

# On Advanced Estimation Techniques for Exoplanet Detection and Characterization Using Ground-based Coronagraphs

Peter R. Lawson<sup>a</sup>, Richard Frazin<sup>b</sup>, Harrison Barrett<sup>c</sup>, Luca Caucci<sup>c</sup>, Nicholas Devaney<sup>d</sup>, Lars Furenlid<sup>c</sup>, Szymon Gladysz<sup>e</sup>, Olivier Guyon<sup>f,g</sup>, John Krist<sup>a</sup>, Jérôme Maire<sup>h</sup>, Christian Marois<sup>i</sup>, Dimitri Mawet<sup>j</sup>, David Mouillet<sup>k</sup>, Laurent Mugnier<sup>l</sup>, Marshall Perrin<sup>m</sup>, Lisa Poyneer<sup>n</sup>, Laurent Pueyo<sup>o</sup>, Dmitry Savransky<sup>n</sup>, Rémi Soummer<sup>m</sup>

<sup>a</sup>Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA;

<sup>b</sup>Atmospheric, Oceanic and Space Sciences, Univ. Michigan, Ann Arbor, MI 48109, USA;

<sup>c</sup>Optical Sciences, University of Arizona, Tucson, AZ 85721, USA;

<sup>d</sup>Applied Optics, National University of Ireland, Galway, Ireland;

<sup>e</sup>Fraunhofer Institute, Gutleuthausstrasse 1, 76275 Ettlingen, Germany;

<sup>f</sup>Steward Observatory, The University of Arizona, Tucson, AZ 85721, USA;

<sup>g</sup>Subaru Telescope, National Astronomical Observatory of Japan, Hilo, HI 96720, USA;

<sup>h</sup>David Dunlap Inst., Univ. of Toronto, 50 St George St., Toronto, ON M5S 3H4, Canada;

<sup>i</sup>NRC, Herzberg Institute of Astrophysics, Victoria, BC V9E 2E7, Canada;

<sup>j</sup>European Southern Observatory, Alonso de Córdova 3107, Vitacura, Casilla 19001, Chile;

<sup>k</sup>IPAG, 414 rue de la Piscine, Domaine Univ., BP 53, 38041 Grenoble Cedex 09, France;

<sup>l</sup>ONERA, Division Optique Théorique et Appliquée, BP 72, 92322 Chatillon Cedex, France;

<sup>m</sup>Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA;

<sup>n</sup>Lawrence Livermore National Laboratory, 7000 East Ave, Livermore, CA 94550, USA;

<sup>o</sup>JHU Department of Physics and Astronomy, 3400 N. Charles St, Baltimore, MD 21218, USA;

## ABSTRACT

The direct imaging of planets around nearby stars is exceedingly difficult. Only about 14 exoplanets have been imaged to date that have masses less than 13 times that of Jupiter. The next generation of planet-finding coronagraphs, including VLT-SPHERE, the Gemini Planet Imager, Palomar P1640, and Subaru HiCIAO have predicted contrast performance of roughly a thousand times less than would be needed to detect Earth-like planets. In this paper we review the state of the art in exoplanet imaging, most notably the method of Locally Optimized Combination of Images (LOCI), and we investigate the potential of improving the detectability of faint exoplanets through the use of advanced statistical methods based on the concepts of the ideal observer and the Hotelling observer. We provide a formal comparison of techniques through a blind data challenge and evaluate performance using the Receiver Operating Characteristic (ROC) and Localization ROC (LROC) curves. We place particular emphasis on the understanding and modeling of realistic sources of measurement noise in ground-based AO-corrected coronagraphs. The work reported in this paper is the result of interactions between the co-authors during a week-long workshop on exoplanet imaging that was held in Squaw Valley, California, in March of 2012.

Keywords: Exoplanets, coronagraphs, adaptive optics, instrumentation, theory

---

Further author information: (Send correspondence to P.R.L.)

P.R.L. : E-mail: Peter.R.Lawson@jpl.nasa.gov, Telephone: +1 (818) 354-0747

## 1. INTRODUCTION

Several of the co-authors of this paper took part in the conference Seeing the Future with Imaging Science, organized by the US National Academies Keck Futures Initiative (NAKFI). The format of that workshop was to bring together scientists and engineers from diverse backgrounds, including radiologist, astronomers, and physicists, to see if they could learn from each other's experience and envisage new approaches to problems of detection and characterization in imaging. One of the study groups at the workshop discussed possible new approaches to exoplanet imaging, and in particular adapting methods used in medical imaging to the post-processing of extreme adaptive-optics corrected coronagraphic images. Of particular interest was the concept of the ideal observer and the Hotelling observer<sup>1-5</sup> and how these approaches might be applied to exoplanet detection. Although these methods have been examined previously, to the authors' knowledge they have not yet been applied to real astronomical data.

Through a seed grant from NAKFI, a follow-on workshop was organized to bring together experienced astronomers as well as experts in imaging theory. Of particular interest was to formally assess and compare the performance of techniques based on the Hotelling observer with the state of the art, represented by angular or spectral differential imaging using the Locally Optimized Combination of Images (LOCI) and its variations.<sup>6-9</sup> It was hoped that in so doing that we might enable the detection of exoplanets at least an order of magnitude fainter than what is currently possible.

Direct imaging of a planet around another star is exceedingly challenging. For even the closest stars observed with the largest ground-based telescopes, the angular separation between star and planet will be near the classical diffraction limit. Moreover, a typical star will be about a billion times brighter than the planet to be imaged. The planetary image is also buried in "speckle noise," which is the result of uncorrected wavefront errors that propagate through the atmosphere and optical system. This speckle noise has a complex spatio-temporal-spectral structure, which is different from a planetary signal. While algorithms now exist that exploit some properties of the signal and noise, there has been little effort to address the full problem in a rigorous and comprehensive way. The aim of our proposal was to begin to address this need.

