

APPLICATION OF A DAMPED LOCALLY OPTIMIZED COMBINATION OF IMAGES METHOD TO THE SPECTRAL CHARACTERIZATION OF FAINT COMPANIONS USING AN INTEGRAL FIELD SPECTROGRAPH

LAURENT PUEYO^{1,2,8}, JUSTIN R. CREPP³, GAUTAM VASISHT⁴, DOUGLAS BRENNER⁵, BEN R. OPPENHEIMER⁵, NEIL ZIMMERMAN⁵,
SASHA HINKLEY³, IAN PARRY⁶, CHARLES BEICHMAN⁷, LYNNE HILLENBRAND³, LEWIS C. ROBERTS, JR.⁴, RICHARD DEKANY³,
MIKE SHAO⁴, RICK BURRUSS⁴, ANTONIN BOUCHEZ³, JENNY ROBERTS³, AND RÉMI SOUMMER²

¹ Department of Physics and Astronomy, Johns Hopkins University, 366 Bloomberg Center, 3400 N. Charles Street, Baltimore, MD 21218 USA; lap@pha.jhu.edu

² Space Telescope Science Institute, 3700 San Marin Drive, Baltimore, MD 21218, USA

³ California Institute of Technology, 1200 E. California Blvd., Pasadena, CA 91125, USA

⁴ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA

⁵ American Museum of Natural History, Central Park West at 79th Street, New York, NY 10024, USA

⁶ Institute of Astronomy, University of Cambridge, Madingley Rd., Cambridge, CB3 0HA, UK

⁷ NASA Exoplanet Science Institute, 770 S. Wilson Avenue, Pasadena, CA 91225, USA

Received 2011 April 22; accepted 2011 November 3; published 2012 February 22

ABSTRACT

High-contrast imaging instruments are now being equipped with integral field spectrographs (IFSs) to facilitate the detection and characterization of faint substellar companions. Algorithms currently envisioned to handle IFS data, such as the Locally Optimized Combination of Images (LOCI) algorithm, rely on aggressive point-spread function (PSF) subtraction, which is ideal for initially identifying companions but results in significantly biased photometry and spectroscopy owing to unwanted mixing with residual starlight. This spectrophotometric issue is further complicated by the fact that algorithmic color response is a function of the companion's spectrum, making it difficult to calibrate the effects of the reduction without using iterations involving a series of injected synthetic companions. In this paper, we introduce a new PSF calibration method, which we call “damped LOCI,” that seeks to alleviate these concerns. By modifying the cost function that determines the weighting coefficients used to construct PSF reference images, and also forcing those coefficients to be positive, it is possible to extract companion spectra with a precision that is set by calibration of the instrument response and transmission of the atmosphere, and not by post-processing. We demonstrate the utility of this approach using on-sky data obtained with the Project 1640 IFS at Palomar. Damped LOCI does not require any iterations on the underlying spectral type of the companion, nor does it rely on priors involving the chromatic and statistical properties of speckles. It is a general technique that can readily be applied to other current and planned instruments that employ IFSs.

Key words: instrumentation: adaptive optics – techniques: high angular resolution – techniques: spectroscopic

Online-only material: color figures

1. INTRODUCTION

Current instruments dedicated to the imaging of exoplanets have a contrast floor that is set by the imperfections on the surfaces of the optical train. In the case of a space observatory, as shown, for instance, by Krist et al. (2005) and Kalas et al. (2008), when the optical surface corrugations are stable in time, the instrument response can be calibrated and subtracted. However, in the ground-based regime, the point-spread function (PSF) varies over timescales that range from minutes to hours (e.g., Hinkley et al. 2007; Fitzgerald & Graham 2006). Since these so-called quasi-static speckles are the limiting factor of ground- and spaced-based observations with regard to faint companion detectability, their properties have been studied extensively over the past few years. It was found that their theoretical properties (Perrin et al. 2003; Aime & Soummer 2004; Soummer & Ferrari 2007; Soummer et al. 2007a) match the statistics observed on-sky (Fitzgerald & Graham 2006; Hinkley et al. 2007).

