Closing the Technology Gap for the Extremely Large Telescope's Planetary Camera and Spectrograph (ELT-PCS): Development of Polarimetric Imaging and Spectroscopy for Exoplanet Characterisation

D. Smith[a], S. Y. Haffert[a], S. Maharana[b], M. Tecza[b], M. Lucas[c], M. Langlois[d], and F. Snik[a]

[a]Leiden Observatory, Leiden University, Leiden, The Netherlands
[b]Department of Physics, University of Oxford, Oxford, United Kingdom
[c]Steward Observatory, University of Arizona, Tucson, AZ, USA
[d]Centre de Recherche Astrophysique de Lyon (CRAL), Université de Lyon, Lyon, France

ABSTRACT

The Planetary Camera and Spectrograph (PCS) will be the Extremely Large Telescope's (ELT) dedicated instrument for exoplanet characterisation, and polarimetry is one of its core diagnostic capabilities. We report on the status of Work Package 3.4 (WP3.4), which is developing the polarimetric imaging and spectroscopy strategy for PCS. Polarimetry offers a unique window into circumstellar disk structure, planetary atmosphere composition, and rocky-planet surface mineralogy, complementary to the broader-purpose capabilities of the ELT's other first-generation instruments. We summarise the science drivers and contrast requirements across four representative science cases (protoplanetary disks, debris disks, young self-luminous planets, and reflected light from giant and rocky planets), distinguish polarimetric differential imaging (PDI) – our adopted baseline contrast technique – from full polarimetry, which additionally recovers the fractional polarisation needed for microphysical characterisation. We synthesise design lessons from seven heritage high-contrast polarimeters (SPHERE/IRDIS, SPHERE/ZIMPOL, GPI, SCExAO, CHARIS, and VAMPIRES) and identify the ELT-specific challenges that go beyond this heritage: a fast F/0.9 beam that amplifies polarisation aberrations relative to F/4 systems by roughly an order of magnitude, Nasmyth instrumental polarisation and a K-mirror derotator whose orientation strongly affects polarimetric efficiency, and a segmented, time-variable primary mirror. We present the current leading candidate architecture – an achromatic half-wave plate modulator, a choice of pupil- or image-plane beam-splitting strategies, and an IRDAP-heritage Mueller-matrix calibration pipeline – while noting that the PCS instrument concept, and several of these design choices, remain open. We close with a roadmap of near-, medium-, and long-term work toward a Phase A design contribution.

Keywords: extremely large telescope, polarimetry, exoplanet characterisation, high-contrast imaging, polarimetric differential imaging, instrumental polarisation, Mueller matrix calibration, circumstellar disks

1. INTRODUCTION

The Extremely Large Telescope (ELT) Planetary Camera and Spectrograph (PCS) is conceived as the ELT's flagship instrument for direct, high-contrast characterisation of exoplanets, spanning reflected-light imaging of rocky and giant planets through to spectroscopy of their atmospheres. Among the ELT's first-generation instrument suite, PCS occupies a specific niche: where HARMONI provides general-purpose integral-field spectroscopy, METIS targets the thermal infrared, and ANDES delivers ultra-high-resolution spectroscopy, PCS specialises in the combination of extreme contrast, high spatial resolution, and precision polarimetry that none of the others is designed to deliver.* Polarimetry sits at the centre of that specialisation. Because starlight is intrinsically unpolarised while light scattered by circumstellar dust or a planetary atmosphere or surface is not, polarimetric differential imaging (PDI) suppresses the stellar halo far more effectively than intensity-only differential imaging, and the resulting polarised intensity and degree of polarisation directly encode the microphysics – grain size and composition, cloud structure, or surface mineralogy – of the scattering medium.

Further author information: (Send correspondence to D.S.) D.S.: E-mail: dsmith@strw.leidenuniv.nl

*Other ELT instruments and their primary specialisations are summarised for context; PCS is intended to complement, not duplicate, their capabilities.

This paper reports on Work Package 3.4 (WP3.4) of the PCS conceptual design study, which is developing the polarimetric imaging and spectroscopy capability for PCS. No final PCS instrument design yet exists, and the polarimetric modulation scheme, beam-splitting strategy, and detailed calibration approach remain open trade studies; this paper should be read as a snapshot of that ongoing study rather than a frozen design. Our approach has been to (1) establish the science-driven requirements (Sec. 2), (2) draw explicit design lessons from seven heritage ground-based high-contrast polarimeters (Sec. 3), (3) characterise the specific ways in which the ELT's 39 m segmented aperture, Nasmyth platform, and F/0.9 beam make polarimetry harder than at any current 8–10 m telescope (Sec. 4), and (4) translate those lessons into a current leading-candidate architecture and a practical calibration strategy (Sec. 5). We close with a near- to long-term roadmap (Sec. 6) toward a Preliminary Design Review contribution.

2. SCIENCE DRIVERS AND REQUIREMENTS

2.1 A fundamental distinction: PDI versus polarimetry

Two related but distinct techniques motivate the PCS polarimetry sub-system, and we find it important to state the distinction explicitly because it shapes both the science case and the calibration strategy. *Polarimetric differential imaging* exploits only the polarised flux, Q and U: subtracting orthogonal polarisation states cancels the unpolarised stellar halo to first order and reveals the polarised signal from circumstellar or planetary scattered light at contrasts the halo would otherwise hide. PDI is a *contrast* technique, and it is our adopted baseline – robust, flight-proven on SPHERE and GPI, and the first capability PCS will implement. *Polarimetry*, by contrast, additionally recovers the total intensity I, giving the fractional polarisation $q = Q/I$ and $u = U/I$ (or the degree of polarisation $p = \sqrt{q^2 + u^2}$). It is this fractional polarisation, not the polarised flux alone, that maps onto microphysical scattering parameters – dust grain size, porosity, and composition for disks; cloud particle size, atmospheric structure, or surface mineralogy for planets – through radiative-transfer and reflectance models. PDI tells us *where* the scattered light is and how strong it is relative to the star; full polarimetry tells us *what the scattering medium is made of*. PCS is being designed to deliver both: PDI as the baseline contrast technique from day one, and a path to full polarimetry once the telescope's own polarimetric response – effectively its Mueller matrix – can be modelled and removed. An open question we carry forward into the design (Sec. 4.4) is the required granularity of that telescope model: a single model applied uniformly across the field, a pixel-dependent model built via an inversion-code approach, or a differential model referenced to a calibration field.