For this effort, a balance was sought between theorists and astronomers, and invitations were extended to and accepted by participants representing all the major near-term ground-based coronagraph projects, namely P1640,<sup>10</sup> the Gemini Planet Imager<sup>11</sup> (GPI), the SPHERE instrument of the Very Large Telescope, and the Subaru SCEAO project.

Our exoplanet imaging workshop took place at Squaw Valley, CA, during the week of 25-30 March 2012. Each day of the workshop two different broad topics related to imaging were discussed as the subject of possible future research. These included the topics of detection, astrometry, spectroscopy, future instrumentation, lessons learned, and an optimal observing campaign design. The sessions were run on the same model as the NAKFI sessions. The participants were split into two groups and each group decided what question they would tackle, and each reported separately at the end of each session. This way, the most compelling research topics were brought to the fore and several collaborations were formed for future work and study. An informal count of the proposed work suggests that approximately a half-dozen possible refereed publications will be written as a direct result of the workshop.

The participants readily agreed to undertake an exoplanet imaging challenge, for which simulated blind data sets would be prepared and distributed for the testing of new algorithms. This was seen as a longer-term project, extending beyond the workshop, that would form the basis for a refereed publication. It was also agreed to provide a preliminary report on the imaging study at this SPIE Conference, as described in the following sections.

## 2. CURRENT PRACTICE

Figure 1 summarizes the current state of the art in coronagraphic imaging and data post-processing. Shown in the figure are the  $5\text{-}\sigma$  contrast limits versus apparent angular separation for observations of 1 hour after post-processing of data and using various coronagraphs. The current generation of instruments is represented by 1) the NIRC2 instrument on the Keck II telescope and 2) the NACO instrument at the European Southern Observatory's Very Large Telescope (VLT), which have been used for example to detect the planets around HR

8799<sup>12, 13</sup> and the planet around  $\beta$  Pictoris,<sup>14, 15</sup> whose K-band delta magnitudes are also shown in the figure.<sup>16</sup> The upcoming generation of instruments, coming on line in 2012–2013 is represented by the curves for 1) VLT-SPHERE and 2) the Gemini Planet Imager (GPI). Possible future instruments on telescopes larger than 10-m are represented by the curve for the Thirty Meter Telescope's Planet Finding Instrument. The data for these curves were obtained from the literature and the individual authors themselves and rescaled for 5- $\sigma$  and 1 hour observations. Also shown in the figure are data for solar system planets, as they would appear in reflected light around a Sun-like star at a distance of 10 pc. We would hope through the work discussed here to extend the performance limits of all these instruments.

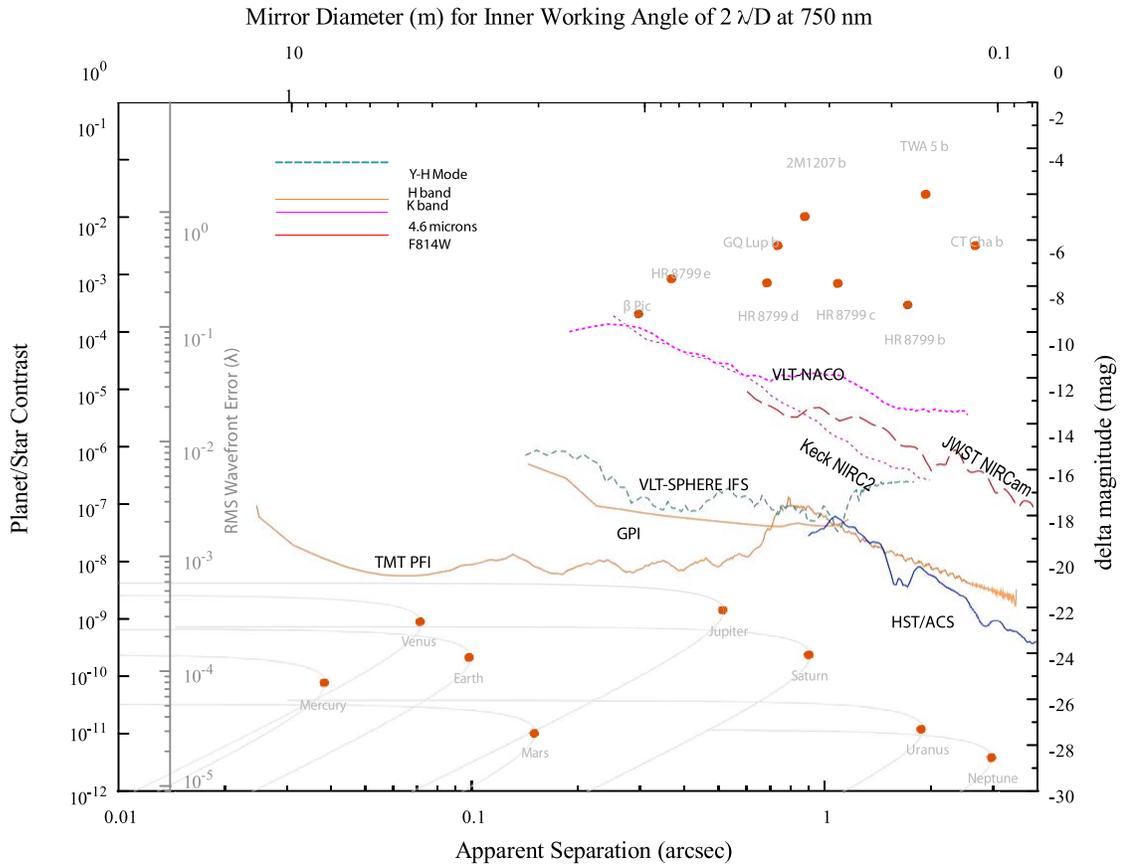


Figure 1:

