

Coronagraph Design Survey for Future Exoplanet Direct Imaging Space Missions: Final Report

Ruslan Belikov^a, Christopher Stark^b, Nick Siegler^c, Emiel Por^d, Bertrand Mennesson^c, Susan Redmond^f, Pin Chen^c, Kevin Fogarty^a, Olivier Guyon^e, Roser Juanola-Parramon^b, Jeremy Kasdin^g, John Krist^c, Dimitri Mawet^{f,c}, Rhonda Morgan^c, Camilo Mejia Prada^c, Laurent Pueyo^d, Garreth Ruane^c, Dan Sirbu^a, Karl Stapelfeldt^c, Armen Tokadjian^c, John Trauger^c, Neil Zimmerman^b, Mary Angelie Alagao^h, Alex Carlottiⁱ, Jamal Chafi^j, David Doelman^k, Jessica Gersh-Range^g, Lorenzo König^e, Lucille Leboulleuxⁱ, Dwight Moody^c, AJ Riggs^c, Eugene Serabyn^c, Frans Snik^k, and Kent Wallace^c

^aNASA Ames Research Center, Moffett Field, CA 94035, USA

^bNASA Goddard Space Flight Center, Greenbelt, MD, USA 20771

^cJet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Dr, Pasadena, CA 91109, USA

^dSpace Telescope Science institute, Baltimore, MD, USA 21218

^eUniversity of Arizona, Tucson, AZ, USA 85721

^fCalifornia Institute of Technology, Pasadena, CA, USA 91125

^gPrinceton University, Princeton, NJ, USA 08544

^hNational Astronomical Research institute of Thailand, Chiang Mai, Thailand

ⁱUniversité Grenoble Alpes, Grenoble, France

^jOukaimeden Observatory, LPHEA, Cadi Ayyad University, Marrakesh, Morocco

^kLeiden University, Leiden, The Netherlands

1. EXECUTIVE SUMMARY

NASA is embarking on an ambitious program to develop the Habitable Worlds Observatory (HWO) flagship to perform transformational astrophysics, as well as directly image ~ 25 potentially Earth-like planets and spectroscopically characterize them for signs of life. This mission was recommended by the Astro2020 decadal survey, which additionally recommended a new approach for flagship formulation based on increasing the scope and depth of early, pre-phase A trades and technology maturation. A critical capability of the HWO mission is the suppression of starlight. To inform future architecture trades, it is necessary to survey a wide range of candidate technologies, from the relatively mature ones such as the ones described in the LUVOIR and HabEx reports to the relatively new and emerging ones, which may lead to breakthrough performance. In this paper, we present a summary of an effort, funded by NASA's Exoplanet Exoploration Program (ExEP), to survey potential coronagraph options for HWO. In particular, our results consist of: (1) a database of different coronagraph designs sourced from the world-wide coronagraph community that are potentially compatible with HWO; (2) evaluation criteria, such as expected mission yields and feasibility of maturing to TRL 5 before phase A; (3) a unified modeling pipeline that processes the designs from (1) and outputs values for any machine-calculable criteria from (2); (4) assessments of maturity of designs, and other criteria that are not machine-calculable; (5) a table presenting an executive summary of designs and our results. While not charged to down-select or prioritize the different coronagraph designs, the products of this survey were designed to facilitate future HWO trade studies.

A table showing the currently identified coronagraph designs and metrics is shown in Figure 1. The columns in the matrix correspond to the 19 coronagraph designs collected from the community so far, which are organized into categories depending on the basic type of architecture. The rows correspond to metrics and are

Send correspondence to Ruslan Belikov

E-mail: ruslan.belikov@nasa.gov, Telephone: 1 650 604 0833

by the interface is now automatically generated by our pipeline. The CDS pipeline also standardized metrics for comparison, e.g., by calculating contrast consistently for all designs. This standardization and automation sped up the CDS analyses, providing rapid turnaround of calculations while enabling apples-to-apples comparisons. HWO could benefit from further pipeline development with a focus on establishing standardized interfaces between key codes and automated execution.

The CDS pipeline also automated calculations of sensitivity to wavefront aberrations. CDS finds that there are designs that, in theory, can provide adequate robustness to several of expected HWO-like wavefront aberrations.

The CDS exoplanet science yield calculations showed that there are a number of coronagraph designs that are, in theory, capable of achieving or exceeding the Astro2020 Decadal Survey’s recommended science goals with a ~ 6 m inscribed diameter telescope, notably both for on- and off-axis options. Additional designs are within “spitting distance” of these goals. We expect future study of these designs to result in improvements in science yield and find that premature down-selection based on these results could negatively impact the design trade space of HWO.

The CDS exoplanet science yield efforts also illustrated the complex trade space of coronagraphy. No single instrument parameter fully captures instrument performance. Further, optimized target lists vary dramatically among coronagraph mask designs and evaluating coronagraph performance via fiducial stars can lead to biased interpretations; target lists need to be jointly optimized with science goals and instrument capabilities. Tracing simulated instrument response to science yield helps us understand the available trade space.

Yield calculations also teach us how to efficiently operate and design a coronagraphic instrument. For example, yield calculations often result in exoplanet detections occurring inside of the nominal IWA, suggesting that this parameter by itself is not a good description of science yield. Combinations of coronagraph masks often result in higher yields than any one mask alone, as some masks are, e.g., better suited for spectral characterization while others are suited for broadband detection. HWO may benefit substantially from carrying forward a number of coronagraph design technologies and approaching instrument design as a “suite” of coronagraphic options.

Of primary importance is the demonstration of high contrast in the laboratory. However, gains in science returns become increasingly marginal as contrast approaches 10^{-10} level, and so demonstrations of improvements in other “knobs” available to coronagraphs (such as throughput, IWA, bandwidth and post-processing) will likely become increasingly important. In fact, it is possible that contrasts already demonstrated in the laboratory[?] will be sufficient to meet HWO science goals, as long as these demonstrations can be reproduced with a higher-performance coronagraph design that simultaneously demonstrates the required bandwidth, off-axis core throughput, and sensitivity to WF aberrations.

Our report represents a snapshot of the ever-evolving state of coronagraph technology, rather than a final assessment. We did not intend or attempt any kind of down-select or prioritization of different coronagraph designs. A meaningful downselect requires, at the very least, iterating and fully optimizing multiple promising architectures, as well as joint trade studies between the coronagraph, telescope, and other aspects of HWO. We hope that our results and tools will be useful in facilitating such optimizations and trades. In particular, our coronagraph database and pipeline are designed to be regularly updated with new coronagraph designs and new metrics, to make it easier for coronagraph designers to rapidly evaluate and optimize their future designs according to the most important metrics determined by future HWO design trades. Conversely, our pipeline is designed to enable HWO teams to rapidly assess the impact of latest advances in coronagraph design and technology development, and maximize the impact of investments in these areas.

2. INTRODUCTION AND MOTIVATION

NASA has recently embarked on an ambitious program to develop the flagship mission recommended by the Astro2020 decadal survey,[?] now known as the Habitable Worlds Observatory (HWO). (See Figure 4.) The driving science goal of HWO, as recommended by Astro2020, is to directly image ~ 25 potentially Earth-like planets and spectroscopically characterize them for signs of life.[?] In addition, HWO will detect and spectroscopically characterize a plethora of other planet types and disks around stars, and perform transformative science in many other areas of astrophysics. HWO is expected to consist of a ~ 6 m inscribed diameter ultra-stable space telescope with a coronagraph capable of reaching $\sim 10^{-10}$ contrast.[?] It builds on numerous previous mission

concept studies, most notably HabEx[?] and LUVOIR,[?] and will build on the upcoming Coronagraph Instrument on the Roman Space Telescope, a flagship mission scheduled to launch by 2027 that will demonstrate many coronagraph technologies potentially valuable for HWO.

In addition, Astro2020 recommended a new approach for flagship formulation, which involves increasing the scope and depth of early, pre-phase-A trades and technology maturation.[?]

A critical capability requiring thorough trades and technology maturation is starlight suppression. To inform future architecture trades, it is necessary to survey a wide range of technologies that could be viable for HWO, from the relatively mature ones, such as the ones on the Roman Coronagraphic Instrument, to the relatively new and emerging ones, which may lead to breakthrough performance.

It is natural to ask why such a survey is needed, given the scope of the LUVOIR and HabEx reports, which already described feasible designs on which the Astro2020 recommendation was based. There are several reasons for conducting a new survey. First, in the 5 years since the HabEx and LUVOIR reports, coronagraph technology has improved, both in terms of designs and lab demonstrations. For example, there are coronagraph designs today which increase the expected exo-Earth yield by 60% relative to the designs in the LUVOIR and HabEx concepts. In terms of lab demonstrations, although the in-lab contrast record of $4e-10$ in a 10% band has stood since 2019,[?] we now have a demonstration of this contrast in a 20% band along with a greater understanding of the limiting factors,[?] as well as demonstrations of more efficient coronagraphs at $1e-9$ and $1e-8$ contrast, and these are expected to improve. A key goal of our survey is to provide an update that captures this new landscape.

Second, Astro2020 recommends thorough, early, well-funded trade studies and technology development in order to lower the risks and costs of HWO (a similar theme also runs through the Large Mission Study Report, 2019). Coronagraph designs span a very rich trade space (Figure 2) and represent several potentially powerful levers to lower the risks and costs of HWO. For example, the greatest risk and cost drivers on HWO may be in demonstrating $1e-10$ contrast and 10pm telescope stability, so any opportunity to reduce these requirements should be investigated. Improvements in coronagraph designs provide just this kind of opportunity. For example, HWO contrast requirements may potentially be relaxed by improving other coronagraph parameters, such as throughput and inner working angle (IWA), while still retaining the same yield of exo-Earths.[?] Similarly, telescope stability requirements can be relaxed by making the coronagraphic instrument more robust to wavefront aberrations.[?] In the past 4 years, new coronagraph designs have appeared that achieve better throughput, IWA, and robustness to low-order aberrations. In addition, new wavefront and post-processing techniques promise to achieve better robustness to all aberrations. If these technologies can be developed at lower cost and risk than developing an ultra-stable space telescope and/or demonstrating $1e-10$ raw contrast (as opposed to, say $3e-10$), then overall HWO mission cost and risk could be reduced. In addition, the performance of coronagraphs for on-axis apertures has been advancing, both in terms of expected science yield and in terms of laboratory performance, and it is possible that in the future, the performance of coronagraphs on obstructed apertures will be comparable to unobstructed apertures.[?]

A third reason is that improving the performance of an instrument is a very powerful and flexible lever to reduce risk and cost. For example, if one wishes to increase the photon flux at the science detector by 10%, one can either improve the instrument efficiency by 10%, or the telescope collecting area by 10%. On a flagship, it is (probably) easier and less costly to do the former. In general, the larger the mission, the more leverage one gets from investments in technology improvements. We hope that CDS will expose opportunities for such improvements, in ways that reduce mission risk, cost, and/or improve science.

A fourth motivating reason for our survey is to help future mission designers understand the impact and sensitivities of engineering trades to science yield and margin. There are several considerations that make this especially important for HWO. For example, the expected science yields of early mission concepts are often optimistic, and erode as mission concepts mature. HWO in particular may be vulnerable to this “yield erosion” because of the uncertainty on η_{Earth} . Another consideration is that the science of HWO needs to remain competitive relative to the expected science landscape of the 2040s, with Extremely Large Telescopes (ELTs) likely delivering some forms of exo-Earth imaging from the ground. An expected benefit of our study is that it will identify opportunities that would enhance yields, provide science margin, guard against yield erosion, and ensure relevance in the science landscape of the 2040s. Of course, opportunities that increase science and

lower science risk often (but not always!) come at the expense of increased mission cost and technology risk. The existence of such opportunities and options would allow future trade studies to better optimize HWO and balance different risks against each other, likely resulting in a better mission overall.

Early, well-funded coronagraphic trade studies may lower the risks and costs of HWO.[?] Multiple lever arms, or "knobs" have been identified that could improve coronagraph performance in novel ways. For example, HWO contrast requirements may potentially be relaxed by improving other coronagraph parameters, such as throughput and inner working angle (IWA), while still retaining the same yield of potentially Earth-like planets.[?] Similarly, telescope stability requirements can be relaxed by making the coronagraphic instrument more robust to wavefront aberrations.[?] In addition, new wavefront control and post-processing techniques promise to achieve better robustness to all aberrations. It is also important to clarify the effects of different coronagraph parameters on science yield and margin, and thus enable mission and coronagraph designers to design a better mission.