Figure 1 illustrates the PDI contrast gain achieved on SPHERE/IRDIS, where the polarised intensity image (or its azimuthal counterpart Q_ϕ) recovers a circumstellar signal that is essentially invisible in total intensity I at the same radius.[5]

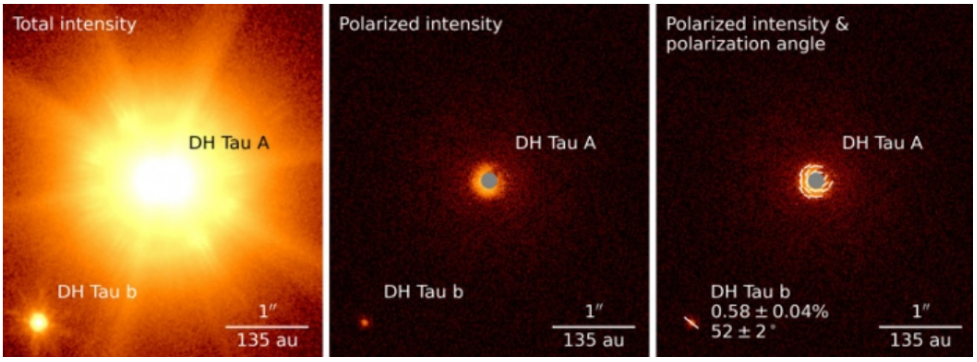

Figure 1. SPHERE/IRDIS H-band images of the young companion DH Tau b: total intensity (left) versus polarised intensity (centre) and polarised intensity with polarisation angle (right). The companion's circumstellar disk signal, revealing a moon-forming disk around DH Tau b, is essentially undetectable in total intensity but clearly recovered in polarised light. Reproduced from van Holstein et al.[5]

2.2 Science cases and contrast requirements

We have considered four representative science cases that together span the contrast, spectral, and polarimetric requirement space PCS must cover.

Protoplanetary disks. Scattered-light PDI using the Q_ϕ/U_ϕ azimuthal Stokes decomposition is already the standard technique at contrasts of 10^{-3}–10^{-4} in Q_ϕ, with SPHERE alone having imaged over 100 systems.[8] These observations

reveal grain properties, disk structure, spirals, gaps, and signatures of planets shaping the disk. PCS, at 5× the angular resolution of an 8 m telescope, extends this to fainter and more distant targets. The relevant spectral range is visible through near-infrared (NIR).

Debris disks. The same PDI technique applies to older, more tenuous debris systems, which are correspondingly fainter in Q_ϕ than young protoplanetary disks; survey campaigns of order 100 stars, such as the Gemini Planet Imager Exoplanet Survey's polarimetric debris-disk campaign, set a useful benchmark for the scale PCS should improve upon. Debris disks trace dust morphology – rings, clumps, and warps – often sculpted by otherwise undetected planets, which raises a specific analysis challenge we flag explicitly: distinguishing a genuinely polarised planet signal from polarised structure intrinsic to the disk itself requires careful joint modelling of both components, not just disk subtraction.

Young, self-luminous planets (thermal). Hot, young giant planets emit primarily in the NIR, with $L \sim 10^{-3}$ to $10^{-5}\,L_\star$ and intensity contrasts of $\sim 10^{-5}$ to 10^{-6}. Their polarisation fraction from atmospheric scattering is at most a few percent, so PDI is not the primary detection technique for this case, though it remains useful for detecting any associated circumstellar disk. PCS's gain here is access to lower-mass planets at smaller separations than current 8 m-class facilities.

Reflected light from giant and rocky planets. This is the most ambitious case and spans two regimes with very different near-term feasibility. Gas giants in reflected light are a realistic near-term PCS target, with intensity contrasts of order 10^{-7}; rocky, potentially habitable-zone planets around M dwarfs, at contrasts of 10^{-10} and polarised contrasts as faint as 10^{-8},[9] are the aspirational case that would require PCS to be fortunate in target availability as much as in instrument performance. Because the planet is intrinsically polarised while the star is not, polarimetry recovers roughly two orders of magnitude of contrast gain relative to direct imaging at the same wavelength. Reaching this regime requires instrumental polarisation corrected to the 10^{-3} level, a dual-beam architecture, an upstream half-wave plate (HWP), and a complete Mueller-matrix calibration of the telescope and instrument; the relevant spectral range is visible light, where rocky-planet reflected-light signatures peak.

Figure 2 illustrates, at a qualitative level, the resulting achievable-contrast picture: simplified contrast curves for a representative 40 m ELT and an 8 m VLT-class telescope at 10 pc, $\lambda = 600$ nm, after extreme adaptive optics, coronagraphy, and angular-differential-imaging (ADI) post-processing, with illustrative Jupiter- and Earth-analogue contrast and separation points overlaid for a Sun-like host star. We stress that this comparison illustrates how ELT aperture improves on the VLT case and is not a realistic end-to-end PCS simulation; by now it is understood that realistic rocky, potentially habitable planets reachable by PCS are expected to orbit M dwarfs rather than Sun-like stars, for which the relevant separations and contrasts differ from those shown here. ELT aperture alone improves both the achievable contrast and the inner working angle relative to the VLT case by roughly an order of magnitude at small separations, underscoring why PCS is being designed from the outset for the most demanding of the four science cases above.

2.3 Key polarimetric requirements

Table 1 summarises the polarimetric requirements distilled from the four science cases above. These are current working requirements rather than formally reviewed top-level requirements, which are themselves a deliverable of this work package.

Table 1. Key polarimetric requirements for PCS, distilled from the science cases of Sec. 2.

Requirement	**Target value**
Instrumental polarisation (IP) correction	$< 0.1\%$ after calibration (absolute accuracy)
Spectral range	NIR ($J/H/K$ bands) and visible (600–1000 nm)
PDI contrast	Q_ϕ/U_ϕ imaging; $U_\phi \approx 0$ used as a calibration quality check
Spectropolarimetry	$R \sim 100$ in $J{+}H$; companion polarisation spectrum to $\pm 0.5\%$
Rocky-planet goal (long-term)	$\Delta p \sim 10^{-3}$ at 10^{-7} intensity contrast

2.4 Observing modes

To support the science cases above without duplicating the broader-purpose capability of HARMONI or METIS, PCS is expected to offer direct imaging, coronagraphic imaging, imaging polarimetry, integral-field-unit (IFU) spectroscopy, high-dispersion spectroscopy (HDS), and spectropolarimetry at both low and high spectral resolution as flexible, combinable observing modes, with the visible and NIR channels described above feeding the imaging-polarimetry, IFU, and HDS paths as appropriate.