Several observing techniques and data reduction algorithms have been devised over the past few years in order to optimize the sensitivity of current imaging instruments to faint exoplanets. Observations are generally carried out by acquiring a series of images in which the observer deliberately chooses to introduce some diversity in the imaging method. The purpose of this

diversity is to force the relative motion of the speckles with respect to the image of the sky to vary from frame to frame. This is the case in angular differential imaging (ADI; Marois et al. 2006; Leconte et al. 2010), spectral differential imaging (Racine et al. 1999; Marois et al. 2004; Biller et al. 2007), chromatic differential imaging (Sparks & Ford 2002), and coherent imaging (Guyon et al. 2010). Post-processing techniques involve a manipulation of this series of images whose purpose is to remove the artifacts due to speckles and leave intact the signal of a companion. The powerful combination of ADI data processed using a Locally Optimized Combination of Images (LOCI) PSF subtraction algorithm has been used (Marois et al. 2008) to obtain one of the first images of a planetary system. The LOCI PSF subtraction technique was devised by Lafrenière et al. (2007) and is quickly becoming a standard in exoplanet searches (Marois et al. 2008; Thalmann et al. 2009; Metchev et al. 2009; Crepp et al. 2011). Other techniques currently discussed in the literature involve a priori models of the PSF and the telescope (Mugnier et al. 2009; Burke & Devaney 2010) or spectral deconvolution based on priors about the underlying spectrum of the companion (Mesa et al. 2011). On the contrary, LOCI only requires minimal priors about the behavior of the instrument and builds a composite reference PSF solely on the data obtained during observations.

In a previous paper we reported a recent implementation of LOCI that is applied to process data from an integral

⁸ NASA Sagan Fellow.

field spectrograph (IFS; Crepp et al. 2011). Our results were obtained using Project 1640. The full system is composed of the Palomar Adaptive Optics (PALAO) followed by a stellar coronagraph (Soummer 2005) and the science IFS. Further details about the instrument are found in Hinkley et al. (2011). The raw data from the infrared detector are calibrated and re-arranged in order to produce a data cube of images at 23 different wavelengths across the J and H bands ($R \sim 45$). The spatial and spectral data extraction procedures are detailed in Zimmerman et al. (2011). The P1640 configuration prevents the field rotation that is necessary to carry out ADI observations but allows both chromatic diversity. Using an IFS behind a high-contrast starlight suppression system is key to the two main science objectives of P1640: (1) detection of faint companions around nearby stars and (2) in situ low-resolution spectral characterization of any companions discovered.

The detection of faint point sources is facilitated by the chromatic diversity of the IFS and allows the observer to apply advanced static speckle calibration methods. Indeed, if all the errors are in an instrument pupil plane, diffractive optics predicts that the speckle pattern will stretch linearly with wavelength, while the scale of the astronomical image is independent of wavelength. In an IFS, the series of measurements of the quasi-static speckle pattern is simultaneous. This alleviates the problem of corrugations evolving temporally between sequential observations. On the P1640 camera the speckles located at $1''$ are at the 5×10^{-3} level and de-correlate in intensity by roughly 1% per minute. This implies that purely sequential observations, separated by 5 minutes in time, can at best reach a contrast of 2.5×10^{-4} after subtraction. Using an IFS alleviates this problem as it provides reference PSFs that are contemporaneous realizations of the speckle pattern present in the science exposure. Each wavelength slice can be interpolated in order to produce rescaled data cubes in which the speckle pattern's scale remains constant. In this rescaled space the companion moves radially. If the variations of the speckle's brightness as a function of wavelength are known a priori, then these rescaled images can be subtracted from one to another to detect a faint companion. This approach was first discussed in Sparks & Ford (2002) and has been successfully applied to both simulated and on-sky IFS data (Vigan et al. 2008; Thatte et al. 2007).