The 5- $\sigma$  contrast limits after post-processing of 1 hour's worth of data for various coronagraph instruments. As can be seen in the plot, there are roughly three grouping of curves: 1) the current state of the art represented by Keck NIRC2 and VLT-NACO; 2) the next generation of instruments, represented by GPI and VLT-SPHERE; and 3) future extremely large telescopes, represented here by TMT PFI. The contrast curves for JWST NIRC2 and HST/ACS are shown for reference. On the top right of the figure are plotted the K-band contrasts of 9 of the approximately 14 exoplanets imaged to date. In the lower part of the figure are plotted our solar system planets as they would appear in reflected light around a Sun-like star at a distance of 10 pc. On the left-hand of the plot is shown the corresponding RMS wavefront error for a coronagraph using a  $64 \times 64$  element deformable mirror.

### 3. EXPLANET IMAGING CHALLENGE DATA

#### 3.1 Description of Simulated Images

Much of the early discussion amongst the participants centered around the organization of an exoplanet imaging challenge that would allow a fair comparison of the performance of different algorithms, specifically including LOCI and methods using the Hotelling observer. The data sets were prepared by Lisa Poyneer with the help of Marshall Perrin and Jérôme Maire. This section is Lisa's overview of the data sets. The data themselves can be downloaded at <http://olbin.jpl.nasa.gov/nakfi/>.

The challenge data were produced using a suite of simulation tools developed for the Gemini Planet Imager (GPI) project. However, the data do not represent a prediction of actual GPI performance, and should not be used to draw conclusions about GPI observational quality.

The goal for this data challenge is to provide a large data set containing many (but not all) important error sources in a way that will allow testing and comparison of alternate planet finding algorithms. For this initial challenge set the group decided to consider bright stars (so as to not be photon-limited), to observe with field rotation (so that ADI could be used) and to have typical atmosphere and quasi-static speckles, though the quasi-static speckles are due only to phase errors, not amplitude errors. These quasi-static errors were generated to be correlated with pointing, so that a phenomenon that may occur in a real system (correlation of quasi-static errors across different observations) could be explored.

Finally, GPI simulation tools were used to directly produce a noisy IFS datacube, as opposed to dispersing the image onto a detector and then reassembling a datacube via a wavelength solution.<sup>17</sup> This provides a simpler first case where the noise in the data is white; in a later data challenge we could address colored noise due to the dispersion and reassembly. In the following paragraphs we provide more description and explanation of various aspects.

The data challenge scenario contains 100 unique observations, drawn from a simulated GPI observing campaign. The GPI campaign scheduler [ref Savransky] was used with a fake star catalog to produce observation parameters (ra, dec, ha) for each target as observed from Cerro Pachon. For this data set we wanted to limit photon noise so all targets are post-fact assigned a stellar magnitude of 5 (bright), even though most were fainter. Each star retains a specific stellar spectral type.

Along with a star and observation parameters, the campaign scheduler software also produces a random assortment of planets around that star, drawn from our models for planet mass, orbit, etc. Some stars have no planets, some have several. Each planet is given a specific angular separation, position angle in the field of view, and brightness. All planets are assumed to have the spectral type L8. In order to probe contrasts from  $10^{-4.5}$  to  $10^{-6.5}$ , the true planet magnitudes are made brighter by 2 to adjust to the artificially bright stars.

A modified version of the GPI AO simulator<sup>18</sup> was used to produce the instrument PSFs at five wavelength sampled in H-band. This simulator was coupled with a vastly simplified model of the GPI calibration system<sup>19</sup> and instrument optical design to provide quasi-static errors. The GPI AO simulator handles both common-path phase aberrations (usually atmospheric, which the AO system can correct) and non-common-path (NCP) aberrations (which the AO system does not correct, but which should be seen by the Calibration system).

Before the simulation begins, the observation pointing is used to produce a unique time-varying set of NCP errors which produce quasi-static speckle. Completely and time-efficiently modeling all of the optics in GPI, their phase and amplitude variations<sup>20</sup> and the gravity-induced shearing that produces NCP errors, and the Calibration system's measurement and correction of them is far beyond the scope of this effort. Instead, we used a simple model that, while not capturing the true behavior of the system, does provide reasonable levels of NCP phase errors, temporal variations and spatial distributions of speckles.

In short, the pointing information (ra, dec, ha) for the observation is converted to a gravity vector. This gravity vector acts on a very simple cantilever model for the GPI instrument and produces a time-varying and gravity-dependent shearing between optical surfaces during the observation. Two out-of-plane optics were chosen, since these produce chromatic phase and amplitude errors, like those that GPI will actually see. (As noted above, amplitude errors for quasi-static speckles were ignored.) The instrument is assumed to start with a very low level of quasi-static error (due to speckle nulling or the like). As the observation progresses, the optics

shear. The phase errors on each optic are Talbot propagated to the pupil plane (resulting amplitude errors are ignored). At the pupil, a simple model for Calibration performance is assumed, and the corrections are applied in a slow closed loop every 1 minute. This produces a sequence of 3600 files sampled at 1 second, each representing the continually changing NCP error as compensated by a noisy Calibration system at 1-minute intervals.

The main computational stumbling block for the data challenge is the time it takes to fully simulate the AO system. GPI runs at 1 kHz and the full GPI AO simulator tracks AO performance and PSFs generation at each time step. The atmospheric model frozen flow is assumed; the phase screen(s) are translated following wind velocity vectors. Generating PSFs every 1 millisecond is computationally unfeasible for a data challenge of this size. Instead, we chose to use a single short exposure (1 ms) image to approximate a 1-second exposure.