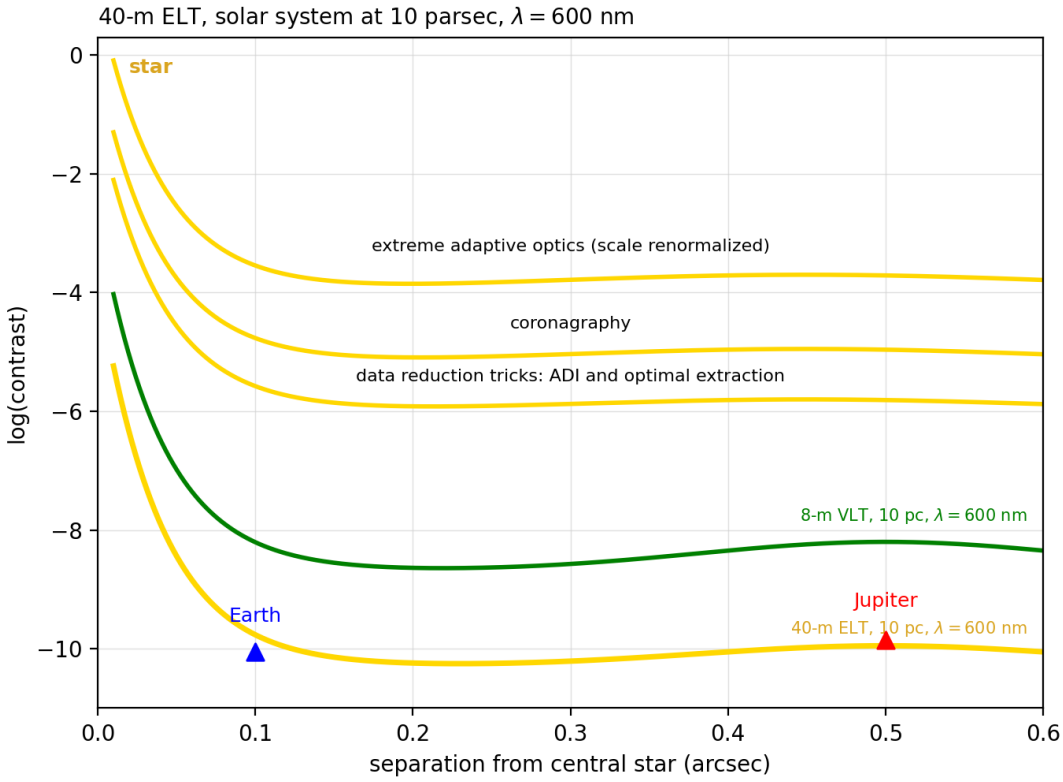

Figure 2. Simplified, illustrative contrast curves for a 40 m ELT-class telescope compared with an 8 m VLT-class telescope, for a Sun-like solar-system analogue at 10 pc and $\lambda = 600$ nm, after extreme adaptive optics, coronagraphy, and ADI/optimal-extraction post-processing. Jupiter- and Earth-analogue contrast/separation points are overlaid for illustration only; this is not a realistic end-to-end PCS simulation, and realistic PCS rocky-planet targets are expected to be M-dwarf habitable-zone planets rather than Earths around Sun-like stars.

3. INSTRUMENT HERITAGE: LESSONS FROM SEVEN HIGH-CONTRAST POLARIMETERS

We base the PCS polarimetry design on explicit lessons drawn from seven existing high-contrast polarimetric instruments across three generations of ground-based extreme-AO systems. For each we summarise what worked, what did not, and the design lesson we carry forward.

3.1 SPHERE and SPHERE/IRDIS

SPHERE's general PDI performance is the field's benchmark: Q_ϕ imaging combined with coronagraphy routinely reaches 10^{-3}–10^{-4} contrast, the dual-beam architecture effectively removes AO speckle residuals, and the azimuthal Stokes parameter U_ϕ – which should be consistent with zero for an optically thin, singly-scattering disk – has become a standard calibration-quality diagnostic that PCS should adopt unconditionally. SPHERE's limitations are equally instructive: the Nasmyth M3 mirror introduces a wavelength-dependent telescope IP of 1.5–3.5%, calibration overhead is large (of order one hour per target plus standards), and the residual IP after calibration ($\sim$0.5%) limits the achievable inner working angle. On SPHERE/IRDIS specifically, the derotator can drop to as little as 5% (H band) or 7% (K_s band) minimum polarimetric efficiency in the worst case, although typical observing strategies that avoid the worst derotator angles perform considerably better in practice; at 45 the derotator removes essentially all polarimetric efficiency, and at other angles it introduces wavelength-dependent linear–circular crosstalk ($|m_{34}| > 0.5$). IRDIS's polarising beam-splitter places regular polarizers behind a regular beam-splitter rather than using a polarising cube, at a cost of a factor of two in photons; in addition, the polarimetric properties of the Nasmyth M3 and M4 mirrors are not, in practice, identical despite nominally matched coatings, and this mismatch drifts further as the coatings age. The resulting design lessons are direct: PDI must be included in PCS from day one; ELT Nasmyth IP from M4 and M5 will resemble SPHERE's M3 problem, so placing the modulator

as far upstream as practicable is essential; and a full Mueller-matrix model of ELT-PCS must exist before first light, with U_ϕ reported as a standard calibration-quality diagnostic in all science products.