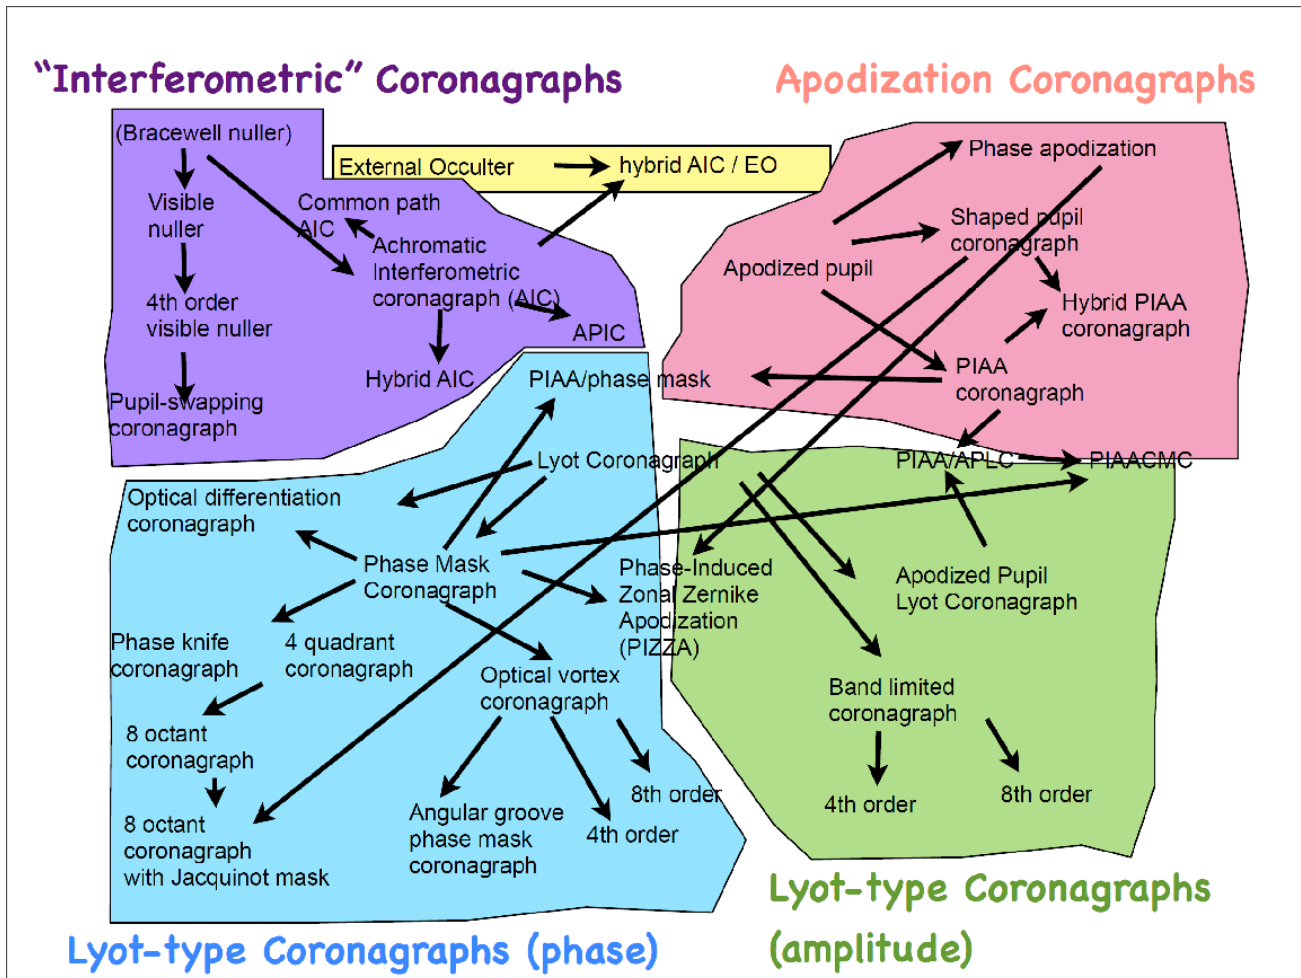


Figure 2: Coronagraph designs have a rich and varied trade space. (Image courtesy of Olivier Guyon.)

Examples what we mean by "knobs" are shown as a simplified cartoon representation in Figure 3, which also shows why coronagraphic knobs are especially important for HWO. In particular, we show examples of 8 mission parameters. Some of them are mostly independent of the coronagraph design (telescope diameter, mirror coatings, detector quantum efficiency), while some of them are directly dependent on the coronagraph design (throughput, inner working angle (IWA), robustness to telescope aberrations, spectral bandwidth). A lot of the non-coronagraphic knobs are already close to saturation on HWO, indicated by the value of "10" on the knob scale (where "turning to 11" represents herculean efforts to go past what is currently considered feasible). For example, given NASA's budget, the diameter that is currently baselined for HWO primary mirror (6.5m) is

close to what is probably the maximum practically possible, and going beyond it may result in a mission that is too expensive. Coating reflectivity, and detector QE are also somewhat close to their maximum value of 100% (and “going to 11” is impossible since it would violate energy conservation). By definition, improvements in parameters that are close to saturation can only be marginal, though occasionally even marginal improvements are valuable. On the other hand, many of the coronagraph parameters (such as coronagraph throughput, IWA, robustness to aberrations, and bandwidth) are still far from fundamental physics limits,⁷ and therefore have a lot of potential to improve – at least in theory. As will be shown later, potential benefits of optimizing coronagraph parameters include: significant improvements in science yield; reducing mission risk and cost by relaxation of key requirements such as telescope stability, contrast, etc.; and enabling a potentially simpler and more cost-effective on-axis aperture.

Of course, just because a parameter has a lot of room to improve does not necessarily mean that these improvements are easy or that the benefit/cost ratio is high. Our survey did not directly evaluate benefit/cost ratios of turning various knobs, but it did capture the benefits as well as the maturity and challenges of our surveyed technologies. This data can be used to facilitate more detailed future studies for HWO, including benefit/cost and benefit/risk analyses of different options.

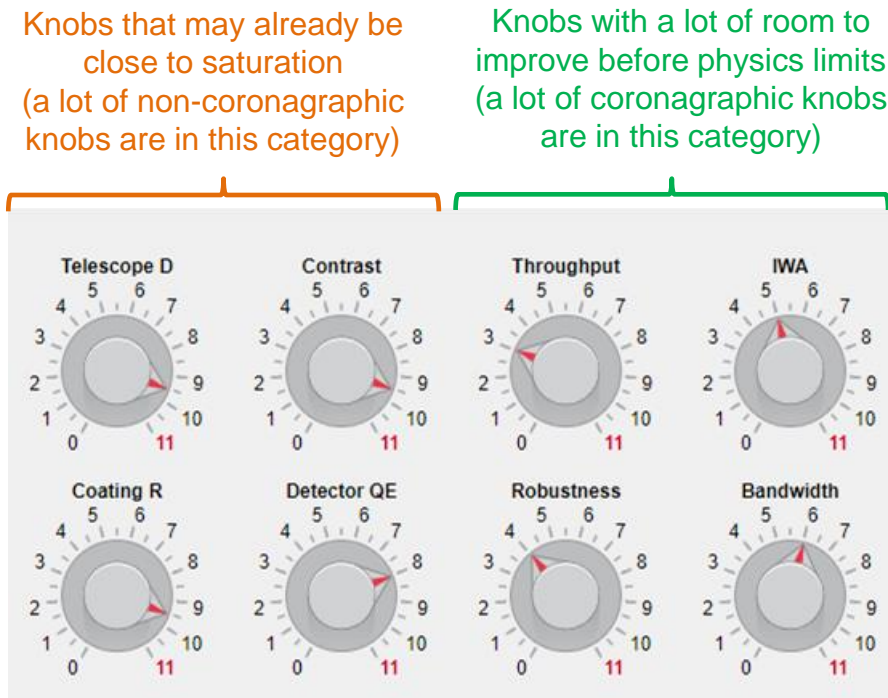


Figure 3: Examples of “knobs” that can be varied in a mission’s trade space. Many of the coronagraphic knobs (right half) currently have a lot of potential to improve before they reach a saturation point. Note: the knobs on this figure are meant as a high-level cartoon for illustration purposes only. Values of the knobs are unitless and are normalized to 10 as the maximum theoretical or practically achievable value for that parameter.

To update and capture the current state of a broad range of coronagraph designs and their potential performance, we conducted the Coronagraph Design Survey (CDS), initiated by the Exoplanet Exploration Program Office (ExEP). Here we report the findings of our survey. We detail the potential science benefits, performance robustness, and technological readiness for a wide variety of community-contributed coronagraph technologies, from the relatively mature ones such as the ones in the Roman Coronagraph Instrument, to the relatively new and emerging ones, which may lead to breakthrough performance. We hope that this survey will help mission designers understand the range of possible coronagraph designs, the sensitivities of these coronagraphs to different engineering trades, the effects of design choices on science yield, and help budget for exoplanet science

margin. We also hope that this design survey inspires additional opportunities for instrument improvement in ways that reduce mission risk, cost, and/or improve science.

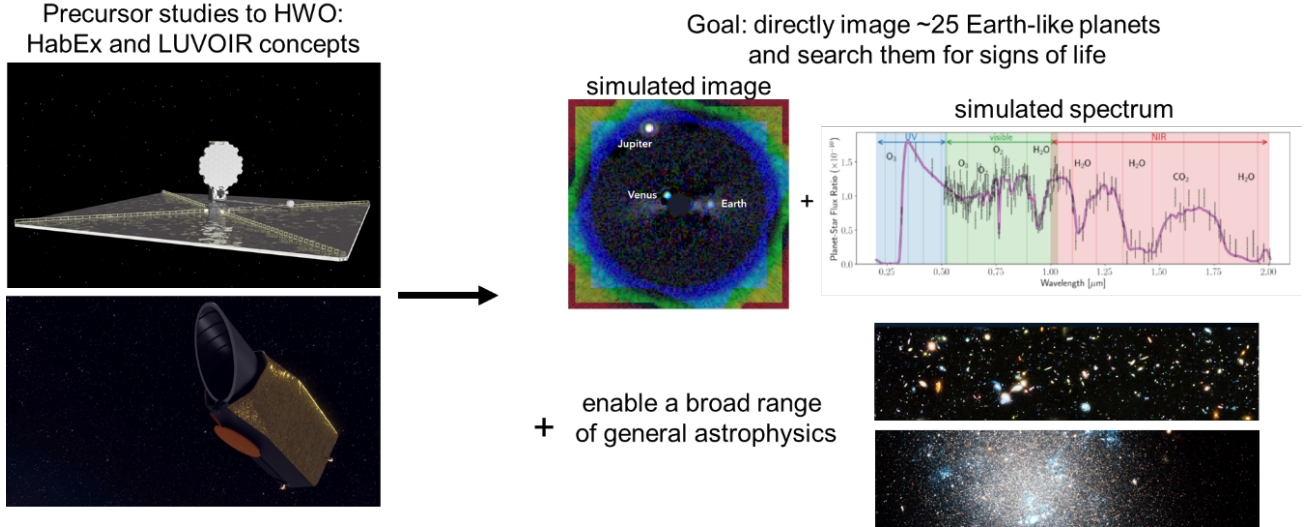


Figure 4: Left: the Habitable Worlds Observatory will build on the LUVOIR and HabEx mission concepts as well as the upcoming Coronagraphic Instrument on the Roman Space Telescope. Right: the driving science goal of HWO is to directly image ~ 25 Earth-like planets and search them for spectroscopic biomarkers.

3. SURVEY SCOPE AND STRUCTURE

The Coronagraph Design Survey was an effort that ran roughly from January 2023 through June 2024. It consisted of 21 members from a diversity of institutions, backgrounds, and career levels, and led by 2 co-chairs (Belikov and Stark). ExEP chief technologist (Siegler) served as an advisor. In addition, many interested coronagraph researchers worldwide participated in our survey by contributing their designs (see Table 1). In addition, CDS coordinated its activities with other parallel studies related to HWO, such as the Coronagraph Technology Roadmap Team (CTR), Coronagraph Deformable Mirror Roadmap (DMTR), the Ultra-Stable Observatory Roadmap Team (USORT), and the HWO START and TAG teams.