Since the typical wind speed is 10 m/s (more on that below), doing a single short exposure per second approximates the clearing time of phase across the 8-meter aperture. Though not fully accurate (the speckles will be sharper and more variable), this is reasonable. Previous work by ourselves with the AO simulator<sup>17</sup> and by others for non-AO applications<sup>21</sup> show that significant speckle averaging occurs only after a characteristic period dependent on the clearing time (pupil diameter divided by wind velocity).

So to generate a 60-minute observations, we use 3600 individual runs of the AO simulator, each producing a 1-ms short exposure PSF that approximates a 1-second PSF. Each short exposure uses a phase screen generated with a unique random seed.

The wind is the major variability for the AO run. For each observation, a unique time-varying atmospheric profile is generated, comprising the turbulence strength  $r_0$  and wind velocity vector. The mean  $r_0$  is 14 cm; the mean wind speed is 10 m/s. Simple models (e.g. smoother random noise, slow random walks) were used to produce plausible evolution of atmospheric statistics over the one hour observation time. Furthermore, since the instrument observed with field rotation to enable ADI processing, the direction of the wind in the sky changes due to the parallactic angle. This is calculated and tracked and fed into the AO simulator such that the atmospheric errors due to wind velocity (e.g. a butterfly pattern in the dark hole) will move correctly through the observation as the parallactic angle changes.

Note that this approach of using 1-ms AO short exposure has huge computationally savings, but will produce images where the atmospheric speckle noise (non-smoothness) is higher than a full 1 kHz simulation.

Also note that this approach has both the AO residuals and the quasi-static NCP errors in the complex field, so that the generated PSFs fully capture the non-linear interactions between the two types of errors. (A much faster, but less accurate, method would be to use power spectral density (PSD) descriptions of the errors to produce random PSFs reflecting atmospheric and quasi-static PSFs, and then add the terms)

After the 3600 short AO (with NCP) error simulations are run and produce 3600 5-wavelength PSFs, the images are composited via addition into 30-second exposures. These 120 images form the basis of each observation for the data challenge.

For the final step of converting to an idealized IFS data cube of the full astronomical scene, we used the GPI Data Simulation Tool.<sup>17</sup> First, each AO output file is converted to the proper input format for the DST. Next, a script file is produced using the observation information (star brightness and spectrum, observations ra, dec, and ha, planet parameters such as magnitude, separation and position angle in the field.) The data simulator uses the occulted star and un-occulted planet PSFs (produced by the AO simulator) to construct the scene for each 30-second IFS exposure. The details of how this is done, and its verification, are described elsewhere.

Finally, we have added a custom mode to the IFS to simulate an idealized IFS. In this case we use only two of the many noise sources (photon noise and detector read noise). Since units are properly carried in the code, we apply photon noise directly in the ideal IFS cube ( $x-y-\lambda$ ). To approximate the read noise on the detector, while making it (artificially) spatially white the data cube. This was done to match the expected level of read noise in a dispersed and re-assembled data cube. This works out to be 1.75 DN rms per ideal IFS pixel in a 30-second exposure.

During this process the keywords in the FITS files describing appropriate knowledge for data challenge participants (parallactic angle, star brightness, etc) are preserved. The final outputs are saved in a compressed format and collected in an archive.

All 100 observations, each with 120 30-second frames, were provided to all data challenge participants.

### 3.2 Results to Date

Our implementation of the Locally Optimized Combination of Images (LOCI) algorithm is based on the original description presented by Lafrenière<sup>6</sup> wherein we seek to minimize the residuals between a target image and a linear combination of a set of reference images. Using the mathematical formalism of Pueyo,<sup>9</sup> we wish to find the set of coefficients  $\{c_k\}$  given by

$$\operatorname{arg\,min}_{c_k} \left\| T - \sum_k c_k R_k \right\|_O, \quad (1)$$

where  $T$  is the target image,  $\{R_k\}_{k=1}^N$  is the set of reference images, and  $O$  is the optimization zone—a (possibly proper) subset of the target and reference image pixels. The solution to this problem is equivalent to the solution of

$$M_O c - V_O = 0 \quad (2)$$

where

$$M_O[i, j] \diamond = \int_O R_i(x, y) R_j(x, y) dx dy, \quad (3)$$

$$V_O[i] \diamond = \int_O R_i(x, y) T(x, y) dx dy, \quad (4)$$

and  $c$  is the column vector containing the values of  $\{c_k\}$ . The reference PSF constructed from this linear combination is subtracted from a subtraction zone  $S$  of the target image that is smaller than the optimization zone, to minimize the risk of subtracting the planet signal along with the background.

While schemes exist for simultaneously using both angular and spectral diversity images in the reference set,<sup>10</sup> the small number of wavelength bands in our data make the inclusion of spectral diversity in a given set of references marginally useful. Instead, we apply the LOCI algorithm to each wavelength independently, using only rotated images in the same wavelength for references, and then sum the final output images to increase signal-to-noise. For the optimization zones, we use radial annuli with height of 5 pixels and angular extent of 30°, with subtraction zones of one third the height and extent, centered in each subtraction zone. The subtraction zones are defined such that they cover the entire space between the inner and outer working angles of each target image, and overlap by 2% of their annular extents so that no uncorrected areas remain in the final subtracted images (areas where subtraction zones overlap are mean combined). The solution to the linear least-squares (LLS) problem is calculated using standard LLS LAPACK routines. It is important to note that the solution is exact only when matrix  $M_O$  is invertible, which only occurs in the case where the set of references forms a complete basis set of the optimization zone, which is not necessarily guaranteed, depending on the specific realization of the speckle field in the reference images. For our implementation and data set, we find that  $M_O$  is well conditioned for the majority of optimization zones, with singularities occurring only near the inner working angle.