3.2 SPHERE/ZIMPOL

Figure 3 shows the SPHERE common-path optical layout and the ZIMPOL detector arm in detail; while no equivalent block diagram yet exists for PCS, this architecture remains the closest structural analogue to what is being considered, and several of its design choices recur directly in Sec. 5. ZIMPOL's kHz electron-multiplying CCD (EMCCD) modulation measures the two orthogonal polarisation states on the same pixels, removing essentially all atmospheric and speckle temporal variation, and achieves the highest-spatial-resolution optical polarimetry yet demonstrated on an 8 m-class telescope (3.6 mas pixels) with polarimetric fidelity across more than a 3000:1 intensity dynamic range.[7] Several less prominent design choices are equally instructive: ZIMPOL carries multiple half-wave plates to steer the polarisation state through its optical train – some are convenience features, but others (notably the one associated with the derotator) are functionally essential – and one HWP is used together with the ferroelectric liquid-crystal (FLC) modulator in an elegant two-stage modulation scheme. ZIMPOL's single largest limiting factor is not instrumental at all: it is the limited adaptive-optics performance achievable at visible wavelengths. Beyond that, the same VLT Nasmyth M3 and derotator IP problems recur (∼4% and ∼3% respectively, with comparable crosstalk), ZIMPOL offers no pupil-stabilised polarimetry (precluding combined PDI and ADI), is visible-only and therefore unsuited to thermal young-planet detection, and its EMCCD charge-shifting introduces pixel-fault calibration overhead. The key design lesson is more nuanced than "faster modulation is always better": it is always better in principle, but kHz-class modulation requires a specialised detector that may introduce more complexity and low-light-performance risk than the polarimetric benefit justifies. A true dual-beam system, as IRDIS demonstrates, can plausibly perform just as well – ZIMPOL's own two detectors were never designed to have identical optical performance, making the system closer to a two-times single-beam design in practice than a true simultaneous dual-beam one. PCS should retain ZIMPOL's two-stage modulation idea where useful, but should not treat kHz modulation as a requirement, and must provide a NIR channel that ZIMPOL itself cannot.

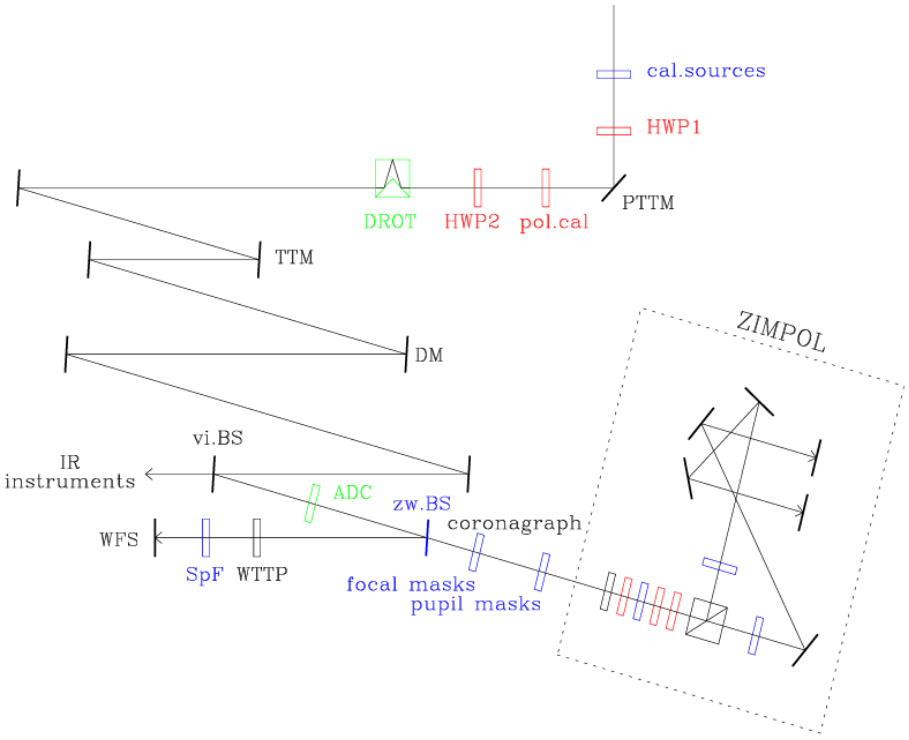

Figure 3. SPHERE common-path optical layout, from the Nasmyth focus (left) through the derotator (DROT), calibration and modulation optics (HWP1, HWP2, pol.cal), deformable mirror (DM), and coronagraph, to the IR-instrument and ZIMPOL (dashed box) arms. Reproduced from Schmid et al.[7]

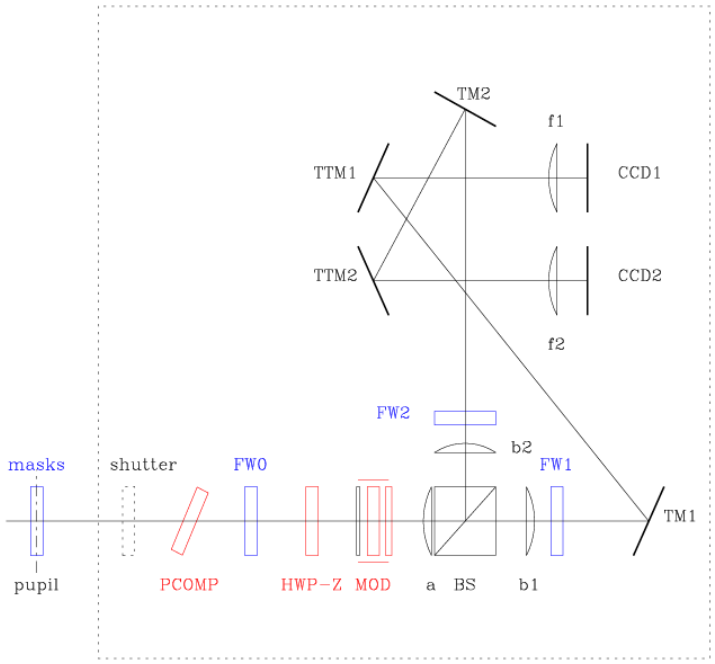

Figure 4. ZIMPOL detector-arm layout (detail of the dashed box in Fig. 3), showing the polarisation compensator (PCOMP), modulator (HWP-Z, MOD), beam-splitter (BS), and dual-beam CCD1/CCD2 detector pair. Reproduced from Schmid et al.[7]

3.3 GPI and SCExAO

The Gemini Planet Imager (GPI) combines a lenslet-based integral field spectrograph with a Wollaston prism, giving simultaneous $J/H/K_s$ spectropolarimetry from a Cassegrain mount that entirely avoids Nasmyth fold-mirror IP;[10] its public data-reduction pipeline and the >500-star GPI Exoplanet Survey are both models PCS should emulate. GPI illustrates two further lessons beyond its headline capability: it offers IFS spectroscopy *or* IFS polarimetry, never both simultaneously, and its dual-beam approach works well specifically because both split polarisation states are sampled identically by the microlens array before any spectral dispersion – a sampling-symmetry principle PCS should preserve in any IFU-based polarimetric mode. GPI's own polarised point-spread function shows clear evidence of polarisation aberrations introduced by the primary mirror, foreshadowing the more severe ELT-scale effects we discuss in Sec. 4.1. Its lenslet sampling does cost spatial resolution relative to a dedicated imager, and complex IP calibration through many optical surfaces remains a real overhead. SCExAO, the modular extreme-AO platform at Subaru hosting CHARIS, VAMPIRES, and other instruments, demonstrates the best NIR Strehl ratio yet achieved on an 8 m telescope and validates low-order wavefront sensing for post-coronagraphic stabilisation, but at a complexity cost: the platform's optical train is now sufficiently convoluted that linear polarisation is no longer an eigenvector of the full mirror train everywhere, and VAMPIRES requires two additional quarter-wave plates simply to route polarisation cleanly through its periscope, on top of three distinct stages of modulation overall. The lesson we draw is correspondingly blunt: SCExAO is a vital platform for experimentation, but PCS needs an architecture with far less inherent complexity, not more.