3.1 Goals and products

The top-level goals and accomplishments of the CDS were as follows:

1. Survey and document viable coronagraph designs across the world that can inform the Habitable Worlds Observatory about their capabilities and technology readiness.
2. Facilitate future evaluation and comparison of the coronagraph designs to advance, based on a set of technical and programmatic assessment criteria.
3. Identify novel coronagraph technologies that could mature rapidly for which NASA’s technology development investments could be efficiently leveraged.

Note that the CDS was a fact-finding effort and did not attempt any technology down-selects. Rather, the focus of CDS was to gather information and coronagraph design options, provide a unified set of criteria, and assess coronagraph options along those criteria.

Key products of the survey are the following:

- a. An information matrix summarizing expected science yields, technical performance metrics, and maturity for a variety of coronagraphs. This was designed to be a living document, periodically updated as coronagraphs mature, new designs become available, and as metrics evolve during HWO studies.

- b. A new software pipeline, enabling a mostly automated workflow from coronagraph designs to performance metrics and science yields. This was designed to be a useful tool to facilitate rapid coronagraph trade studies by HWO teams, as well as facilitate rapid evaluations of coronagraphs for coronagraph designers, accelerating iterations of coronagraph designs and trade space explorations.
- c. Reports such as this one, as well as an upcoming final report with more comprehensive details on the survey.

A draft table showing the currently identified coronagraph designs and metrics is shown in Figure XX. This figure is meant to show a high-level organization, with details of rows and columns being covered in later sections. The columns in the matrix correspond to the 19 coronagraph designs collected from the community so far, which are organized into categories depending on the basic type of architecture. The rows correspond to metrics and are grouped into three categories: science yield; coronagraph performance and robustness; maturity, compatibility, and programmatic considerations. The CDS team consists of three working groups, each tasked with computing or evaluating each of these three groups of metrics for all coronagraph designs to provide a consistent set of evaluations. The majority of the science yield metrics, as well as coronagraph performance and robustness metrics, will be computed by a mostly automated pipeline which will provide a high degree of consistency across all coronagraph designs. The category of maturity, compatibility, and programmatic considerations does not lend itself to machine computations and is therefore somewhat subjective, but the CDS attempted to apply a uniform process to evaluate those metrics for all coronagraphs.

These products are available at the following link: <https://exoplanets.nasa.gov/exep/resources/documents/>.

3.2 Survey Process

The first step for an interested coronagraph designer to participate in the CDS was to fill out a quad chart with some basic high-level information about their design. This quad chart is shown in Figure 5. For purposes of the survey, a “coronagraph design” was defined as anything that can be reasonably described by this quad chart (including cases where a couple of the bullet points may not be applicable). At a high level, this meant a design of the part of the high-contrast imaging system between the entrance and exit pupils of the instrument. The intent was to mainly focus on the design of the coronagraphic masks, optics, and their geometric layout. However, designers could optionally include coronagraphic subsystems such as shapes on deformable mirrors as part of their design (but not detector or post-processing, which would be handled by separate working groups). The design fidelity ranged from a very simple and abstract unfolded layout with idealized components, to designs with more realistic layouts that also include realistic manufacturing errors. This allowed capturing a diversified portfolio of designs, ranging from already mature ones that have the lowest technical risk, to less mature ones that provide opportunities to reduce the risk and cost of the telescope and/or enhance the science. The CDS made the following additional assumptions:

1. Designs needed to have a plausible TRL path on time scales relevant to the Habitable Worlds Observatory (TRL5 by early 2030s, given adequate funding). However, since robotic servicing is expected to be possible, we did not turn away less mature designs.
2. Baseline telescope apertures are shown in Figure 7 (created by the USORT and meant to be similar to the JWST pupils). CDS prioritized the off- and on-axis apertures shown in Figure 7, but also considered other custom apertures submitted by coronagraph designers. These apertures were provided to all interested designers.
3. Telescope diameter was assumed to be 6m (inscribed).
4. The telescope, but not necessarily the coronagraph, was assumed to have IR, V, and UV capabilities. The initial focus of our survey was on visible coronagraph designs, but it was designed to be compatible for IR and UV-specific designs.

5. The CDS accepted designs for technologies and coronagraph subsystems that were not complete systems, as long as this subsystem functioned somewhere between the entrance pupil and the exit pupil of the coronagraph instrument. For example, a new design for a low-order wavefront sensor was within scope, but new detector technologies were outside – those will be handled by a separate detector working group.

Name of Design or Technology

Name(s) of designer(s) and contact email

<p>Basic description</p> <ul style="list-style-type: none"> • Driving motivation(s) for this design • Key advantages, disadvantages, and distinguishing characteristics • Anything else you want to point out 	<p>Current maturity and TRL</p> <ul style="list-style-type: none"> • Estimated TRL of components (see Technology Readiness Level Definitions) • Brief description of current lab performance (provide references, if any) • Anything else you'd like to point out
<p>Optical layout diagram or technical description</p> <ul style="list-style-type: none"> • Between the exit pupil of telescope and science plane • Can be very simple and unfolded – we are not considering details of packaging, relays, etc. • We will assume a standard baseline LOWFS, DMs, detector, etc. However, if one of those subsystems is special or unusual in your design, please include it. • If you wish to submit a coronagraph “technology” rather than complete “design”, CDS can provide a baseline design for you to pair with your technology 	<p>Risks and next steps</p> <ul style="list-style-type: none"> • Narrative of next steps, including tall poles, risks, and long-lead items • Feasibility of getting to TRL 5 (if applicable), as well as any rough estimates about cost and time to TRL 5 (if available)

Figure 5: Coronagraph quad chart questionnaire: the first step for prospective coronagraph designers to enter the design survey.

The submission of a coronagraph quad chart initiated the workflow diagrammed in Figure 6. The CDS team (blue boxes, each corresponding to one of the three working groups within the CDS) worked with the coronagraph design teams (green box, who can be anyone in the world) to obtain two items: the “coronagraph” operator (section 3.2.1) and information on coronagraph maturity and compatibility with HWO (section 3.2.3). The coronagraph operator was processed by a mostly automated pipeline, which produced two sets of outputs: the “yield input package” (a set of standard files needed for the computation of yields, see section 3.2.1), and a standard set of coronagraph performance metrics (contrast, throughput curves, robustness, etc. — see section 3.2.3). A mostly automated yield code (Altruistic Yield Optimizer, or AYO²) computed the expected mission yield (see Section 3.2.2.).

This subsection is a brief summary of the survey process (please see⁷ for more details, including more precise definitions and assumptions). First, coronagraph designers across the world were invited to submit coronagraph designs to the survey (green box in Figure 6). The submitted information consisted of a quad-chart description, a “coronagraph operator” (see Section ??), and a filled out maturity questionnaire (see Section 4.6). A coronagraph operator was a machine-executable function that implemented the (static) transformation between the entrance and exit pupil of a given coronagraph, designed to produce high contrast for the case of no wavefront errors at the entrance pupil. For purposes of the CDS, a “coronagraph design” was essentially the collection of above three items.

Designers were encouraged (but not required) to produce designs for versions of HWO-like pupils which existed in 2023 (created by the Ultrastable Observatory Roadmap Team, or USORT – see pupils highlighted in green in Figure 7). However, because producing a coronagraph design for a new telescope pupil often involves

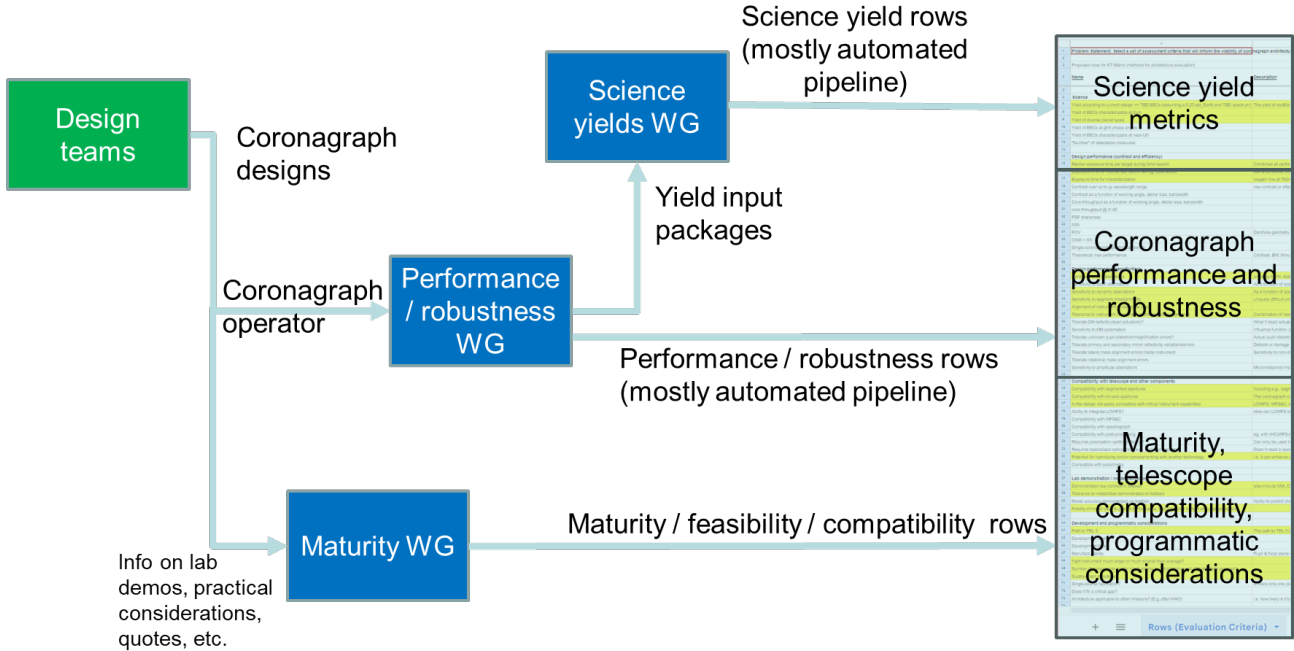


Figure 6: Workflow and pipeline of the Coronagraph Design Survey. Design teams across the world (green box) worked with the CDS working groups (blue boxes) to provide the coronagraph operator and information on the maturity of their coronagraph technology. The CDS working groups computed and evaluated different metrics of that design, including yields, coronagraph performance and robustness, maturity, telescope compatibility, and programmatic considerations (see section 3 for details).

significant effort (beyond the scope of our survey), CDS accepted designs for other pupils that coronagraph designers could produce more easily, or already had on hand. Figure 7 shows a list of telescope pupils for which coronagraph operators were submitted, along with the names of the coronagraph designs that were submitted for each pupil. These designs are described more fully in the next section.

Each coronagraph operator was processed by a mostly automated pipeline (see Figure 6), which produced two sets of outputs: the “yield input package”, and a standard set of coronagraph engineering metrics (contrast, throughput curves, robustness. A mostly automated yield code (Altruistic Yield Optimizer, or AYO⁷) computed the expected mission yield.

4. RESULTS

4.1 Coronagraph design options and description

The CDS surveyed a total of 15 coronagraph concepts, as well as a “2N-th Order Optimal Coronagraph” which represents a theoretical limit for all coronagraphs, and 5 enhancing technologies. These concepts are listed in the table below. Across these 15 concepts, we processed 28 different designs through our pipeline, as well as 16 additional designs for 2N-th Order Optimal coronagraph. See Figure 7 for a list of designs broken down by telescope pupil.