In Crepp et al. (2011), we presented an algorithm that combines the chromatic differential imaging in Sparks & Ford (2002) with the LOCI PSF subtraction approach. While LOCI was first presented in Lafrenière et al. (2007) as a method to reduce ADI images, P1640 does not take advantage of field rotation. Instead, the baseline P1640 reduction presented in Crepp et al. (2011) relies on the IFS's radial chromatic stretch of the PSF and on temporal diversity. It integrates both of them in an LOCI-type reduction scheme. We showed that this quasi-static speckle calibration method provided a gain greater than an order of magnitude in the contrast between raw and processed images and an ultimate detection level $\simeq 1 \times 10^{-5}$ at $1''$ in a 20 minute H -band exposure on a third magnitude star. Further improvements in contrast, while maintaining reasonable exposure times, require coherent wave front calibrations, such as the ones presented in Bordé & Traub (2006), Guyon et al. (2010), Give'on et al. (2007), and Pueyo et al. (2009) for the case of space-based observations and Wallace et al. (2010) and Sauvage et al. (2007) for the case of ground-based telescopes.

A second advantage of the P1640 IFS is that, once a companion is detected by means of the aforementioned chromatic diversity, its spectral information is an immediate by-product. Over

the past 10 years, IFS have been used previously to characterize companions with favorable contrast ratio (e.g., McElwain et al. 2007; Patience et al. 2010). When the companion is fainter than the speckle floor, reductions such as the ones presented in Crepp et al. (2011) ought to be carried out to first detect the companion. However, blind applications of LOCI post-processing tailored for detections to spectral data cubes present the major caveat of altering the spectrophotometry and possibly introduce artificial features in the low resolution of the detected object. Recent results obtained using the OSIRIS IFS at the Keck observatory show a depletion of 60% of the companion flux that is constant over the spectral window (Bowler et al. 2010). In this example the depletion is gray and as such has little impact on the scientific conclusions of the observations. However, P1640 seeks faint objects at close angular separations whose signal is buried within the bright speckle field, thus requiring more aggressive PSF subtraction. Moreover, P1640 entirely relies on chromatic diversity, which can potentially lead to a poor conditioning of the inverse problem associated with LOCI (see below). In practice, when testing our reduction on synthetic objects injected within the control radius of the AO system, we observe that the post-LOCI retrieved spectral information is considerably biased with a non-gray algorithmic response. This biasing of the spectral information is entirely due to the PSF subtraction algorithm and can be an order of magnitude larger than the next most important uncertainty, owing to the telluric calibration described in Zimmerman et al. (2011). It stems from two independent phenomena: a bias caused by partial fitting of the companion flux with the starlight in the reference frames and a spectral cross-talk between actual companion flux and neighboring channels. The combination of these errors is very often the dominant source of uncertainties in spectroscopic measurements carried out using post-LOCI-subtracted IFS data cubes.

In this paper we introduce a new method that significantly lowers the error in the measured spectral energy distribution (SED) due to the reduction algorithm. We first introduce a modified version of the LOCI algorithm that alleviates the problem of flux depletion of faint companions that is common to most locally optimized PSF subtraction algorithms. The ‘‘damped LOCI’’ algorithm introduced in this communication seeks to maximize the residual flux of a putative companion while minimizing the noise due to quasi-static speckles, whereas standard LOCI implementations only focus on the latter objective. This is achieved by re-formulating the least-squares problem associated with finding a locally optimal composite reference PSF: we augment the cost function associated with the determination of the composite PSF coefficients with a penalty term that scales with the flux of a potential companion. ‘‘Damped LOCI’’ thus solves the following optimization problem: ‘‘*minimize the least-squares fit to the local speckle field while maximizing the least-squares fit in the area where a known companion is present.*’’

The formalism and the notations underlying LOCI and ‘‘damped LOCI’’ are detailed in Appendix A. Section 2 focuses on the specifics and the implementation of each algorithm. Both approaches yield comparable performances in terms of band-averaged faint companions' signal-to-noise ratio (S/N). Appendix B discusses the second-order impact of each method on the detectability problem, and in particular the advantages of ‘‘damped LOCI’’ when the inverse problem associated with the PSF reconstruction is ill conditioned. In Section 3 we discuss the core matter of this paper, namely, the measurement of the spectrum of faint companions. When companions have