3.4 CHARIS and VAMPIRES

CHARIS provides simultaneous $J+H+K$ NIR lenslet-IFS coverage (1.1–2.4 μm) and has demonstrated spectropolarimetry via an added Wollaston prism, but its Wollaston sits *in front of* the lenslet array, so the lenslet's own fixed sampling pattern, combined with the resulting sampling mismatch between polarisation states, can itself mimic a polarisation signal – on top of the $\sim$20–30% throughput penalty and very limited field of view inherent to the lenslet design. The lesson is architectural: a dedicated polarimetric imaging channel is preferable to an IFS-based one for disk imaging, with IFS-style polarimetry retained as a valuable secondary mode for companion spectropolarimetry, not the primary disk-imaging path. VAMPIRES, the visible dual-beam fast-modulation channel of SCExAO, uses two simultaneous cameras to eliminate temporal atmospheric noise and FLC modulation at $\sim$100 Hz to suppress speckle variation, achieving polarisation contrast ratios exceeding 1000:1 routinely; it has also added quarter-wave plates to recover polarimetric efficiency lost to derotation,

an approach directly relevant to the PCS K-mirror trade discussed in Sec. 4.2. Its FLC modulator, however, carries real maintenance overhead – temperature sensitivity, long-term drift, and chromatic retardance requiring wavelength-dependent calibration – compounded by a complex, time-varying Mueller matrix when combined with a derotator. The resulting lesson on modulation speed is more nuanced than a single number: modulate as fast as the detector can keep up with; $\sim$100 Hz FLC modulation is still a credible option, and a rapidly spinning HWP may be as well, but a true dual-beam system does not need to out-run the speckle lifetime to perform well, since simultaneous sampling already removes most of the time-variable error term that fast modulation is otherwise compensating for.

3.5 Synthesis: design lessons for PCS

Drawing the seven instruments together, the lessons that most directly shape the PCS design are: (1) modulate as fast as is practical and beneficial – kHz-class modulation is not obviously required, and the right speed depends on the adopted AO and dark-hole performance, not on a fixed target frequency; (2) a dual-beam, simultaneous architecture, with both polarisation states sampled as identically as possible, is the single most consistently successful design choice across all seven instruments; (3) the Mueller-matrix calibration model of ELT-PCS will need to capture the full time variability of the M4–M5–derotator–M6 chain, its wavelength dependence, and the long-term ageing of mirror coatings, not a single static model; (4) the derotator (K-mirror) introduces strong, state-dependent crosstalk and must be modelled at every angle, with operational avoidance of its most damaging orientation; (5) a separate visible and NIR channel is required to cover the full science case, with the caveat that, if implemented with physically separate polarimetric optics per channel, the two channels cannot be used simultaneously; and (6) placing the modulator very far upstream – at an intermediate focus ahead of the telescope's later fold mirrors – is attractive in principle but, as we discuss next, may not be available to PCS at all.

4. ELT-SPECIFIC DESIGN CHALLENGES

Beyond the heritage lessons above, four characteristics of the ELT itself make polarimetry harder than at any current 8–10 m-class facility.

4.1 Fast focal ratio and polarisation aberrations

When a converging or diverging beam reflects from a curved mirror, rays across the beam strike the surface at a range of incidence angles, and the Fresnel equations give each ray a slightly different retardance and diattenuation. The result is a spatially varying polarisation pattern across the pupil – a polarisation aberration – whose effect on the point spread function depends on its specific form. It is useful to distinguish two regimes precisely, since they are sometimes conflated. A *beam shift* is polarisation tip/tilt: the PSF for one polarisation state is laterally displaced from the other, producing gradient-like structure directly in Q and U. The dominant effect of the ELT primary mirror M1, by contrast, is *polarisation astigmatism*, combining a diattenuation (differential amplitude) component and a retardance (differential phase) component: this produces polarised rings and wings in the PSF together with polarised speckles, where the rings are azimuthally polarised in a pattern that can mimic a genuine circumstellar disk signal, while the speckles are randomly polarised but average toward zero only slowly. Segmentation compounds both effects: segment-to-segment coating variation and missing or replaced segments add further pupil-polarisation structure beyond the smooth-mirror case. Because this effect scales as the inverse square of the focal ratio, the ELT's F/0.9 final beam is approximately 20$\times$ worse than an F/4 system such as the VLT for a fixed aperture geometry – not the order-of-magnitude-smaller factor that a naive linear $(1/f)$ scaling would suggest. Anche et al. quantify the resulting impact directly: polarisation aberrations alone limit ELT coronagraphic contrast to of order 10^{-5} in the H and K bands, with the dominant residual having a defocus-like spatial form rather than a speckle-like one, making it harder to subtract with standard speckle-suppression techniques; the same residual structure affects polarimetric Q/U images by mimicking scattered light from a circumstellar disk.[2] Figure 5 shows the resulting PSF polarisation pattern as a function of band and aberration order.