The designs that were collected range from traditional coronagraphs that were already planned for flight on the Roman Coronagraph Instrument (Shaped Pupil Coronagraph, or SPC, Hybrid Lyot Coronagraph, or HLC), and the coronagraphs baselined for the LUVOIR and HabEx mission concepts (Vector Vortex Coronagraph, or VVC, Apodized Pupil Lyot Coronagraph, or APLC) to more modern designs, including the emerging class of photonic chip coronagraphs. Also included were several enhancing technologies, such as better low-order wavefront sensors.^{?, ?, ?}

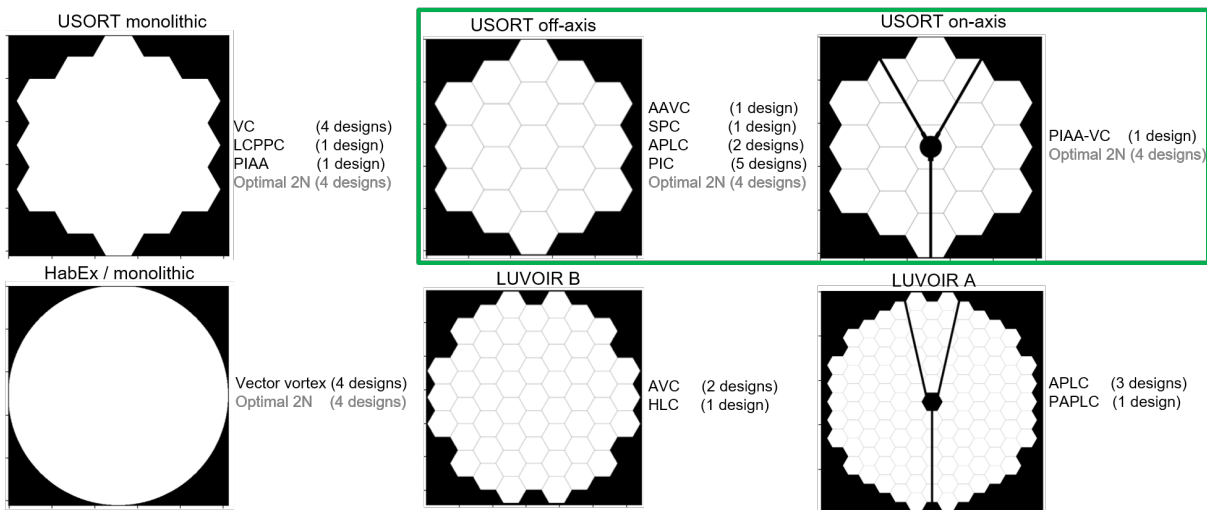


Figure 7: Apertures for which coronagraph operators were submitted to CDS. The two apertures highlighted in green were an early version of an HWO-like aperture, created by the USORT team and meant to be similar to JWST. Next to each aperture is a list of coronagraph operators that were submitted by coronagraph designers and processed by CDS. Apertures highlighted in green were prioritized, but CDS did not turn away designs for other apertures.

Table 1: Coronagraph design submission breakdown indicating which family the design falls into as well as the team that submitted the design.

Coronagraph Family	Coronagraph Design	Submitted by
Primary focal-plane coronagraphs	HLC: Hybrid Lyot Coronagraph	Trauger, Moody, Krist, et al.
	VC: Vortex Coronagraphs	
	AVC: (Apodized VC)	Ruane, Riggs, et al.
	VVC: (Vector VC)	Mawet, Serabyn, Ruane, et al.
	LCPPC: Liquid-crystal Phase-plate Coronagraph	Doelman, Snik, et al.
	EvWaCo: Evanescent Wave Coronagraph	Alagao et al.
	MSPM: Metasurface Scalar Phase Mask	Koenig et al.
Primary pupil-plane coronagraphs	SPC: Shaped Pupil Coronagraph	Kasdin et al.
	PIAA Phase Induced Amplitude Apodization	Sirbu et al.
	IAH: Interferometric Apodization by Homothety	Chafi et al.
Hybrid (pupil+focal) coronagraphs	APLC: Apodized Pupil Lyot Coronagraph	Pueyo, Riggs, et al.
	PAPLC: Phase Apodized Pupil Lyot Coronagraph	Por et al.
	PIAA-VC: PIAA-Vortex Coronagraph	Fogarty et al.
Photonic coronagraphs and theoretical limits	AstroPIC: Astrophysics Photonic Integrated Circuit	Sirbu, Fogarty, Jewell, et al.
	Hybrid PIC	Por et al.
	Optimal 2N-th Order Coronagraph (theoretical limit)	Belikov et al.
Enhancing Technologies	Dual-Purpose Lyot Coronagraphs	Wallace et al.
	Redundant Apodized Pupil	Leboulleux et al.
	Integrated Dynamix Low-Order Wavefront Controller	Trauger et al.
	Single-Mode Fiber Nulling Coronagraph	Serabyn et al.
	Adaptive Apodization for Fiber-Fed Spectroscopy	Carlotti et al.

This list represents a wide variety of coronagraph options spanning many different dimensions. Coronagraph designs vary in which plane they primarily operate in, or even if they operate in free space or photonically. They also represent a wide variety of teams and institutions. Finally, they represent a wide variety of performance and maturity. As will be shown later, higher performing coronagraphs also tend to be less mature, especially for segmented and/or obstructed apertures, hinting at a potential future trade of science performance and margin vs. technical maturity. More generally, the rich diversity of coronagraphs could enable better-informed HWO trade studies, ultimately leading to mission improvements such as increasing science yield and/or margin, and lowering overall mission risk and cost.

For better visualization, the designs were grouped into ones where the key component is: a focal plane element (e.g. HLC, VVC); pupil plane element (e.g. SPC, PIAA, or Phase Induced Amplitude Apodization); hybrid designs where pupil-plane and focal-plane elements are equally important (e.g. APLC, PIAA-vortex); an emerging class of coronagraphs based on photonic chips; and enhancing technologies (e.g. LOWFS, fiber nulling). Photonic designs also include an optimal n -th order coronagraph, which is not a specific design, but a theoretical limit for the performance of all coronagraphs,⁷ but in theory is implementable by a photonic chip. This theoretical limit serves to show how much improvement remains possible with further advances in coronagraph design and technology. In particular, by improving coronagraph throughput, inner working angle, and robustness, HWO exo-Earth yields can be improved by a factor of 2-4 relative to the LUVUOIR and HabEx final reports. (A similar level of improvement can additionally be gained by improvements in certain other coronagraph parameters, such as bandwidth and end-to-end quantum efficiency.) Although normally technologies with TRL as low as photonic chips would be too risky to consider for a flagship, the difference in this case is that the development of photonic chips is driven by a \$1 trillion dollar telecommunications industry. With such a strong technology driver, photonic chips may revolutionize optics in a similar way that electronic chips revolutionized the electronics industry in the 1960s, or CCDs revolutionized imaging (note: CCDs were TRL1 in 1970, 20 years before Hubble launch). Of course, it remains to be seen whether photonic chips will advance on time scales relevant to HWO, and whether development of photonic chips by the telecom industry would proceed in a direction relevant to HWO, but they do provide a potentially game-changing opportunity.

This list represents a wide variety of coronagraph options spanning many different dimensions, ranging from ones that are already implemented for flight on Roman (HLC, SPC), to other relatively mature ones with TRLs of 3-4 (VVC, APLC, PIAA), to less mature but higher performing designs (e.g., PIAA-Vortex, and many of the options on the enhancing technologies group), to the least mature but maximally performing designs (photonic chips and optimal n -th order coronagraph). Coronagraph designs vary in which plane they primarily operate in, or even if they operate in free space or photonically. They also represent a wide variety of teams and institutions. Finally, they represent a wide variety of performance and maturity. As will be shown later, higher performing coronagraphs also tend to be less mature, especially for segmented and/or obstructed apertures, hinting at a potential future trade of science performance and margin vs. technical maturity. More generally, the rich diversity of coronagraphs could enable better-informed HWO trade studies, ultimately leading to mission improvements such as increasing science yield and/or margin, and lowering overall mission risk and cost.

4.2 Key Metrics

Our survey evaluated coronagraphs according to three types of metrics: (1) coronagraph performance and robustness; (2) expected science yields; and (3) maturity, feasibility, and compatibility of coronagraphs with HWO. See Figure 6 for a list of metrics/criteria (rows of the table), as well as coronagraphs (columns of the table, which are the same as the coronagraph list in Table 1). These are described in more detail below.

4.3 Pipeline and Results

4.3.1 Coronagraph performance / robustness

The CDS has developed an automated pipeline which inputs a “coronagraph operator” (provided by a coronagraph designer) corresponding to a given coronagraph design and outputs performance and robustness metrics for that design. The definition of a coronagraph operator bears some discussion, because it reflects many of the simplifying assumptions our study has made. In particular, our aim was to maximally simplify and streamline the evaluation process, so that a large number of coronagraph designs can be processed, while still retaining at

least a minimally meaningful degree of fidelity. This required boiling down the functionality of the instrument to its most important and fundamental essentials.

The CDS defines a coronagraph operator as a linear passive operator $C : E_{in}(x, y; \lambda) \rightarrow E_{out}(x, y; \lambda)$ that transforms the electric field at an input plane into some other electric field at the output plane. This operator may itself be a function of user-specified parameters (such as different masks in a filter wheel, or different settings of the DM), or vary randomly in time (for example, due to thermal or mechanical instabilities). But, fundamentally, everything that there is to know about the performance, robustness, and expected yield of any coronagraph can be derived from its mathematical operator C . The implementation details of the coronagraph, such as masks, layout, etc., are hidden inside the operator and are irrelevant to calculation of coronagraph performance metrics, as long as the mathematical transformation C is known. For example, the broadband response can be calculated by integrating over the λ parameter. A response to extended sources can be computed by integrating point sources over that extended source. Throughput curves can be computed by calculating the response to off-axis point sources at different separations. Sensitivities to aberrations can be computed by passing that aberration as an input to the coronagraph operator. Note that we do not include wavefront control algorithms as part of the definition of C , but C itself can be a function of DM actuators: $C(a)$, where a is a 2D map of DM actuator settings. Thus, any wavefront control algorithm can be wrapped around the coronagraph C , and is in that sense decoupled from the coronagraph definition. A more detailed and formal description is available in the documentation of our pipeline.

Within our pipeline, a coronagraph operator C is represented as a function with a standard definition (written in either Python or MATLAB). This function has the following input parameters: a 2D electric field at the coronagraph input plane, wavelength, and optional parameters representing any variable parts of the coronagraph the designer wishes to include. It outputs a 2D electric field in the coronagraph output plane. For simplicity, we defined the input and output planes as the entrance and exit pupil planes. This means that the coronagraph operator does not actually propagate all the way to the dark zone, but it allows our pipeline to do the final propagation to the science detector in the same way for all coronagraphs. Full details about coronagraph operator definition are documented separately and are available upon request.

Once a coronagraph operator is received by CDS from designers, it is passed to the CDS performance and sensitivity calculator, which produces engineering metrics (see next subsection), as well as a file input package needed by science yield calculators such as the Altruistic Yield Optimizer, or AYO (see Science Yield subsection below). This CDS calculator is a new python toolset designed to facilitate automated batch processing of coronagraph designs, and is one of the key products of CDS that can be leveraged by future teams. It is currently available upon request, with a public release expected in the future. Its key features are: automation and rapid-turnaround (minutes or hours per coronagraph for a complete set of metrics); standardized interface to coronagraph designs and to yield calculators, enabling batch processing and apples-to-apples comparisons; simplicity and low barrier of entry to coronagraph and mission designers. In addition to computing engineering and science metrics for each coronagraph, its output can also be used in the future by coronagraph designers to iterate and optimize their designs as HWO evolves.

The pipeline metrics include: throughput curves, inner working angle, contrast curves, as well as sensitivity to stellar diameter, low-order aberrations, high-order aberrations, and segment-level errors. Note that these plots were produced assuming a fully static system with ideal telescope optics and no manufacturing errors. We expect the initial iteration of our designs and pipeline results to be similarly idealized, and when submitting the first versions of coronagraph operators, coronagraph designers do not need to include DMs in their coronagraph operators. That said, the pipeline architecture is compatible with adding wavefront control to correct for wavefront and manufacturing errors. Future versions of coronagraph operators can either include DMs as part of their coronagraph operator (with DM maps as input parameters), or rely on the wavefront control system model provided by CDS.

In addition to the metrics shown in Figure 7, the pipeline itself can be made publicly available as one of the products of CDS. This will enable post-CDS HWO teams to, if they wish, build on our results to add more coronagraphs and/or more metrics, while maintaining consistency and continuity with original CDS results.