already been detected, we show that a series of binary PSF weightings, or masking, can be used in conjunction with classical LOCI methods to retrieve the spectrum of relatively bright objects. We illustrate, however, that the performance of such methods degrades rapidly with increasing contrast and show that using masking in conjunction with “damped LOCI” decouples spectral precision and contrast. We show that this method yields in all cases better absolute and relative spectrophotometric accuracy when compared with classical LOCI reductions. In particular, for objects as bright as the residual speckles, we obtain estimates of the SED of the order of the telluric calibration uncertainties of the spectrograph. For faint objects “damped LOCI” yields SED uncertainties consistent with higher S/N (less “photon starved”) state-of-the-art band-averaged photometric estimates (Marois et al. 2008; Thalmann et al. 2009; Metchev et al. 2009). As the architecture of future exoplanet imagers and P1640 are similar (Macintosh et al. 2006; Beuzit et al. 2008), the results presented here are broadly relevant to forthcoming exoplanet imaging endeavors.

2. STANDARD LOCI AND DAMPED LOCI IMPLEMENTATIONS

2.1. Standard LOCI

The LOCI algorithm involves reconstructing a reference PSF from IFS data. It is based on an inverse problem that seeks to minimize the residuals between a given astronomical image, the target image to be reduced, and a calibration image. The algorithm, as introduced by Lafrenière et al. (2007), seeks to find an optimal combination of images, within an ensemble of reference frames, that locally minimizes the least-squares residual between the target frame and the composite reference, in a large optimization zone of the image, \mathcal{O} zone. This weighted reference is then subtracted from the reduced image for only a small subset of pixels, the subtraction (S) zone. Crepp et al. (2011) presented how such an approach can be utilized to reduce IFS data. Appendix A details the specifics of the frame selection and the notations underlying the present manuscript. In this section we focus on the kernel of the LOCI algorithm, namely, the inverse problem associated with the determination of the optimal reference weights.

The search for the coefficients of the composite reference is based on a least-squares cost function integrated over the zone \mathcal{O} :

$$\min_{\{c_k\}} \left\{ \int_{\mathcal{O}} dx dy w(x, y) \left[T(x, y) - \sum_k^{N_{\mathcal{R}}} c_k R_k(x, y) \right]^2 \right\} \\ \Leftrightarrow \min_{\{c_k\}} \left\{ \left\| T - \sum_k^{N_{\mathcal{R}}} c_k R_k \right\|_{w, \mathcal{O}} \right\}, \quad (1)$$

where T stands for the target image, $\{R_k\}$ is the ensemble of reference images, $\{c_k\}$ is the ensemble of LOCI coefficients, $w(x, y)$ is a weighting function that can be used to mask certain pixels, $N_{\mathcal{R}}$ is the number of reference frames in \mathcal{R} , and $\|\cdot\|_{w, \mathcal{O}}$ denotes the L^2 norm, weighted by w , over the optimization zone. In matrix notation we rewrite Equation (1) as

$$\left\| T - \sum_k^{N_{\mathcal{R}}} c_k R_k \right\|_{w, \mathcal{O}} = c^T \mathbf{M}_{\mathcal{O}}^{(R, R)} c - 2\mathbf{V}_{\mathcal{O}}^{(R, T)} c + \mathbf{S}_{\mathcal{O}}^{(T, T)}, \quad (2)$$

where c is an $N_{\mathcal{R}}$ -dimensional vector, $\mathbf{M}_{\mathcal{O}}^{(R, R)}$ is an $N_{\mathcal{R}} \times N_{\mathcal{R}}$ matrix, $\mathbf{V}_{\mathcal{O}}^{(R, T)}$ is an $N_{\mathcal{R}}$ -dimensional vector, and $\mathbf{S}_{\mathcal{O}}^{(T, T)}$ is a scalar. They are defined by

$$c = [c_1, c_2, \dots, c_k, \dots, c_{N_{\mathcal{R}}}]^T \quad (3)$$

$$\mathbf{M}_{\mathcal{O}}^{(R, R)}[p, q] = \int_{\mathcal{O}} dx dy w(x, y) R_p(x, y) R_q(x, y) \quad (4)$$

$$\mathbf{V}_{\mathcal{O}}^{(R, T)}[p] = \int_{\mathcal{O}} dx dy w(x, y) R_p(x, y) T(x, y) \quad (5)$$

$$\mathbf{S}_{\mathcal{O}}^{(T, T)} = \int_{\mathcal{O}} dx dy w(x, y) T(x, y) T(x, y). \quad (6)$$