We are pursuing two complementary mitigation strategies. The first is calibration: the systematic polarisation structure of the PSF can in principle be measured using ordinary, unpolarised ("boring") reference stars, modelling the pupil-plane polarisation map and subtracting it from science data; the open question is the level to which this calibration can be pushed before it is limited by time-variable coating ageing or photon noise. The second is hardware compensation: a spatially varying retarder in the pupil plane – conceptually a "polarisation apodiser" – could in principle cancel the retardance aberration map contributed by M1 through M5, most plausibly implemented with patterned liquid-crystal (LC) technology. The central engineering challenge for this approach is that such a compensator must work broadband while LC retarders

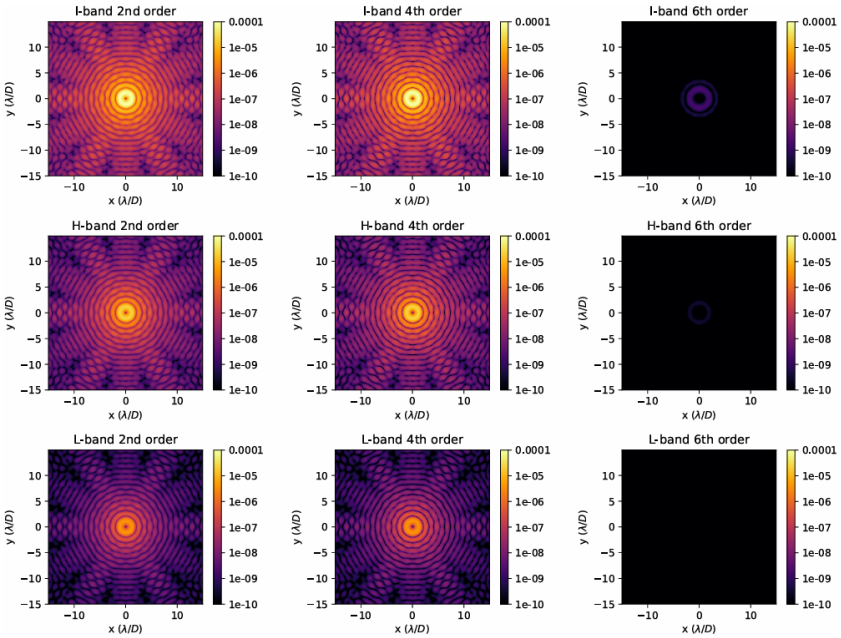

Figure 5. Stellar residuals at *I*-, *H*-, and *L*-band for second-, fourth-, and sixth-order ELT coronagraphs, illustrating the azimuthally polarised ring/wing structure that can mimic a circumstellar disk signal. Reproduced from Anche et al.[2]

are intrinsically chromatic, and must additionally compensate for variable retardance and diattenuation across the optic, whereas current coating processes typically produce only uniform retardance. We regard this as a credible Phase A trade study – quantifying achievable calibration depth against compensator complexity – rather than a committed design element.

4.2 Nasmyth instrumental polarisation and the K-mirror

PCS will be located at a Nasmyth focal station, where the inclined fold mirrors M4 and M5 generate linear instrumental polarisation of order 2–3.4% in the Stokes Q direction.[1] The behaviour of this IP depends critically on the mirror configuration: in a straight-through M4+M5 configuration the IP is independent of zenith angle and correspondingly straightforward to calibrate, whereas in a lateral configuration that adds the facility derotator mirror M6, the IP varies with zenith angle and must instead be modelled per telescope pointing. A two-stage modulation scheme applied at an intermediate focus, ahead of M4 and M5, has been shown to bring the resulting IP below requirement on comparable systems,[1] and that intermediate focus is itself "polarisation-free," making it the textbook-ideal location for a modulator. In practice, however, that intermediate-focus location is not available to PCS: current expectations are that PCS will sit at the ELT's folded Nasmyth station behind the facility derotator M6, where the beam is large and the available interfaces are constrained. The open design question is therefore not whether the modulator can go upstream of M4 and M5 – it cannot – but whether it can be placed upstream of M6 at all; if not, the calibration model must instead absorb the variable M4–M5–rotation–M6 geometry directly, with full wavelength dependence and long-term coating-ageing terms included. Figure 6 reproduces the relevant heritage result for instrumental polarisation correction via an intermediate-focus modulator switch.

The K-mirror field derotator compounds this picture. A three-flat-mirror image derotator used to compensate field rotation at alt-azimuth telescopes, the K-mirror behaves as a quarter-wave plate rather than a half-wave plate; critically, it is not the case that every K-mirror angle is problematic – if the derotator is perfectly aligned in s/p orientation with the rest of the mirror train, there is no issue at all. In practice, however, misalignment at certain angles – most notably near 45 – can drive polarimetric efficiency to zero and introduces wavelength-dependent linear–circular crosstalk that can exceed $|m_{34}| = 0.5$ at other angles. The operational consequence is straightforward to state even though it is not yet formally adopted: avoid scheduling science observations with the K-mirror near 45, model its full Mueller matrix at each operationally used angle via dedicated calibration sequences, log the K-mirror angle in every science header, and optimise its mirror coatings jointly for diattenuation and retardance rather than for reflectivity alone – a joint optimisation that Anche et al. show can reduce ELT diattenuation below 1% without sacrificing reflectivity.[2] A further idea worth carrying into the