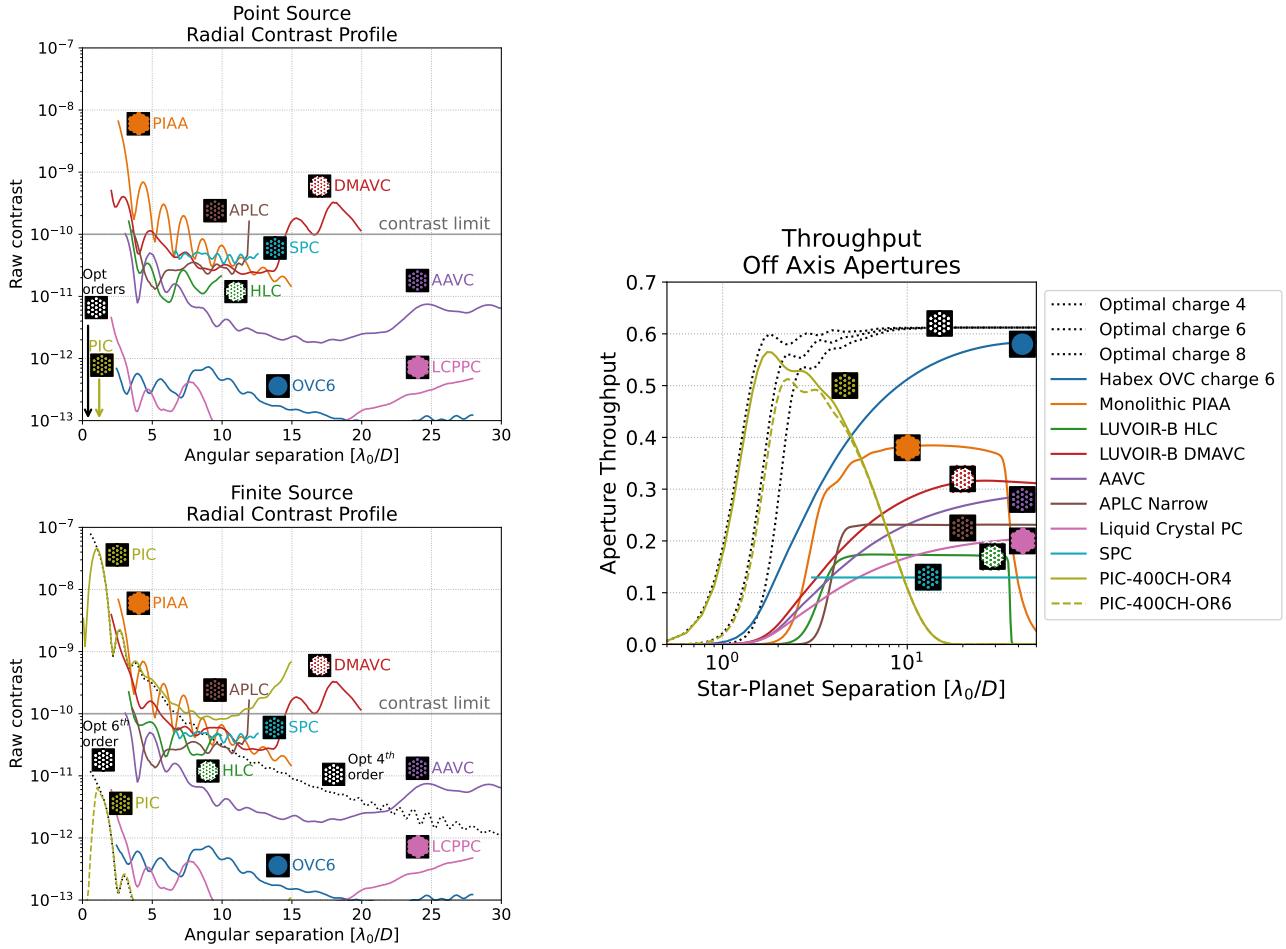


Figure 8: Contrast and throughput curves for off-axis aperture coronagraph designs. Top: point source. Bottom: finite star of $0.05 \lambda/D$ diameter star with no limb-darkening. Spectral bandwidth was 20%. Note that these plots represent a “snapshot”, and many coronagraph designs can be optimized further during future trade studies. Therefore, caution should be taken when comparing performance of different coronagraph designs at this early stage.




Figure 9: Contrast and throughput curves for on-axis aperture coronagraph designs. Top: point source. Bottom: finite star of $0.05 \lambda/D$ diameter star with no limb-darkening. Spectral bandwidth was 20%. Note that these plots represent a “snapshot”, and many coronagraph designs can be optimized further during future trade studies. Therefore, caution should be taken when comparing performance of different coronagraph designs at this early stage.

4.4 Engineering Metrics

Example outputs of the CDS pipeline are shown in Figures 8, 9, and 10. Before we discuss these figures, an important disclaimer is in order. Although these figures compare different coronagraph designs according to the same metrics, most coronagraph designs submitted to our survey were not fully optimized for HWO. Such optimization requires effort significantly beyond the scope of our study. We hope, however, that these plots can serve as a useful starting point and catalyst for such optimizations. In addition, note that we plot coronagraphs of different maturity on the same plots. As a rule of thumb, coronagraphs of lower maturity tend to have more optimistic performance because they may not include some of the practical limitations that more mature coronagraph designs tend to include in their designs. Therefore, caution should be taken when comparing the performance of coronagraphs at different maturity levels.

Figure 8 shows raw contrast and core throughput curves for all coronagraph designs we received for off-axis apertures. Contrast curves are provided for both a point source and a $0.05 \lambda/D$ diameter star. The stellar diameter was chosen based on the HWO ExEP Precursor Science Stars[?] where 68% of the stars are ≤ 1.172 milliarcseconds which translates to $0.05 \lambda/D$ at 650 nm for a 6 m aperture. When we consider the potential HWO target list produced by Mamajek and Stapelfeldt 2023,[?] $0.05 \lambda/D$ is equal or larger than 84% of the targets. The aperture for each coronagraph is shown as a thumbnail next to the coronagraph label. These were computed for a point source (top row) and a $0.05 \lambda/D$ diameter star (bottom row); both use 20% spectral bandwidth. Most of the coronagraph designs have contrasts better than 10^{-10} over their respective designed working angle ranges, even for such large stellar diameters. Some designs, such as PIAA and DMAVC do not have 10^{-10} contrast over the entire range, but their higher throughput makes up for it in terms of science yield (see Sec. 4.5).

The throughput curves in the right panel of Fig. 8 show that there is a lot of range in throughput, depending on the design. Theoretical limits on throughput for the USORT off-axis aperture are represented by dotted lines for different order optimal coronagraphs. The optimal coronagraphs correspond to a fundamental trade as higher order optimal coronagraphs have better robustness to low order aberrations and stellar angular size, at the expense of IWA degradation. The only design that approaches these optimal limits, at least for small working angles, is the PIC (Photonic Integrated Chip, yellow), although its maturity level is low (TRL of about 2). Note that the outer working angle (OWA) for the PIC coronagraphs is set by the number of channels and the IWA is set by the order. The next high-throughput design is the OVC (optical vortex coronagraph, blue), but its throughput drops by about a factor of 2 when the aperture is segmented as shown by the red DMAVC (DM-Apodized Vortex Coronagraph) curve. The most mature coronagraphs on our list (HLC and SPC, which are scheduled for flight on Roman Space Telescope in 2026) are also the ones with the lowest throughput. Thus, two general observations can be made from these plots: (1) coronagraphs have a lot of theoretical potential to improve in terms of throughput and IWA, without sacrificing robustness to low order aberrations; (2) generally speaking, coronagraph designs with greater throughput on segmented apertures are also the ones with the lowest TRL. Thus, coronagraph throughput and IWA is one of the most powerful knobs we have on HWO, but turning them requires concerted investments in coronagraph technology maturation.

Contrast and throughput curves for on-axis aperture coronagraph designs are shown in Figure 9. Although at present, there is a paucity of coronagraph designs for on-axis segmented apertures, the PAPLC and APLC by themselves make on-axis apertures potentially viable. Furthermore, the PAPLC concept makes the coronagraph performance for on-axis apertures potentially competitive with off-axis apertures given its high throughput at the IWA, at least for smaller diameter stars. In addition, there is an emerging concept called the PIAA-Vortex^{?,?} which promises to serve as another viable high-performing option for on-axis apertures, but the design has not been optimized in time for this report. Note that the paucity of coronagraph designs for on-axis apertures does not necessarily reflect a fundamental limitation of coronagraphs for on-axis apertures as compared to off-axis ones. In fact, it is known that aperture shape does not appreciably affect the fundamental physics limits on coronagraph contrast or throughput.[?] This can be seen by comparing the optimal coronagraph curves in Figures 8 and 9. Rather, the paucity of coronagraphs for on-axis apertures may simply be due to them being more challenging to design and test. Thus, it appears that with a sufficiently strong technology development effort, the gap between on-axis and off-axis coronagraphs can be closed.

Another type of plots produced by our survey is shown in Fig. 10. In particular, two examples are shown. The left panel shows the sensitivity of one of our coronagraph designs (LUVOIR-B HLC) to various low order aberrations. These can be used to evaluate several trades, such as the aforementioned relationship between sensitivity and IWA for different order coronagraphs. In the right panel of Fig. 10 we show the sensitivity of the LUVOIR-B DMAVC design to segment piston/tip/tilt instabilities. This type of data, in conjunction with post-processing and dark zone maintenance performance estimates, could potentially be used to inform telescope stability requirements and for joint telescope-coronagraph trades. In addition to these examples, our pipeline is capable of generating a number of additional sensitivity metrics for a variety of aberrations.




Figure 10: Examples of sensitivity plots produced by CDS pipeline. Left: Low order sensitivity plot for the LUVOIR-B HLC design. Note that the further right on the plot a line is, the more robust the coronagraph is to that type of aberration. In this case, the design is very robust to tip and tilt but more sensitive to other low order aberrations. Right: Segment-level sensitivity for the LUVOIR-B DMAVC design. In this case, the design is robust to a wavefront error (WFE) of $< 10^{-5}$ waves but becomes increasingly sensitive above a WFE of 10^{-5} waves.

Finally, our pipeline automates the generation of a “yield file package”, which is a set of files representing a coronagraph that allows exoplanet yield calculations by tools such as AYO and EXOSIMS. We describe these calculations in the next section.

4.5 Science Yield

Exoplanet science yield, or the number of planets that can be detected and/or characterized during a given survey time, is a useful metric to study the performance of a future high-contrast imaging mission and has been used extensively to study previous mission concepts.^{?, ?, ?, ?} Whereas coronagraph metrics like contrast, throughput, and IWA provide a direct connection to instrument performance, exoplanet yield maps these performance metrics onto the astrophysical universe to estimate scientific productivity. Previous studies have shown that the relationships between coronagraph performance metrics and exoplanet yield are not necessarily intuitive and have highlighted the importance of connecting engineering trades to scientific productivity.[?]

To estimate an exoplanet yield, one must define a minimum set of observational criteria that all exoplanets must meet to count toward the yield. Some observational criteria, like exoplanet detection, highlights the need for a larger dark hole to find the planets, while other criteria, like characterization at long wavelengths, highlight the need for a small IWA—the choice of metric can therefore affect coronagraph trade analyses. As such, the CDS adopted four yield metrics to track the potential scientific productivity of each coronagraph design option that would highlight different aspects of coronagraphy (e.g., throughput vs. IWA). Briefly, these three yield metrics are:

1. Broadband detection only (no spectral characterizations or constraints on number of visits)

Table 2. Baseline Astrophysical Parameters

Parameter	Value	Description
η	0.24	Fraction of Sun-like stars with an exoEarth candidate
R_p	$[0.6, 1.4] R$	ExoEarth candidate radius range
a	$[0.95, 1.67] \text{ AU}$	ExoEarth candidate semi-major axis range ^a
e	0	Eccentricity (circular orbits)
$\cos i$	$[-1, 1]$	Cosine of inclination (uniform distribution)
Ω	$[0, 2\pi)$	Argument of pericenter (uniform distribution)
M	$[0, 2\pi)$	Mean anomaly (uniform distribution)
Φ	Lambertian	Phase function
A_G	0.2	Geometric albedo of exoEarth candidate at 0.55 and 1 μm
z	$23 \text{ mag arcsec}^{-2}$	Average V band surface brightness of zodiacal light ^b
z'	$22 \text{ mag arcsec}^{-2}$	V band surface brightness of 1 zodi of exozodiacal dust ^c
n	3.0	Exozodi level of each star

^aFor a solar twin. The habitable zone is scaled by $\sqrt{L_*/L}$.

^bVaries with ecliptic latitude.

^cFor Solar twin. Varies with spectral type and planet-star separation—see Appendix C in Ref. ?.