This implementation of LOCI was used in the following manner. Ninety-five of the datasets were analyzed, producing a single processed ‘image’ for each. The thresholding the image was changed to different levels and standard image processing techniques were used to identify as candidate detections features that were of the right size and shape (e.g. do not select as a detection a single pixel, do not select as a detection a region which is a long streak.) This was all done in an automated and fully blind mode. Then the results were compared to the known planet locations to determine if the detections were true positives or false positives.

Our initial results are presented in Figure 2. The planet separation from star is plotted along the x-axis. The planet contrast is plotted on the y axis, expressed as powers of 10. A magnitude 20 planet in the challenge data has a contrast of  $10^{-6}$ . Each black dot is a planet that is hidden in the data. Each colored circle represents a true detection. Bare black dots are either misses, or lie in areas that were not searched (grey regions). The size of the circle indicates the maximum threshold in the LOCI image at which the planet was detected: bright planets at the bottom have higher thresholds. The color of the circle indicates the number of false positive detections in the entire observation for that threshold. Dark red in this case is no false positives. Orange is 5 false positives

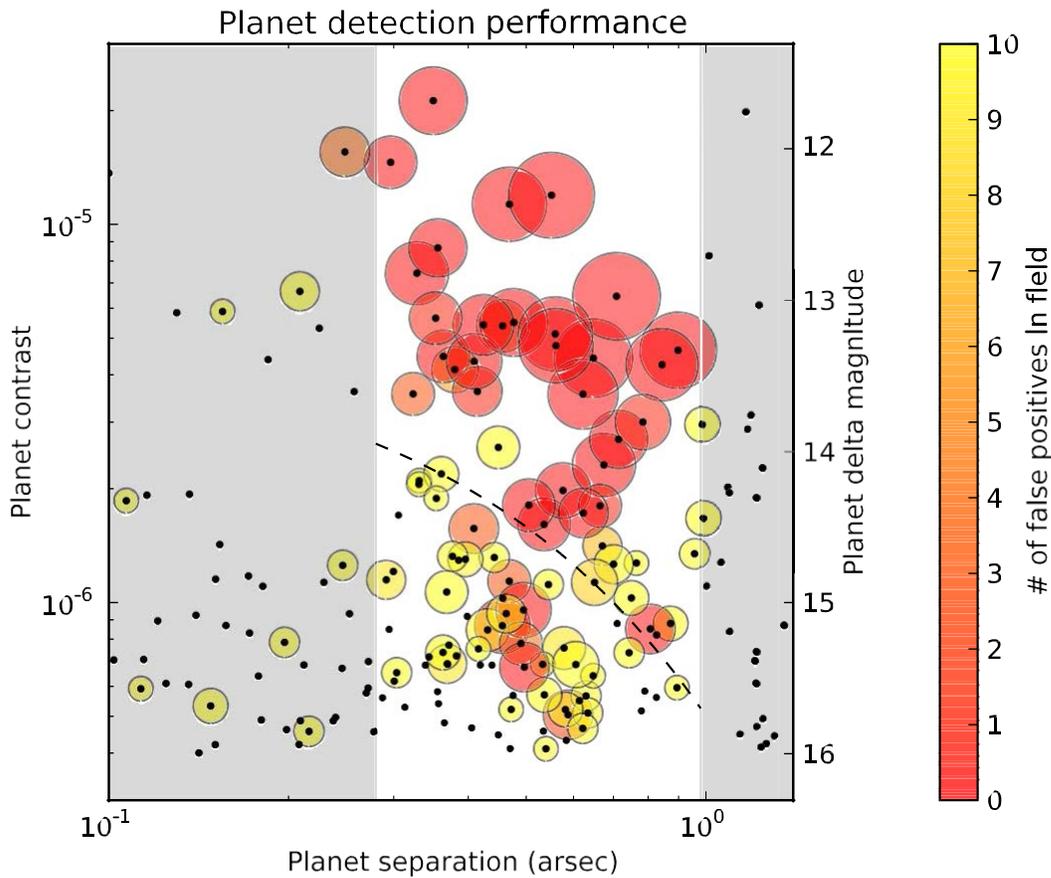


Figure 2:

The planet separation from star is plotted along the x-axis. The planet contrast is plotted on the y axis, expressed as powers of 10. A magnitude 20 planet in the challenge data has a contrast of  $10^{-6}$ . Each black dot is a planet that is hidden in the data. Each colored circle represents a true detection. Bare black dots are either misses, or lie in areas that were not searched (grey regions). The size of the circle indicates the maximum threshold in the LOCI image at which the planet was detected: bright planets at the bottom have higher thresholds. The color of the circle indicates the number of false positive detections in the entire observation for that threshold. Dark red in this case is no false positives. Orange is 5 false positives also detected at that threshold. Yellow is 10 false positives, and aqua is 20 false positives, both of which are quite high levels.

also detected at that threshold. Yellow is 10 false positives, and aqua is 20 false positives, both of which are quite high levels.

There are several ways to investigate this further. Particularly interesting would be:

1. Searching closer in than 0.29 arcsec (our default IWA) and farther out than 0.98 arcsec (our default OWA, which is a circle inscribed in the "dark hole").
2. Improving our sensitivity limit to making true detections with fewer false alarms on fainter planets. This is roughly indicated by the dash diagonal line, which we want to move up by turning the orange and yellow circles above to dark red.
3. Presenting our results in an FROC formalism

4. Quantifying “Prob Detection” and “Prob False Alarm” as a function of planet location and brightness and chosen threshold.