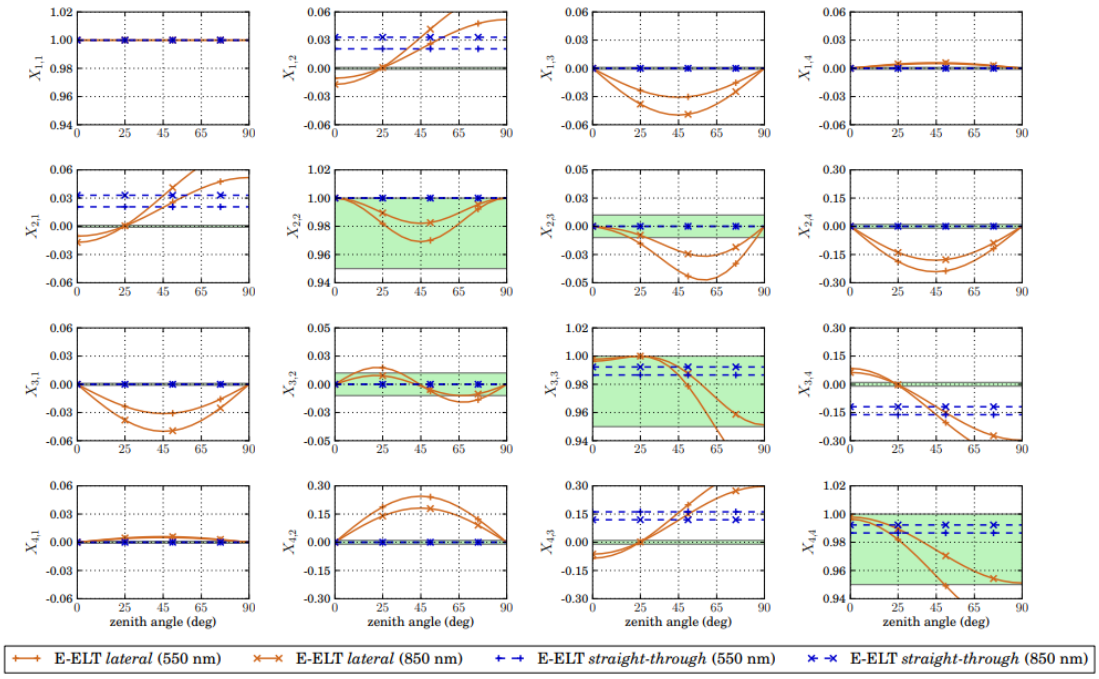

Figure 6. Normalised response matrices versus zenith angle for the E-ELT's two Nasmyth configurations considered, comparing lateral and straight-through M4/M5 arrangements with and without intermediate-focus modulation. Reproduced from de Juan Ovelar et al.[1]

trade study is the placement of half-wave plates immediately before and after the K-mirror specifically to manage the polarisation state entering and leaving it, independent of the primary modulation scheme.

4.3 Segmented primary mirror

The ELT's 39 m segmented M1 introduces polarisation structure that a monolithic primary would not: segment-to-segment coating variation produces non-uniform pupil polarisation, and missing or replaced segments, together with their edge diffraction, add further structure on top of that. Anche et al. explicitly include segment errors and coating non-uniformity in their simulations,[2] with a dedicated follow-up quantifying the segment-to-segment coating-variation contribution specifically,[3] and accommodating this effect may ultimately require spatially resolved modelling of the pupil polarisation map rather than a single segment-averaged term, adding further complexity to the calibration strategy of Sec. 4.4.

4.4 Calibration floor

The achievable calibration floor differs sharply between the visible and NIR. At optical wavelengths, the supply of suitable unpolarised and polarised standard stars is ample, and we do not expect a fundamental accuracy limitation at the 1% level. In the NIR, by contrast, suitable infrared polarimetric standards may be difficult to find in adequate numbers, which could impose a fundamental accuracy floor of order 1% unless a dedicated study identifies a workable calibrator strategy; variable dust and disk properties around candidate calibrators further complicate the picture. Establishing a realistic NIR calibration floor, and the associated choice between a list of multiple calibration sources versus a minimal pair of one unpolarised and one polarised reference source, is accordingly one of the work package's near-term deliverables, with validation against prototype systems planned for early 2028.

5. PROPOSED PCS POLARIMETRY ARCHITECTURE

We stress again that PCS's overall instrument concept is not finalised, and the elements below represent the current leading candidate informed by Secs. 3 and 4, not a frozen baseline.

5.1 Modulation strategy

The half-wave plate remains our preferred modulator over a ferroelectric liquid crystal, on the basis of simplicity, achromaticity, and long-term stability; following the heritage lessons of Sec. 3, we no longer specify a fixed rotation frequency such as 2 Hz, but instead require modulation as fast as is practical, with kHz-class rates not obviously necessary and the precise requirement depending on the adopted AO and dark-hole performance. Placement upstream of M4 and M5 will not be possible; whether the modulator can be placed upstream of the facility derotator M6 at all is, as discussed in Sec. 4.2, a genuinely open question rather than a settled design choice. Separate, independently optimised HWPs are planned for the visible and NIR channels, each targeting a retardance of 180 ± 2 across its respective band.

5.2 Beam-splitting strategy

Table 2 summarises four beam-splitting architectures under evaluation. Options are not mutually exclusive: the choice is expected to depend on science priority, with disk imaging favouring options A or D, companion spectropolarimetry favouring option C, and survey efficiency favouring option B.

Table 2. Beam-splitting architecture options under evaluation for PCS.

Option	Architecture	Notes
A	Wollaston prism, pupil plane	Both polarisation states land simultaneously on the same detector, giving excellent flat-field cancellation; validated on SPHERE/IRDIS and GPI. Preferred for a dedicated NIR polarimetric imaging channel.
B	Cube beam splitter, image plane	Transmits one state and reflects the other onto two separate detector arms; better PSF sampling at the cost of detector co-registration and a chromatic splitting ratio. IRDIS instead uses regular polarizers behind a regular beam-splitter, not a polarising cube, in this role.
C	IFU + Wollaston (GPI-style)	A lenslet IFU with the dispersing grating replaced or supplemented by a Wollaston prism (or a single polarisation grating combining both functions), giving simultaneous spectral and polarimetric information at the cost of throughput.
D	Polarisation grating + quarter-wave plate	A QWP converts linear Q/U to circular polarisation ahead of a liquid-crystal polarisation grating that separates left- and right-circular states. Offers a wider acceptance angle and lower sensitivity to input beam angle than a Wollaston prism – though the ELT's fast beam itself is not the deciding factor in this comparison – at the cost of less flight heritage.

The flight-heritage gap noted for Option D is being actively narrowed: the same broadband (>1 octave), high-diffraction-efficiency liquid-crystal polarisation-grating technology is currently being matured and validated within our group through an unrelated, parallel hyperspectral polarimetric-imaging effort outside astronomy, which will provide useful independent performance heritage ahead of any PCS-specific qualification campaign.

Wherever multiple beam-splitters are used at the back ends of different instrument arms, the split polarisation states must be eigenvectors of whatever optical train lies between the modulator and that beam-splitter – a constraint the SCExAO/VAMPIRES experience (Sec. 3) shows is easy to violate inadvertently as system complexity grows.

5.3 Derotator strategy

We adopt the K-mirror operational and calibration strategy of Sec. 4.2 as our working baseline: avoid the 45 region operationally, use optimised silver or dielectric coatings selected jointly for diattenuation and retardance, maintain a full Mueller-matrix model per rotation state, log the angle in every science and calibration header, and evaluate dedicated half-wave plates immediately before and after the K-mirror to manage its input and output polarisation states.

5.4 Calibration pipeline

Following the IRDAP precedent,[4,6] we plan a Mueller-matrix model of the full ELT-PCS optical path from M1 to the detector, calibrated using a combination of an internal calibration source and unpolarised and polarised standard stars, validated first against SPHERE before being adapted for ELT-PCS, and released as a public pipeline from first light. The model must capture the full time variability of the M4–M5–rotation–M6 chain, wavelength dependence, and long-term coating ageing (Sec. 4.4); a pupil-plane liquid-crystal aberration compensator (Sec. 4.1) remains under study as a means of relaxing the calibration burden this model must otherwise carry.

