scientific interests lie in the direct detection and characterization of extrasolar planets and debris disks.

Eric E. Mamajek is a deputy program chief scientist of the NASA Exoplanet Exploration Program, officed at the Jet Propulsion Laboratory. His research interest is the formation and evolution of exoplanetary systems and their host stars.

Peter Plavchan is an associate professor of physics and astronomy at George Mason University and studies the frequency and formation of planets around cool dwarf stars. He received his PhD from UCLA in 2006 and undergraduate degree from Caltech in 2001. He is the PI of the NASA Probe-class mission concept, EarthFinder, and discovered the young transiting exoplanetary system orbiting AU Mic, the second closest premain sequence star to the Sun.

Tyler D. Robinson is an assistant professor of astronomy and planetary science at Northern Arizona University. Prior to his current position, he was a NASA Carl Sagan fellow at the University of California, Santa Cruz, and a NASA postdoctoral program fellow at NASA Ames Research Center. He uses sophisticated radiative transfer and climate tools to study the atmospheres of Solar System worlds, exoplanets, and brown dwarfs. he also develops instrument models for exoplanet direct imaging.

Biographies of the other authors are not available.