## 4. FUTURE DIRECTIONS AND COLLABORATIONS

### 4.1 Derive Practical Methods for Computation of Covariance Matrices:

Barrett, Caucci, Furenlid

The activity we identify as Task # 1 is to develop a practical and robust methodology for implementing the calculation of the data covariance matrix. This activity includes computing covariance matrices from simulation or experimental data, their inverses, and the square root of the inverses that comprises the prewhitening filter. The major challenge for the activity just identified is that data structures are very large, hence attention is required to storage, usage of efficient low-dimensional representation, recursive updates, and other mathematical tools that we will develop as we progress in our research.

Knowledge of the data covariance matrix directly and indirectly leads to many benefits, such as the optimal linear observer (Hotelling observer), the prewhitening of the data to enhance the performance of other algorithms (such as the matched filter), the simple calculation of ROC/LROC/FROC curves that characterize the performance of the algorithm and the imaging system combination. By bringing statistical rigor to data analysis and enabling the quick calculation of scalar metrics of performance (detection SNR, AUC, AULROC, AUFROC), a tool is created for algorithm, acquisition strategy, and hardware optimization.

Our objectives are to develop a complete theory for the decomposition and manipulation of the data covariance matrices that can be practically realized in state-of-the-art computing hardware. We will take advantage of the latest developments in CPU/GPU parallel architectures, inherent symmetries in the covariances, sparsity, the matrix-inversion lemma, the Shur complement, tools for simultaneous diagonalization, and recursive methods.

### 4.2 Gain Experience Analyzing Existing Data: Gladysz

The superiority of the Hotelling approach, in terms of the statistically derived quantities like the area under the receiver operating curve, has been demonstrated in several papers on simulated data. The assumptions that went into these simulated tests like the availability of noise-free training data might be optimistic, although there is a push to collect and process auxiliary data in future high-contrast instruments, that would translate, in theory, to the situation where the various components of the covariance matrix can be tracked separately facilitating the subsequent inversion process. The other approach is to try to extract the covariance matrix directly from the noisy data, available, for example, from the archive. Unfortunately there are serious obstacles to this approach, most notably strong requirements on the number of the available noisy images (or sequences in the case of the spatio-temporal Hotelling observer). The number of additional data for the computation of a (non-singular) covariance matrix must be greater than the order of the covariance matrix. This implies that, although this approach is conceptually simple (brute-force) and requires only noisy science-camera images, it relies on millions to tens of millions of image sequences to process one image sequence of moderate size (64 by 64 pixels by 10-50 frames).

We hope to use this second approach, namely generation of the covariance matrix directly from noisy science-camera images. Some strides will have to be made to reduce the problem, for example by stripping off the diagonal (readout + Poisson) component of the matrix and thus separating the speckle part. If this is possible then the matrix inversion lemma could then be employed. Perhaps the reduction of dimensionality of the matrix could be possible by exploiting the behaviour of temporal correlations seen in real data.

We will apply the spatio-temporal Hotelling observer to real data from the 3m Shane Telescope at the Lick Observatory. The data consists of 100 objects, 10,000 short-exposure images per object, which might be enough to construct the covariance matrix using the brute-force approach. We would use the signal-known-exactly paradigm first. Four algorithms would be used: spatial Hotelling observer on long exposures, spatio-temporal Hotelling observer on the sequences of short exposures, statistical speckle discrimination on sequences of frames, and a combination of statistical speckle discrimination (as a pre-processing, object-amplifying step), and a subsequent application of the spatial Hotelling observer. The ROC curves would be constructed for each of these methods. The final plot would show these curves. We hope to show that the exploitation of the temporal information is

beneficial, and that advanced image processing methods, like the spatio-temporal Hotelling observer can actually be used on real data and produce a benefit in terms of detectability. The crucial question is: given the error in the sample-based estimate of the covariance matrix do we still gain by employing the Hotelling observer? At what level of error (or sample size) does the benefit, in terms of AUC, disappear?

### 4.3 Compare how different algorithms make use of information in the data: Mouillet and Mugnier

High contrast multi-lambda imaging data is very rich and in particular includes significant redundancy whereas the number of parameters of interest is relatively small. The optimal detection method will need to fully benefit from data covariances (through e.g. hotelling approaches, possibly including pre-whitening). Various approaches already exists or are proposed to reduce such high contrast data. They have been developed in different context and communities, and use part of the covariance of the data more or less explicitly.

The data challenge will make possible a comparison of the performance of these various approaches in common and well-controlled conditions (with common simulated data sets).

We propose to discuss these various approaches based on their principle and based on their performance in the data challenge, in terms of their of the various redundancies and a priori known properties of the data. Practically, we can list what we think is the most useful information that helps discriminating a planet from speckles (information on planet itself and on speckle halo properties), and identify whether and how each approach uses it.

This discussion will also help us propose evolutions of the existing approaches, possibly

1. understanding their behaviour, possibly different depending on the observing conditions (where the dominant limitation may vary and the interest of various information will weigh differently).
2. merging the best properties of each of them.

In particular, each method may estimate differently the prior information needed to do the detection, such as the different pieces of the data covariance matrix, which may be either calibrated prior to the observation, or from the telemetry if available, or most often estimated from the data.

### 4.4 Derive more accurate models of the point-spread functions: Devaney

Together with Eric Thiébaud (Obs. de Lyon) I have developed code which will fit spline functions to multi-wavelength images in order to build an estimate of the PSF (including residual speckle). The fitting is inherently local and so could be equivalent to a multi-spectral LOCI. The difference is that a model of the PSF is built, rather than using a combination of the actual data. We believe that using the model may bring a gain in contrast ratio. (Perhaps it could be used to provide the noise-free AO covariance term?)