2. Six broadband detection observations (to account for orbit determination) and subsequent spectral characterization to search for H₂O on all exoEarth candidates (typically at $\lambda = 1.0 \mu\text{m}$ with SNR=5 and R=140)
3. Six broadband detection observations and subsequent spectral characterization to search for O₂ on all exoEarth candidates (at $\lambda = 0.8 \mu\text{m}$ with SNR=10 and R=140)
4. Six broadband detection observations and subsequent spectral characterization to search for CO₂ on all exoEarth candidates (at $\lambda = 1.65 \mu\text{m}$ with SNR=12 and R=70)

In order to compute the yields of these science metrics, we leverage the existing Altruistic Yield Optimizer (AYO)^{?, ?, ?} and Exoplanet Open-Source Imaging Mission Simulator (ExoSIMS)^{?, ?} tools. Both tools adopted the same set of astrophysical inputs and mission parameters, and the same standard format of coronagraph model files that simulate the on- and off-axis coronagraphic PSFs established in previous studies.[?] In studies prior to CDS, coronagraph designers were typically asked to produce those files themselves. With CDS, we simplified this process by incorporating the production of these files into the CDS pipeline. This reduces the burden on the coronagraph designers, who are now only required to produce the coronagraph operator function, as well as the yield modelers, who no longer need to verify the interface with each individual coronagraph designer.

To perform yield calculations, one must make assumptions about the telescope and instrument design that can significantly impact the absolute yields. To avoid casting a pessimistic light on any coronagraph design submitted to the survey, the CDS defined a nominal set of mission and instrument parameters that, while based on plausible telescope and instrument design layouts, are intentionally more optimistic than the LUVOIR and HabEx assumptions. Specifically, we assumed two parallel visible wavelength coronagraphs that each operate in dual polarization and only two aluminum reflections prior to the visible coronagraphs. Additionally, we adopted zero detector noise as our baseline, with an expectation that detector noise constraints could later be derived by quantifying how detector noise degrades yield. We made identical astrophysical assumptions as the LUVOIR and HabEx studies, with the exception of exozodi, for which we distributed the same surface density of exozodi (3 “zodis”) to all stars instead of randomly drawing individual values on a star-by-star basis. The full list of astrophysical and mission parameters can be found in Tables 2 and 3.

For yield calculations using AYO, we built off of recent work by Refs. ? and ?. Specifically, we included wavelength optimization for all detection observations and for spectral characterization observations when searching for H₂O (we adopt a fixed wavelength, SNR, and spectral resolution when searching for O₂ and CO₂). We updated the AYO code to incorporate two significant changes. First, we modified the treatment of the noise floor. Previous versions of AYO implemented a noise floor as a cut-off where planets fainter than a given Δmag

Table 3. Coronagraph-based Mission Parameters

Parameter	Value	Description
General Parameters		
$\Sigma\tau$	2 yrs	Total exoplanet science time of the mission
τ_{slew}	1 hr	Static overhead for slew and settling time
τ_{WFC}	1.3 hrs ^a	Static overhead to dig dark hole
τ'_{WFC}	1.1	Multiplicative overhead to touch up dark hole
D	7.87 m	Telescope circumscribed diameter (USORT aperture)
D_{ins}	6.5 m	Telescope inscribed diameter (USORT aperture)
A	Per USORT aperture	Collecting area of telescope (USORT aperture)
X	0.7	Photometric aperture radius in λ/D_{ins} ^b
Ω	$\pi(X\lambda/D_{\text{LS}})^2$ radians	Solid angle subtended by photometric aperture ^b
ζ_{floor}	None	Raw contrast floor enforced regardless of coronagraph design
$\Delta m_{\text{mag floor}}$	26.5	Noise floor (faintest detectable point source at S/N= 10)
T_{contam}	0.95	Effective throughput due to contamination applied to all observations
IWA_{min}	$1.25 \lambda/D$	Minimum working angle enforced as hard limit
Detection Parameters		
λ_{d}	$0.5 \mu\text{m}^{\text{c}}$	Central wavelength for detection
$\Delta\lambda_{\text{d}}$	$2 \times \Delta\lambda$	Bandwidth assumed for detection (2 VIS coronagraphs simultaneously)
$\Delta\lambda$	Lesser of 20% and design	Coronagraph design bandwidth
S/N_{d}	7	S/N required for detection
$T_{\text{optical,d}}$	0.56^{c}	End-to-end reflectivity/transmissivity at λ_{d}
$\tau_{\text{d,limit}}$	2 mos	Detection time limit including overheads
$n_{\text{pix,d}}$	8^{c}	# of pixels in photometric aperture for detections at λ_{d}
H₂O Characterization Parameters		
λ_{H2O}	$1.0 \mu\text{m}^{\text{c}}$	Wavelength for characterization
S/N_{H2O}	5^{c}	Signal to noise per spectral bin evaluated in continuum
R_{H2O}	140	Spectral resolving power
$T_{\text{optical,H2O}}$	0.32^{c}	End-to-end reflectivity/transmissivity at λ_{H2O} including IFS optics
$n_{\text{pix,H2O}}$	96^{c}	# of pixels per spectral bin in coronagraph IFS at λ_{H2O}
$\tau_{\text{H2O,limit}}$	2 mos	Characterization time limit including overheads
O₂ Characterization Parameters		
λ_{O2}	$0.8 \mu\text{m}$	Wavelength for characterization
S/N_{O2}	10	Signal to noise per spectral bin evaluated in continuum
R_{O2}	140	Spectral resolving power
$T_{\text{optical,O2}}$	0.3	End-to-end reflectivity/transmissivity at λ_{O2} including IFS optics
$n_{\text{pix,O2}}$	61	# of pixels per spectral bin in coronagraph IFS at λ_{H2O}
$\tau_{\text{O2,limit}}$	2 mos	Characterization time limit including overheads
CO₂ Characterization Parameters		
λ_{CO2}	$1.65 \mu\text{m}$	Wavelength for characterization in coronagraph IFS
S/N_{CO2}	12	Signal to noise per spectral bin evaluated in continuum
R_{CO2}	70	Spectral resolving power
$T_{\text{optical,CO2}}$	0.5	End-to-end reflectivity/transmissivity at λ_{CO2} including IFS optics
$n_{\text{pix,CO2}}$	62	# of pixels per spectral bin in coronagraph IFS at λ_{CO2}
$\tau_{\text{CO2,limit}}$	2 mos	Characterization time limit including overheads
Detector Parameters		
ξ	$0 e^- \text{pix}^{-1} \text{s}^{-1}$	Dark current
RN	$0 e^- \text{pix}^{-1} \text{read}^{-1}$	Read noise
τ_{read}	N/A	Time between reads
CIC	$0 e^- \text{pix}^{-1} \text{frame}^{-1}$	Clock induced charge
T_{QE}	0.9	Raw QE of the detector at all wavelengths
T_{dQE}	0.75	Effective throughput due to bad pixel/cosmic ray mitigation

^aFor the USORT aperture with VC6 coronagraph and baseline throughput. Same value adopted independent of coronagraph throughput.

^b D_{LS} is the diameter of Lyot stop projected onto the primary mirror. AYO optimizes this and the associated encircled energy to minimize exposure time on a planet-by-planet basis.

^cExample provided at most likely bandpass; AYO optimizes bandpass and adjusts values accordingly.

were treated as unobservable. This meant that the exposure times of planets approaching the noise floor were unaffected by the noise and not asymptotic in nature. In our new version of AYO, we incorporate a noise floor term that grows linearly with time and appears in the denominator of the exposure time equation, identical to the implementation of the Roman Coronagraph Instrument.⁷ Second, while we maintained the assumption of aperture photometry, we allowed the size and shape of the aperture to be optimized on a planet-by-planet basis. To do so, we calculated the core throughput at all two dimensional stellocentric offsets as a function of core size, then calculated all planet’s exposure times as a function of core size. The yield code then chooses the minimum exposure time for that planet. In practice, this means that planets far from the background-limited regime are observed with larger photometric apertures (which does not significantly improve yield, as these planets already have short exposure times), planets in the background-limited regime but far from the noise floor are observed with photometric aperture radii $\sim 1 \lambda/D$, and planets close to the noise floor are observed with very small photometric apertures. To prevent arbitrarily small photometric apertures, we required the photometric aperture area to be larger than one detector pixel.

For each yield calculation, we produced a standard set of outputs. These outputs include:

- a complete record of input parameters
- the set of target stars selected for observations
- the calculated observation plan including exposure times and yields
- yields for a wide variety of planet types detected during the EEC survey⁷
- visualizations showing the distribution of target star parameters
- visualizations showing the observation choices on a star-by-star basis
- visualizations showing the numerical convergence of yield calculations
- visualizations showing the yield tool’s interpretation of the coronagraph performance

Figure 11 shows an example of some of the outputs from the yield code for a calculation of the second yield metric (six observations with H₂O detection) using four optimal 2N-th order coronagraphs (for N = 1,2,3, and 4). The top row of the figure shows the distribution of target star properties while the second row shows the distribution of exposure times for both detection and characterization. The bottom half of the figure shows scatter plots of the target stars selected for observation (and those not selected in gray), each color coded by different observational parameters, from HZ completeness and exposure time, to optimized choice of coronagraph, wavelength, and photometric aperture. Each target star used the coronagraph order best matched to its characteristics (see "Det. Coronagraph" and "Char. Coronagraph" panels in Figure 11), simulating an instrument where different coronagraph masks can be swapped in and out for different stars.

Exoplanet yield is not the only important scientific metric to track. Exposure time is another valuable metric and provides a better understanding of whether additional science could be conducted beyond the modeled survey, as well as what quality of data might result. The CDS considered multiple methods for reporting exposure time metrics. Our study initially focused on reporting the exposure times for a handful of fiducial stars, but this ultimately proved to be challenging. Figure 12 illustrates why—shown are three target lists, each optimized by AYO for different coronagraph masks. The ideal target stars vary dramatically between coronagraph mask designs, such that it is difficult to find a sample of stars that serve as good “fiducials” for all mask designs. For example, picking a solar twin at 10 pc would make for a reasonable target for the DM-AVC, but would lead to implausibly long exposure times for the SPC and 16-channel PIC, which do not accurately portray those coronagraphs’ typical exposure times. It is clear that fiducial target stars can lead to biased outcomes and a coronagraph design’s target stars should be the results of optimized yield calculations.

Given this, the CDS opted to report the median exposure times resulting from yield calculations. Specifically, CDS tracked:

Target Properties

$N_{\text{stars}} = 344$, $N_{\text{obs}} = 2060$, Yield = 57.84
 Max Dist. = 25.1 pc, Max Diam. = 8.9 mas