Using these notations, the quadratic residual minimization problem associated with standard LOCI implementations can be solved by taking the derivative of Equation (2) with respect to c ,

$$\min_{\{c_k\}} \left\{ \left\| T - \sum_k c_k R_k \right\|_{w, \mathcal{O}} \right\} \Leftrightarrow \mathbf{M}_{\mathcal{O}}^{(R, R)} c - \mathbf{V}_{\mathcal{O}}^{(R, T)} = 0. \quad (7)$$

If the ensemble of references, $\{R_k \in \mathcal{R}\}$, forms a complete basis set of the \mathcal{O} zone, then the matrix $\mathbf{M}_{\mathcal{O}}^{(R, R)}$ is invertible. If the speckle diversity is only temporal, and if each frame corresponds to an independent realization of the speckle pattern, then the inverse problem in Equation (7) is well conditioned when the number of PSF cores in the optimization zone matches the number of reference images, $N_A = N_{\mathcal{R}}$. In this case the optimal weights are defined as

$$c_{\text{LOCI}} = (\mathbf{M}_{\mathcal{O}}^{(R, R)})^{-1} \mathbf{V}_{\mathcal{O}}^{(R, T)}. \quad (8)$$

However, when using IFS data, the color slices within each exposure correspond to the same realization of the speckle field, albeit seen at different wavelengths. As a consequence, they are highly correlated and this results in poorer conditioning of $\mathbf{M}_{\mathcal{O}}^{(R, R)}$. This conditioning is a measure of how well N_A , the number of PSF cores, matches the number of linearly independent reference frames in the zone \mathcal{O} .

Carrying out a selection of the most relevant slices in \mathcal{R} , so that the number of independent references matches N_A ought to be done locally and iteratively (Marois et al. 2010). It represents a considerable computational overhead, and in practice, LOCI reductions are carried out with N_A constant over large portions of the image. Then, in the zones for which the problem is ill posed, generic pseudo-inverse routines or eigenvalue truncations provide a robust alternative, at the cost of losing track of the physics in the process. Indeed, such inversion methods lead to the subtraction of the least correlated modes in the reference images. If the flux of the companion in the target frame can be described using some of these independent reference modes, then some of its flux is subtracted by the composite reference. The photometric information is then biased. Choosing an optimization zone much larger than the subtraction zone, a larger N_A , mitigates this issue. It increases the contribution of pixels that do not see any companion flux and reduces the likelihood of fitting modes associated with the companion with one of the reference modes. However, this occurs at the expense

of losing some of the local characteristics of the suppression, e.g., larger optimization zones. Such a procedure can be efficient outside of the control radius of the AO system, where the speckle halo is relatively uniform (Bowler et al. 2010). However, when the speckles are bright and structured, optimizing the LOCI parameters in order to obtain the best spectral retrieval quickly becomes an intractable problem. In other words, when the minimization is ill posed, pseudo-inverse routines aim at minimizing the speckle noise in the reduced image without considerations of the flux of a potential companion. Whereas such an approach is adequate for faint companion detection, it can be problematic when quantifying the spectrophotometry of the discovered point sources. We now explore an alternate reduction approach which resolves these issues.