The current approach to planet detection is simply to subtract the model from the data at any single (SDI) or at all wavelengths and look for planets in the residuals. For the data challenge I could use the model PSF in a simple diagonal covariance Hotelling (as in Daniel Burkes thesis). Eric Thiébaud is also interested in applying the DARWIN-type approach i.e. fitting splines and planets to the data. A simple (sub-optimal) approach to estimating the planet spectrum would be to follow up detection with re-fitting the data to provide the planet spectrum plus a new spline model of the PSF.

Questions: will it work on coronagraphic data (more complex wavelength scaling)? How close to the core can it work ? How long will it take ?

#### 4.5 Make Optimal Use of Auxiliary Data Measurable by GPI and SPHERE: Savransky

The actions on this task flow directly from the development of the Hotelling observer formalism for the specific problem of ground-based planet-finding (Task #1) and the experience gained from applying these tools to real and simulated data (Tasks #2 and #3). An initial priority must be to agree upon conventions between the two (or more) groups, especially on data formatting (and the inclusion of synchronized telemetry) and on a detection criteria (i.e., using ROC/LROC as the primary metric).

The projects will independently create lists of all auxiliary data that they currently plan to collect, and identify any telemetry that should be added (in the case of GPI, this includes either raw or processed outputs from the CAL system). The teams will then exchange lists and thereby begin to populate a global list of useful auxiliary data for planet-finding that will be shared with the community in general.

Along with the identification of all auxiliary data, we must decide whether each particular telemetry stream can be stored entirely, or whether we will only store a statistical description in the form of a recursively updated first and second moment estimate. The latter approach may also include using the Schur complement to calculate the conditional statistics of the science and auxiliary data.

During the first year of science operations, the teams will build up experience on applying the techniques and algorithms from Task #1 on real science data. It would be useful for the teams to communicate during this period to compare experiences and any refinements they have made to their methodologies.

At the same time, the groups must begin to prepare for the inclusion of the information provided by the Task #1 tools into their campaign scheduling; for example, incorporating AUROC statistics into dynamic campaign rescheduling after the initial  $K_G$  matrices have been constructed from initial science data. As part of this preparation, it may be useful to construct simple campaign simulations to evaluate the effects of this type of information.

Finally, Laurent Pueyo has a simulated CAL data set which can be used to evaluate maximum likelihood wavefront reconstruction. While the data is GPI CAL specific, it can act as a general test for the efficacy of an MLE in wavefront reconstruction. Furthermore, this data can be folded into an extension of the data challenge as an auxiliary data set.

### 5. CONCLUSION

At the time of writing the exoplanet imaging challenge is still in progress. The simulated datasets have been prepared and distributed for the first round of analysis in the exoplanet imaging challenge. The challenge itself will take place in several phases, with each phase introducing added complexity to the data. A preview of results is presented here. The final results will be published in the astronomical literature. Work with this group is actively being planned through 2012–2013.

Through support from the National Academies Keck Futures Initiative we have initiated several collaborations. This is a field that is still in its infancy. The formal approach that we advocate promises to yield a more profound understanding of correlations in noise in exoplanet data and how they may be used to our advantage. This work promises to help not only ground-based astronomy, but by extension space-based astronomy as well.

### ACKNOWLEDGMENTS

This work was undertaken with support of a seed grant from the National Academies Keck Futures Initiative.

Christian Marois (HIA, NRC) provided the K-band data for Keck NIRC2; Dimitri Mawet (ESO) provided the K-band data for VLT-NACO; The GPI data were downloaded from planetimager.org, corresponding to predictions by McBride;<sup>11</sup> Dino Mesa (INAF) provided the data for the Integral Field Spectrograph of VLT-SPHERE;<sup>22</sup> John Krist (JPL/Caltech) provided the JWST NIRC2 and HST/ACS curves; and Bruce Macintosh (UC Santa Cruz) provided the TMT PFI data.

This work was partly conducted at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. Reference in this paper to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute



Figure 3: Participants at the Exoplanet Imaging Workshop. Left to right: Harrison Barrett, David Mouillet, Szymon Gladysz, Luca Caucci, Lars Furenlid, Laurent Mugnier, Dmitry Savransky, Lisa Poyneer, Laurent Pueyo, Peter Lawson, Dimitri Mawet, Nicholas Devaney, and Richard Frazin. Not shown: Olivier Guyon, John Krist, Christian Marois, and Remi Soummer.

or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

Copyright 2012. All rights reserved.

## REFERENCES

- [1] Barrett, H. H., Myers, K. J., Devaney, N., Dainty, J. C., and Caucci, L., "Task performance in astronomical adaptive optics," in [Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series], Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series 6272 (July 2006).
- [2] Barrett, H. H., Myers, K. J., Devaney, N., and Dainty, C., "Objective assessment of image quality. IV. Application to adaptive optics," *Journal of the Optical Society of America A* 23, 3080–3105 (Dec. 2006).
- [3] Caucci, L., Barrett, H. H., Devaney, N., and Rodríguez, J. J., "Application of the Hotelling and ideal observers to detection and localization of exoplanets," *Journal of the Optical Society of America A* 24, 13 (2007).
- [4] Burke, D., Devaney, N., Gladysz, S., Barrett, H. H., Whitaker, M. K., and Caucci, L., "Optimal linear estimation of binary star parameters," in [Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series], Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series 7015 (July 2008).
- [5] Caucci, L., Barrett, H. H., and Rodríguez, J. J., "Spatio-temporal Hotelling observer for signal detection from image sequences," *Optics Express* 17, 10946 (June 2009).