5.5 Practical implementation

Translating the architecture above into an operational recipe, we have identified four practical building blocks. (1) *Modulation*: rotate the HWP through 0, 22.5, 45, and 67.5 to measure $+Q$, $+U$, $-Q$, and $-U$ in turn, at as fast a cadence as the science case and detector allow. (2) *K-mirror calibration*: enforce an operational exclusion zone around 45, measure the full K-mirror Mueller matrix at each angle used with an internal polarised source, log the angle per exposure, and re-calibrate coating retardance at first light and at least annually thereafter. (3) *Polarisation-aberration mitigation*: map the aberration pattern nightly using unpolarised standard stars, fit it with a Zernike-based polarisation-mode model, and subtract the model from Q and U images before disk analysis, while continuing to evaluate the pupil-plane liquid-crystal compensator as a longer-term hardware alternative. (4) *Calibration pipeline*: build directly on the open-source IRDAP codebase,[4] maintaining a full M1-to-detector model that updates with zenith angle, K-mirror angle, and HWP state, calibrated nightly against internal and standard-star sources and validated on SPHERE ahead of ELT first light.

6. ROADMAP

WP3.4 is structured into three overlapping phases. *Phase 1 – Benchmarking and Requirements* (mid-2026) comprises the instrument benchmarking and trade study reported in this paper together with the translation of PCS science goals into quantitative polarimetric sensitivity and contrast requirements. *Phase 2 – Design Studies* (mid-2026 through early 2027) covers detailed modelling of telescope-induced polarisation systematics across M1–M6 and the K-mirror, the K-mirror waveplate-placement trade study, achromatic-waveplate bandwidth feasibility, the dual-beam-versus-temporal modulation-architecture decision, the infrared calibration-floor study, and a set of polarisation-safe optical design rules for the PCS optical design more broadly. *Phase 3 – Integration and Review* (early 2027 through mid-2028) covers end-to-end AO-coronagraph-detector-polarimetry simulations, a polarimetric error budget, sub-system prototyping of the waveplate and modulator, a detector trade study, feasibility studies for IFU polarimetry and fibre-injected high-dispersion spectropolarimetry, architecture freeze, the Preliminary Design Review package, calibration validation on prototype systems, and long-term stability testing. In the near term, the most immediate priorities are to finalise the science-driven requirements, specify the HWP in detail, begin Mueller-calculus modelling of the ELT optical path, and engage the broader community – including direct contact with the developers of the IRDAP pipeline – on its applicability to an ELT-scale instrument; in-house instrumental-polarisation and polarisation-aberration simulations are now considerably more tractable to begin than they would have been even a few years ago, owing to the availability of capable AI-assisted coding tools.

7. CONCLUSION

Polarimetric differential imaging is the gold standard for high-contrast imaging of circumstellar and planetary scattered light, and PCS will include it from day one; but it is the step beyond PDI – full, spectrally resolved polarimetry recovering Q/I and U/I – that ultimately delivers the microphysical characterisation of disks and exoplanet atmospheres and surfaces that is PCS's primary science goal. The single most important architectural lesson from seven instruments of heritage is a simultaneous, identically sampled dual-beam design combined with a fast, achromatic HWP modulator and a comprehensive Mueller-matrix calibration model. The Nasmyth fold mirrors and the K-mirror derotator are the primary systematic risk and must be modelled at the level of individual exposures; the ELT's F/0.9 beam makes polarisation aberrations of the kind first quantified by Anche et al. roughly $20\times$ worse than at an F/4 telescope such as the VLT, demanding calibration strategies beyond direct heritage transfer. Reaching the 10^{-3} instrumental-polarisation correction required for rocky-planet reflected-light science will require both a complete Mueller-matrix model and a modulator placed as far upstream in the optical train as the final PCS opto-mechanical layout permits. We have presented the current leading-candidate architecture against this backdrop, while emphasising that the underlying PCS instrument concept, and several specific design choices within it, remain open trade studies for the design phases ahead.

ACKNOWLEDGMENTS

The authors thank the wider ELT-PCS consortium for ongoing discussion of the instrument trade study presented here.

REFERENCES

[1] M. de Juan Ovelar, "Instrumental polarization at the E-ELT," *Astronomy & Astrophysics* **562**, A8 (2014).

[2] R. M. Anche et al., "Polarization aberrations in next-generation giant segmented mirror telescopes (GSMTs) I. Effect on the coronagraphic performance," *Astronomy & Astrophysics* **672**, A121 (2023).

[3] J. Ashcraft et al., "Polarization aberrations in next-generation GSMTs II: segment-to-segment coating variations," *Astronomy & Astrophysics* **695**, A28 (2025).

[4] R. G. van Holstein et al., "Polarimetric imaging mode of VLT/SPHERE/IRDIS II. Characterization and correction of instrumental polarization effects," *Astronomy & Astrophysics* **633**, A64 (2020).

[5] R. G. van Holstein et al., "A survey of the linear polarization of directly imaged exoplanets and brown dwarf companions with SPHERE-IRDIS – First polarimetric detections revealing disks around DH Tau B and GSC 6214-210 B," *Astronomy & Astrophysics* **647**, A21 (2021).

[6] R. G. van Holstein, C. U. Keller, F. Snik, and S. P. Bos, "Polarization-dependent beam shifts upon metallic reflection in high-contrast imagers and telescopes," *Astronomy & Astrophysics* **677**, A150 (2023).

[7] H. M. Schmid et al., "SPHERE/ZIMPOL high resolution polarimetric imager," *Astronomy & Astrophysics* **619**, A9 (2018).

[8] M. Benisty, C. Dominik, K. Follette, A. Garufi, C. Ginski, J. Hashimoto, M. Keppler, W. Kley, and J. Monnier, "Optical and Near-infrared View of Planet-forming Disks and Protoplanets," in *Protostars and Planets VII*, ASP Conf. Ser. **534**, 605 (2023).

[9] D. M. Stam, J. W. Hovenier, and L. B. F. M. Waters, "Using polarimetry to detect and characterize Jupiter-like extrasolar planets," *Astronomy & Astrophysics* **428**, 663–672 (2004).

[10] B. Macintosh et al., "First light of the Gemini Planet Imager," *Proceedings of the National Academy of Sciences* **111**, 12661–12666 (2014).