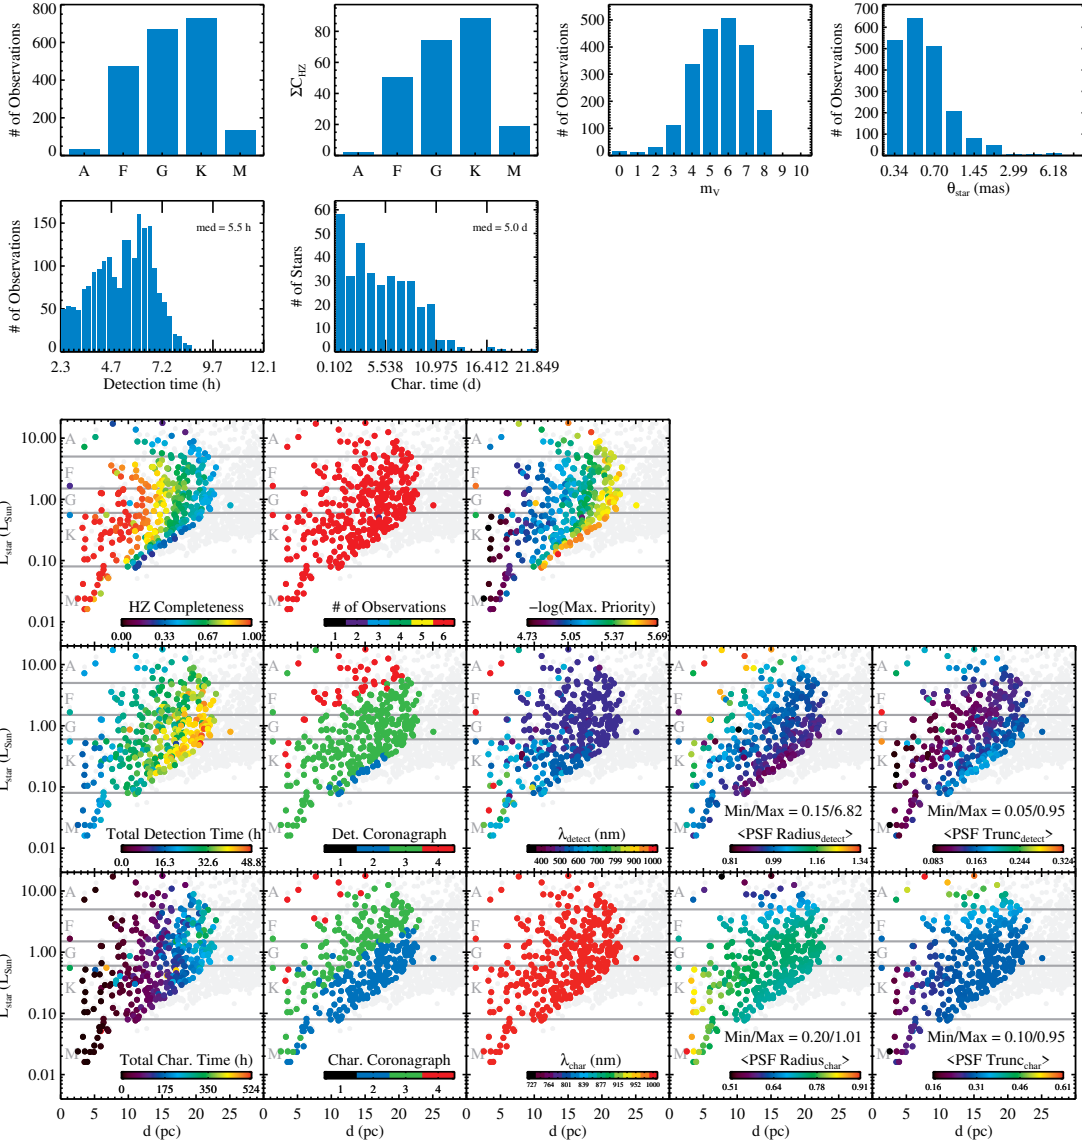


Figure 11: Example of visualizations produced by the CDS implementation of AYO. The results shown are for a simulation of four optimal coronagraph masks evaluated with the second yield metric (six observations plus H₂O search).

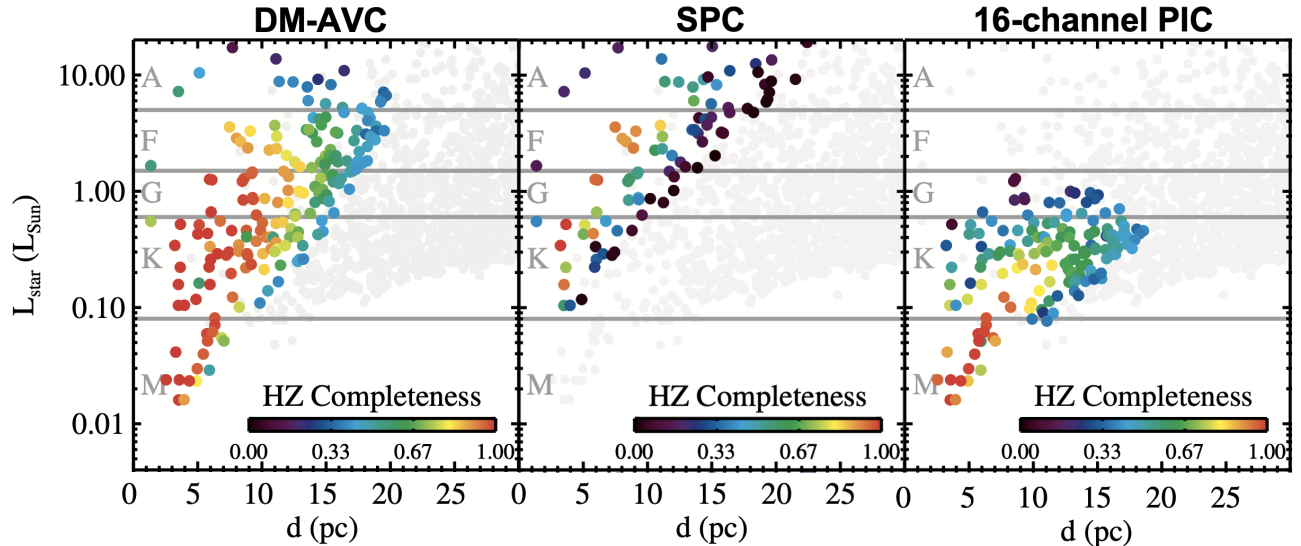


Figure 12: Scatter plots of the targets optimally selected by AYO for three different coronagraph masks. Target lists vary dramatically between coronagraph designs, making any selection of “fiducial stars” difficult.

- The median detection exposure time from EEC survey yields
- The median characterization exposure time to detect H_2O
- The median characterization exposure time to detect O_2
- The median characterization exposure time to detect CO_2

Figure 13 shows the results of all yield and exposure time calculations. All metrics are expressed as percentages, normalized to the results of the optimal coronagraph “limit”. (Note that this limit depends on the assumptions listed above, such as the spectral bandwidth available to the coronagraph instrument, and can be surpassed if these assumptions are varied.) The absolute yields and exposure times are listed parenthetically. The Hybrid Lyot Coronagraph (HLC), Shaped Pupil Coronagraph (SPC), Phase-Induced Amplitude Apodization (PIAA), Amplitude-Apodized Vortex Coronagraph (AAVC), and Full Photonic Chip results all adopted a single coronagraph mask design. When evaluating the performance of the Vortex Coronagraph (VC), Deformable Mirror-Apodized Vortex Coronagraph (DMAVC), Apodized Pupil Lyot Coronagraph (APLC), Phase-Apodized Pupil Lyot Coronagraph (PAPLC), and Optimal Coronagraph Limit, we supplied AYO with multiple coronagraph masks of that “flavor” and allowed the code to optimize coronagraph mask selection on a star-by-star basis—this often lead to the preference of 1-2 mask designs.

As shown in Figure 13, many of the coronagraph designs have yields $\sim 50\%$ that of the optimal limit. Notably this includes designs for on-axis apertures (PAPLC), illustrating the improvement in coronagraph design for obscured apertures that has occurred since the LUVOIR study.⁷ It also suggests that while coronagraph designs have progressed, there is still additional room for improvement in the future with the potential to roughly double yields (or even more, if some of the assumptions in Table 3 can be improved by future technology development).

Median exposure times for many coronagraph designs range from $\sim 1 - 5$ times longer than the optimal limit. In some cases, the median exposure time ratio is < 1 , due to the fact that the optimal coronagraph observes a much larger target list with more challenging stars, resulting in higher yields with a slightly longer median exposure time. While the median exposure time metric is a useful parameter to track, it must be understood in the context of the yield.

Regarding absolute yields, Figure 13 shows that for the baseline mission performance assumptions made in this study, multiple coronagraph designs exist that are consistent with the yield goals recommended by the Astro2020 Decadal Survey.⁷ Given the large number of mission parameters that could be altered to improve yields, we consider the majority of designs at least within “spitting distance” of these goals.

	Focal-Plane Coronagraphs			Pupil-Plane Coronagraphs		Hybrid Coronagraphs			Emerging Technologies	
	HLC	VC (monolith)	DMAVC (segmented)	SPC	PIAA (classic)	AAVC	APLC	PAPLC	Full Photonic Chip	Optimal Coro. Limit
Science Yields										
EEC Yield (VIS detections only)	16% (13)	68% (54)	57% (45)	30% (24)	64% (51)	44% (35)	43% (34)	39% (31)	81% (65)	100% (78)
EEC Yield (Detect + orbits + H2O search)	15% (9)	67% (40)	55% (33)	13% (8)	56% (34)	47% (28)	36% (22)	46% (28)	81% (49)	100% (60)
EEC Yield (Detect + orbit + CO2 search)	5% (2)	50% (15)	32% (10)	2% (1)	29% (9)	28% (8)	15% (5)	44% (13)	86% (26)	100% (30)
EEC Yield (Detect + orbit + O2 search)	14% (7)	69% (33)	57% (27)	14% (7)	61% (29)	51% (24)	41% (20)	53% (25)	84% (40)	100% (48)
Total yield of all planet types	13% (170)	62% (838)	47% (635)	12% (157)	45% (605)	43% (582)	25% (333)	40% (543)	65% (875)	100% (1345)
Exposure Times										
Median detection time for blind survey	5.6x (31 hrs)	1.3x (7 hrs)	1.4x (8 hrs)	3.2x (18 hrs)	1.3x (7 hrs)	2.1x (11 hrs)	2.1x (11 hrs)	2.9x (16 hrs)	1.4x (8 hrs)	1x (6 hrs)
Median char time for H2O	5.7x (26 days)	1.9x (9 days)	2.8x (12 days)	5.8x (26 days)	2.9x (13 days)	2.6x (12 days)	4.2x (19 days)	0.9x (4 days)	0.9x (4 days)	1x (5 days)
Median char time for CO2	1.3x (33 days)	1.2x (30 days)	1.3x (32 days)	0.9x (22 days)	1.2x (29 days)	1.2x (30 days)	1.2x (31 days)	1.3x (32 days)	0.9x (23 days)	1x (25 days)
Median char time for O2	3.1x (33 days)	1.8x (19 days)	2.3x (24 days)	2.4x (25 days)	2.0x (22 days)	2.2x (23 days)	2.7x (28 days)	0.8x (9 days)	1.1x (11 days)	1x (11 days)

Figure 13: Exoplanet yields and exposure times relative to the optimal coronagraph limit listed in the right-most column. Absolute yields and exposure times are listed parenthetically.





	Coronagraph Type	HLC	EWC	VVC	SP(L)C	PIAA(CMC)	PAPLC	Fiber-nulling
	In air or Vacuum?	Vacuum	In air	Vacuum	Vacuum	Vacuum	In air	In air
	Off-axis Monolith	5.2×10^{-10} 3-15 λ/D One-sided	$\sim 10^{-4}$ at $\sim 3\lambda/D$	1.6×10^{-9} 3-10 λ/D One-sided	2.4×10^{-9} 4-10 λ/D 90 deg	1×10^{-8} 2-4 λ/D One-sided		2.5×10^{-5} ~ 0.5 -2 λ/D 360 deg
	Off-axis Segmented			4.7×10^{-9} 3-10 λ/D One-sided			4.2×10^{-8} 2-13 λ/D One-sided*	Would be same
	On-axis Monolith (CGI)	1.6×10^{-9} 3-9 λ/D 360 deg			4.1×10^{-9} 3-9 λ/D 2x 65 deg 3.5×10^{-9} 6.3-19.5 λ/D 360 deg	1.8×10^{-7} 1.3-6.5 λ/D One-sided		Would be same
	On-axis Segmented					1.9×10^{-8} 3.5-8 λ/D One-sided		Would be same

Figure 14: Broadband lab contrast results for different apertures.

4.6 Maturity / compatibility with telescope / programmatic considerations

A final set of metrics the CDS assessed had to do with the coronagraph maturity, compatibility with telescope (as well as with other components), and development/programmatic considerations. Because these types of metrics are typically not machine-computable, we developed a standard questionnaire to collect this information from coronagraph designers. Following the questionnaire, CDS worked with designers to ensure that the information is complete, accurate and consistent between different designs before filling out the relevant rows of the information matrix. A discussion of all the resulting information is outside the scope of this paper, but a key takeaway is that coronagraphs with higher theoretically predicted science yields tend to have challenges in other areas, such as maturity, highlighting the need in continuing investments in coronagraph maturation. The trade of performance vs. maturity is important and non-trivial, and is closely related to the trade of science performance or margin against technical risk. We hope that the data we surveyed about expected science performance and maturity can serve as a good starting point for that trade. (See the information matrix for a comprehensive summary.)