2.2. Flux Preserving Locally Optimized Speckle Suppression: Damped LOCI

“Damped LOCI” relies on the same zonal approach as above. The modification we introduce resides in posing differently the inverse problem that defines the modal coefficients of the composite reference. Instead of solely minimizing the residual least-squares fit between target frame and composite reference in the optimization zone, we constrain the possible combinations of coefficients over which the solution may be sought. We first impose positivity, $c_k \geq 0$ for all k , in order to avoid the pathological case where two adjacent coefficients are large in absolute value but of different signs. Such a situation would lead to high spectral cross-talk from the companion flux at neighboring wavelengths. While this condition might appear too stringent, we show below that ensuring the positivity of all the reference coefficients allows us to introduce concisely in the minimization problem a penalty term that scales with opposite of the flux of a companion. Because we are interested in devising a reduction method that does not modify the spectrophotometry of a detected companion, the quality of the fit in zone S is of critical importance. Ideally, even if the flux from a companion is present in zone S , we want the algorithm to only fit the stellar contribution in that zone, leaving the companion contribution unperturbed. *In essence “damped LOCI” seeks composite references that are the result of a competition between minimizing the fit residual in zone \mathcal{O} and maximizing it in zone S , under the constraint that all the composite coefficients are positive.* Since zone \mathcal{O} is much larger than S , the minimization of its residual fit should prevail, and as a consequence the speckles will be subtracted in a way that preserves more of the companion flux than in classical LOCI implementations.

Formally, the squared difference between T and the composite reference integrated over S can be written as

$$I_S^{(2)} = \|T - cR\|_{w,S} = c^T \mathbf{M}_S^{(R,R)} c - 2\mathbf{V}_S^{(R,T)} c + \mathbf{S}_S^{(T,T)}, \quad (9)$$

where we have replaced the \mathcal{O} subscript by the S subscript for the matrix and vectors. When there is flux from a companion in the S zone, we want this quantity to remain large while the residual in the zone \mathcal{O} is minimized. We can thus use an augmented cost function to write the optimization problem as follows:

$$\min_{\{c_k, \Lambda\}} \{J_0(c_k, \Lambda)\} \Leftrightarrow \min_{\{c_k, \Lambda\}} \left\{ c^T \mathbf{M}_O^{(R,R)} c - 2\mathbf{V}_O^{(R,T)} c + \mathbf{S}_O^{(T,T)} - \Lambda(c^T \mathbf{M}_S^{(R,R)} c - 2\mathbf{V}_S^{(R,T)} c + \mathbf{S}_S^{(T,T)}) \right\}, \quad (10)$$

where Λ is a multiplicative constant that allows us to include the flux preservation constraint in the LOCI minimization. The term $-\Lambda(c^T \mathbf{M}_S^{(R,R)} c - 2\mathbf{V}_S^{(R,T)} c + \mathbf{S}_S^{(T,T)})$ is called the flux preservation penalty term, and $-\Lambda$ is the Lagrange multiplier associated with this term. If one seeks to maximize the flux of a potential companion in the S zone by minimizing $J_0(c_k, \Lambda)$, then Λ ought to be positive. As a consequence, this cost function is not positive definite, making the practical implementation of this optimization difficult. Forcing positivity for the c_k coefficients leads us to ensure the positivity of the second term in Equation (10) and thus facilitates the search for the optimal coefficients.

To do so, we project the LOCI processed image, I^P , over all the modes present in the collection of references.

$$I^P = T - cR \quad (11)$$

$$\sum_{k=1}^{N_R} c_k \int_S w R_k I^P = \mathbf{V}_S^{(R,I^P)} c = \mathbf{V}_O^{(R,T)} c - c^T \mathbf{M}_S^{(R,R)} c, \quad (12)$$

where we have taken the inner product over the S zone of the composite reference with both sides of Equation (11). Replacing the quadratic term in Equation (9) by the quadratic term in Equation (12) yields

$$I_S^{(2)} = \mathbf{S}_S^{(T,T)} - \mathbf{V}_S^{(R,I^P)} c - \mathbf{V}_S^{(R,T)} c. \quad (13)$$

Since the condition $c_k \geq 0$ for all k is being enforced and the actual images only include positive counts, $\mathbf{S}_S^{(T,T)} > 0$ and $\mathbf{V}_S^{(R,T)} c > 0$. In order to devise an augmented cost function that consists of the sum of positive scalar quantities, we need to show that maximizing $I_S^{(2)}$ is equivalent to minimizing $\mathbf{V}_S^{(R,T)} c$. When there is no companion in the S zone, $|\mathbf{V}_S^{(R,I^P)} c| \ll |\mathbf{V}_S^{(R,T)} c|$, since the flux in the processed image is smaller than the flux in the target image. Even if $\mathbf{V}_S^{(R,I^P)} c$ happens to be negative, which means that the subtraction worked too well, the integrated square flux in the subtraction zone is dominated by the behavior of $\mathbf{V}_S^{(R,T)} c$. When there is a companion, $\mathbf{V}_S^{(R,I^P)} c > 0$ and has the same sign as $\mathbf{V}_S^{(R,T)} c$. Thus, the integrated squared residual in the subtraction zone behaves as the opposite of the positive scalar $\mathbf{V}_S^{(R,T)} c$. As a consequence, we modify the augmented cost function shown in Equation (10), in order to avoid the use of a negative Lagrange multiplier associated with the full quadratic form of the least-squares fit in the S zone. Instead, we choose to use $\mathbf{V}_S^{(R,T)}$ in conjunction with a positive Lagrange multiplier as a penalty term.