- [6] Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., and Artigau, É., “A New Algorithm for Point-Spread Function Subtraction in High-Contrast Imaging: A Demonstration with Angular Differential Imaging,” *Astrophys. J.* 660, 770–780 (May 2007).
- [7] Marois, C., Macintosh, B., and Vêran, J.-P., “Exoplanet imaging with LOCI processing: photometry and astrometry with the new SOSIE pipeline,” in [Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series], Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series 7736 (July 2010).
- [8] Galicher, R., Marois, C., Macintosh, B., Barman, T., and Konopacky, Q., “M-band Imaging of the HR 8799 Planetary System Using an Innovative LOCI-based Background Subtraction Technique,” *Astrophys. J.* 739, L41 (Oct. 2011).
- [9] Pueyo, L., Crepp, J. R., Vasisht, G., Brenner, D., Oppenheimer, B. R., Zimmerman, N., Hinkley, S., Parry, I., Beichman, C., Hillenbrand, L., Roberts, L. C., Dekany, R., Shao, M., Burruss, R., Bouchez, A., Roberts, J., and Soummer, R., “Application of a Damped Locally Optimized Combination of Images Method to the Spectral Characterization of Faint Companions Using an Integral Field Spectrograph,” *Astrophys. J. Supp. Ser.* 199, 6 (Mar. 2012).
- [10] Crepp, J. R., Pueyo, L., Brenner, D., Oppenheimer, B. R., Zimmerman, N., Hinkley, S., Parry, I., King, D., Vasisht, G., Beichman, C., Hillenbrand, L., Dekany, R., Shao, M., Burruss, R., Roberts, L. C., Bouchez, A., Roberts, J., and Soummer, R., “Speckle Suppression with the Project 1640 Integral Field Spectrograph,” *Astrophys. J.* 729, 132 (Mar. 2011).
- [11] McBride, J., Graham, J. R., Macintosh, B., Beckwith, S. V. W., Marois, C., Poyneer, L. A., and Wiktorowicz, S. J., “Experimental Design for the Gemini Planet Imager,” *Pub. Astron. Soc. Pac.* 123, 692–708 (June 2011).
- [12] Marois, C., Macintosh, B., Barman, T., Zuckerman, B., Song, I., Patience, J., Lafrenière, D., and Doyon, R., “Direct Imaging of Multiple Planets Orbiting the Star HR 8799,” *Science* 322, 1348– (Nov. 2008).
- [13] Marois, C., Zuckerman, B., Konopacky, Q. M., Macintosh, B., and Barman, T., “Images of a fourth planet orbiting HR 8799,” *Nature* 468, 1080–1083 (Dec. 2010).
- [14] Lagrange, A.-M., Bonnefoy, M., Chauvin, G., Apai, D., Ehrenreich, D., Boccaletti, A., Gratadour, D., Rouan, D., Mouillet, D., Lacour, S., and Kasper, M., “A Giant Planet Imaged in the Disk of the Young Star  $\beta$  Pictoris,” *Science* 329, 57– (July 2010).
- [15] Bonnefoy, M., Lagrange, A.-M., Boccaletti, A., Chauvin, G., Apai, D., Allard, F., Ehrenreich, D., Girard, J. H. V., Mouillet, D., Rouan, D., Gratadour, D., and Kasper, M., “High angular resolution detection of  $\beta$  Pictoris b at 2.18  $\mu\text{m}$ ,” *Astron. Astrophys.* 528, L15 (Apr. 2011).
- [16] Patience, J., King, R. R., De Rosa, R. J., Vigan, A., Witte, S., Rice, E., Helling, C., and Hauschildt, P., “Spectroscopy across the brown dwarf/planetary mass boundary. I. Near-infrared JHK spectra,” *Astron. Astrophys.* 540, A85 (Apr. 2012).
- [17] Maire, J., Perrin, M. D., Doyon, R., Artigau, E., Dunn, J., Gavel, D. T., Graham, J. R., Lafrenière, D., Larkin, J. E., Lavigne, J.-F., Macintosh, B. A., Marois, C., Oppenheimer, B., Palmer, D. W., Poyneer, L. A., Thibault, S., and Vêran, J.-P., “Data reduction pipeline for the Gemini Planet Imager,” in [Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series], Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series 7735 (July 2010).
- [18] Poyneer, L. A. and Macintosh, B. A., “Optimal fourier control performance and speckle behavior in high-contrast imaging with adaptive optics,” *Opt. Exp.* 14, 7499–7514 (2006).
- [19] Wallace, J. K., Angione, J., Bartos, R., Best, P., Burruss, R., Fregoso, F., Levine, B. M., Nemati, B., Shao, M., and Shelton, C., “Post-coronagraph wavefront sensor for Gemini Planet Imager,” in [Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series], Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series 7015 (July 2008).
- [20] Marois, C., Macintosh, B., Soummer, R., Poyneer, L., and Bauman, B., “An end-to-end polychromatic Fresnel propagation model of GPI,” in [Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series], Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series 7015 (July 2008).
- [21] de Vries, W. H., Olivier, S. S., Asztalos, S. J., Rosenberg, L. J., and Baker, K. L., “Image Ellipticity from Atmospheric Aberrations,” *Astrophys. J.* 662, 744–749 (June 2007).

- [22] Mesa, D., Gratton, R., Berton, A., Antichi, J., Verinaud, C., Boccaletti, A., Kasper, M., Claudi, R. U., Desidera, S., Giro, E., Beuzit, J.-L., Dohlen, K., Feldt, M., Mouillet, D., Chauvin, G., and Vigan, A., "Simulation of planet detection with the SPHERE integral field spectrograph," *Astron. Astrophys.* 529, A131 (May 2011).