In this section, we focus on one of the most important aspects of maturity, namely laboratory demonstrations. Figure 14 shows broadband laboratory contrasts for different coronagraphs and different apertures. These

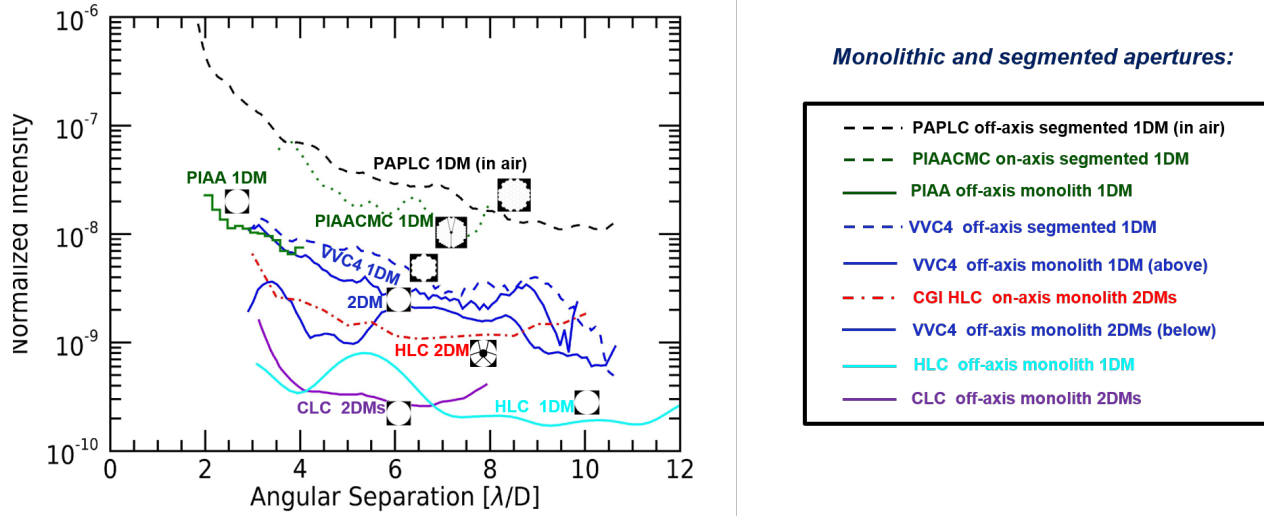


Figure 15: Figure adapted from Mennesson et al. 2024 (<https://arxiv.org/pdf/2404.18036>). Best azimuthal mean contrast (normalized intensity) demonstrated to date by different starlight suppression approaches and laboratory experiments over a $\sim 10\%$ spectral bandwidth. The x-axis shows the angular separation in units of λ/D , where D is the entrance pupil inscribed diameter and λ is the central wavelength of the bandpass. Coronagraphic results were obtained with either one or two DMs, and for different aperture types: off-axis monolith (plain curves); off-axis segmented (dashed curves); on-axis monolith (dashed dotted curve); and on-axis segmented (dotted curve).




Figure 16: Core throughput of the broad-band coronagraph lab setups tested so far. Core throughput is given as a function of angular separation in units of λ/D , where λ is the central wavelength and D is the inscribed diameter of the entrance aperture. At a given separation, the core throughput is computed within a circular aperture of radius $0.7\lambda/D$ for all systems, except for the Roman HLC and SPC-spec coronagraphs, which have highly-spatially-extended PSFs and for which the PSF FWHM region is used instead. *The PIAACMC curve assumes that inverse PIAA optics are used to correct for off-axis PSF distortion (not tested in original lab setup).

represent 7 coronagraph approaches tested via high contrast lab demos in broadband ($>\sim 10\%$). Only 4 were tested in vacuum (HLC, SPC, VVC, PIAA) and 3 only on segmented aperture (VVC4, PIAACMC, and PAPLC). Only one (PIAACMC) was tested on a segmented on-axis aperture. Currently, the best results are found on monolithic apertures. The contrast curves from these results are also shown in Figure 15.

In Figure 16, we also show the theoretical (not lab-measured) throughput for lab setups from Figure 15. As shown in the Science Yields section, throughput is an important driver of science performance and in some sense as important as contrast. However, coronagraphs with currently deepest contrast demonstrated in the lab tend to have lower throughput, and used monolithic apertures. Therefore, advancing lab performance of high throughput coronagraph on segmented apertures is critical.

5. FINDINGS

The Coronagraph Design Survey has completed its ~ 1.5 -year study of coronagraph designs. During this time, the CDS solicited coronagraph designs from the community, assessed their maturity, instrument performance and robustness, and their potential scientific performance. Here we summarize the most noteworthy findings resulting from these efforts.

The CDS acknowledges that the coronagraph design community is expansive, with involvement spanning multiple continents and countries. Continued international involvement in coronagraph design for HWO could result in improved designs in the future. The breadth of submitted coronagraph designs, from traditional Lyot-style coronagraphs to photonic integrated circuits, reflects the creativity of the community and ingenuity that can be expected to address HWO’s challenges.

Substantial progress has been made since the final LUVOIR and HabEx reports. Multiple new coronagraph designs for use with segmented apertures have been invented with promising performance metrics. This is true for both on- and off-axis telescope designs. In fact, this is especially true for on-axis designs, so that on-axis telescopes appear to be a viable option that could achieve the Astro2020 Decadal Survey’s recommended science goals. We expect design improvements to continue. Critically, we note that due to limitations in the scope of the CDS study, it was not possible to iterate and optimize designs. It is quite possible there may already exist design improvements that are “low hanging fruit.”

Many, but not all, of the coronagraph designs submitted to the CDS study adopt a similar optical layout. By adopting an optical layout that enables as many coronagraph mask designs as possible, HWO may be able to extend the instrument development timescale, potentially improving science yield while reducing risk. At the same time, multiple testbeds with different layouts may be needed to develop a properly diversified portfolio of coronagraph architectures.

The CDS pipeline has reduced some barriers to studying HWO coronagraph design. This is in part due to the simplicity of the interface between coronagraph design and exoplanet science yields, as the file package required by the interface is now automatically generated by our pipeline. The CDS pipeline also standardized metrics for comparison, e.g., by calculating contrast consistently for all designs. This standardization and automation sped up the CDS analyses, providing rapid turnaround of calculations while enabling apples-to-apples comparisons. HWO could benefit from further pipeline development with a focus on establishing standardized interfaces between key codes and automated execution.

The CDS pipeline also automated calculations of sensitivity to wavefront aberrations. CDS finds that there are designs that, in theory, can provide adequate robustness to several of expected HWO-like wavefront aberrations.

The CDS exoplanet science yield calculations showed that there are a number of coronagraph designs that are, in theory, capable of achieving or exceeding the Astro2020 Decadal Survey’s recommended science goals with a ~ 6 m inscribed diameter telescope, notably both for on- and off-axis options. Additional designs are within “spitting distance” of these goals. We expect future study of these designs to result in improvements in science yield and find that premature down-selection based on these results could negatively impact the design trade space of HWO.

The CDS exoplanet science yield efforts also illustrated the complex trade space of coronagraphy. No single instrument parameter fully captures instrument performance. Further, optimized target lists vary dramatically

among coronagraph mask designs and evaluating coronagraph performance via fiducial stars can lead to biased interpretations; target lists need to be jointly optimized with science goals and instrument capabilities. Tracing simulated instrument response to science yield helps us understand the available trade space.

Yield calculations also teach us how to efficiently operate and design a coronagraphic instrument. For example, yield calculations often result in exoplanet detections occurring inside of the nominal IWA, suggesting that this parameter by itself is not a good description of science yield. Combinations of coronagraph masks often result in higher yields than any one mask alone, as some masks are, e.g., better suited for spectral characterization while others are suited for broadband detection. HWO may benefit substantially from carrying forward a number of coronagraph design technologies and approaching instrument design as a “suite” of coronagraphic options.

Of primary importance is the demonstration of high contrast in the laboratory. However, gains in science returns become increasingly marginal as contrast approaches 10^{-10} level, and so demonstrations of improvements in other “knobs” available to coronagraphs (such as throughput, IWA, bandwidth and post-processing) will likely become increasingly important. In fact, it is possible that contrasts already demonstrated in the laboratory⁷ will be sufficient to meet HWO science goals, as long as these demonstrations can be reproduced with a higher-performance coronagraph design that simultaneously demonstrates the required bandwidth, off-axis core throughput, and sensitivity to WF aberrations.

Our report represents a snapshot of the ever-evolving state of coronagraph technology, rather than a final assessment. We did not intend or attempt any kind of down-select or prioritization of different coronagraph designs. A meaningful downselect requires, at the very least, iterating and fully optimizing multiple promising architectures, as well as joint trade studies between the coronagraph, telescope, and other aspects of HWO. We hope that our results and tools will be useful in facilitating such optimizations and trades. In particular, our coronagraph database and pipeline are designed to be regularly updated with new coronagraph designs and new metrics, to make it easier for coronagraph designers to rapidly evaluate and optimize their future designs according to the most important metrics determined by future HWO design trades. Conversely, our pipeline is designed to enable HWO teams to rapidly assess the impact of latest advances in coronagraph design and technology development, and maximize the impact of investments in these areas.

5.1 Suggested future trade studies

The CDS identified a number of possible high priority trades related to coronagraph design.

First, HWO could consider the benefits of multiple parallel coronagraph channels that effectively increase the instantaneous bandwidth of the instrument. Such a design change has several benefits that can be reflected in yield calculations and exposure times. Specifically, detection times are shortened due to the faster photon collection rate, and spectral characterization times to search for water vapor can be shortened by a factor of ~ 1.4 as a broader bandwidth covers more water absorption lines⁷ (the latter is a factor we do not account for in the CDS study). There are additional benefits that are not reflected in yield calculations, including improved redundancy, color information obtained during detections that can help identify planets as they move from one epoch of observation to another, etc. Of course additional channels come at the expense of mass, power, and volume, all of which will need to be considered in such a trade.

Another primary trade identified by the CDS is the trade off between raw contrast and other coronagraph performance metrics, such as off-axis throughput. Previous studies have shown that yield can be a stronger function of IWA and throughput than contrast, assuming that the contrast noise floor remains fixed. Given real-world trades between these parameters that occur during coronagraph design, it may be more beneficial to relax the raw contrast from the notional 10^{-10} to achieve higher throughput, relaxed sensitivity requirements, or improved bandwidth. A better understanding of this trade space would likely require a multi-dimensional analysis on a coronagraph-by-coronagraph basis, but may be of great benefit to informing the HWO design space.

Finally, the CDS notes that coronagraph designs for on-axis telescopes have improved substantially, motivating the need to continue evaluating the on- versus off-axis telescope trade. Notably, the PAPLC design appears to provide high throughput, broad bandwidth, and robust contrast, all at a remarkably small IWA, but at the expense of field of view, making it best suited for spectral characterizations. Other on-axis designs are emerging, such as PIAA-Vortex and PICs. We find it likely that a combination of coronagraph designs geared toward different types of observations may continue to improve the expected science return of on-axis telescopes.

5.2 Suggested future analyses

The CDS identified several potential analyses that may be valuable for future HWO studies.

Sensitivity studies are an effective way to identify which parameters are the biggest “lever arms” to pull (or “knobs” to turn) in order to improve science return. However, these studies can mislead if we do not appreciate how many lever arms actually exist. For instance, while exoplanet yield is most sensitive to telescope diameter, there is only one way to increase telescope diameter. On the contrary, while exoplanet yield is only moderately sensitive to end-to-end throughput, there are many dozens of ways to potentially improve throughput without requiring improvements to the coronagraph. While improving the off-axis throughput of the coronagraph is a laudible goal, this may be more challenging than improving the throughput of the overall system. A survey of possible design changes to the system and the potential benefits associated with each could give a fuller understanding of the mission design trade space.

The CDS survey was structured around an information matrix to assess the maturity, robustness, and science performance of *individual* coronagraph designs. However, previous studies have shown (and the CDS confirmed) that performance of the mission can be improved when coronagraphs are not looked at individually, but given the “opportunity” to work together. Future analyses that consider which designs work well together could be a useful path to improving HWO coronagraph performance and science yields.

The CDS pipeline was developed to rapidly assess *potential* coronagraph performance. As such, issues that would degrade performance like WFSC limitations and polarization aberrations are not incorporated in the current version of the pipeline. Upgrading the pipeline to incorporate realistic expectations for WFSC performance as well as polarization aberrations could help with future trade studies for HWO.

ACKNOWLEDGMENTS

This work was supported in part by the NASA Exoplanet Exploration Program, NASA Goddard Space Flight Center, and NASA Ames Research Center. We gratefully acknowledge the contributions of the team members of the Coronagraph Design Survey, as well as all the designers who contributed their coronagraph designs to this work. In addition, we gratefully acknowledge the participation and inputs from the Subject Matter Experts and Observers (Marie B. Levine, Shawn Domagal-Goldman, Julie A. Crooke, Breann N. Sitarski, Arielle Bertrou-Cantou, Armen Tokadjian), as well as invaluable discussions with the CTR, DMTR, and USORT teams. 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). Any opinions, findings, and conclusions expressed in this work are those of the authors and do not necessarily represent the views of the National Aeronautics and Space Administration.