Since $\mathbf{V}_S^{(R,T)}$ scales as the opposite of the residual fit in the S zone, this modification re-formulates the problem of minimizing Equation (10) in a tractable way. It allows us to write the following positive definite augmented cost function:

$$J_1(c, \Lambda) = c^T \mathbf{M}_O^{(R,R)} c - 2\mathbf{V}_O^{(R,T)} c + \mathbf{S}_O^{(R,T)} + \Lambda \mathbf{V}_S^{(R,T)} c, \quad (14)$$

where Λ is the Lagrange multiplier associated with the quality of the least-squares fit in the S zone. By augmenting the cost function using a quantity that scales with the flux suppression of a potential companion, we introduce a damping term in the speckle suppression algorithm. This term forces the overall speckle reduction to compete with flux preservation in the S

zone. Having chosen a positive definite cost function allows us to use standard quadratic programming routines to solve the following inverse problem:

Find c^{DL} and Λ^{DL} that minimize $J(c^{\text{DL}}, \Lambda^{\text{DL}})$ under the constraint that $c_k^{\text{ML}} \geq 0$ for all k ,

where the subscript DL stands for ‘‘damped LOCI,’’ d-LOCI hereafter.

In practice, we implement this optimization problem by nesting a linear least-squares solver within a one-dimensional nonlinear search,

$$\Lambda^{\text{DL}} = \min_{\Lambda} J(c(\Lambda), \Lambda) \quad (15)$$

with

$$c(\Lambda) = (\mathbf{M}_O^{(R,R)})^\dagger (\mathbf{V}_O^{(R,T)} - \frac{\Lambda}{2} \mathbf{V}_S^{(R,T)}), \quad (16)$$

where the superscript \dagger stands for the linear least-squares minimization under the constraint that all coefficients are positive. This algorithm will find the optimal weight between speckle and signal suppression and thus should behave better than regular LOCI with respect to the photometry and spectral features of a potential companion. Note that the damping in d-LOCI occurs in two separate instances: the positivity constraint on the coefficients and the conservation of the companion flux in the \mathcal{S} zone. We used a set of synthetic companions to assess the relative performances of both LOCI and d-LOCI with respect to detectability. Our results are presented in Appendix B and illustrate that both methods yield comparable S/N. In the next section we delve into the problem of estimating the spectrophotometry of companions detected in IFS data using LOCI-type reductions.

3. SPECTRAL EXTRACTION WITH d-LOCI

3.1. Spectral Bias and Contamination

If we write the discrete spectrum of a companion, sampled at the spectral resolution of the IFS, as $\{I_p^C\}_{p=1 \dots N_\lambda}$, then the flux of a potential companion in an LOCI-subtracted image at the wavelength λ_{p_0} can be written as

$$I_{p_0}^{C,\text{LOCI}} = I_{p_0}^C - I_{p_0}^{\text{Bias}} + I_{\text{Star}} - I_{\text{Cross}} \quad (17)$$

$$I_{p_0}^{C,\text{LOCI}} = I_{p_0}^C - I_{p_0}^{\text{Bias}} + I_{\text{Star}} - \sum_{p \in \mathcal{R}} c_p \gamma_{p_0,p} I_p^C. \quad (18)$$

The three biases in Equation (18) are described as follows:

1. $I_{p_0}^{\text{Bias}}$ corresponds to the amount of companion flux that is directly fitted by the PSF subtraction algorithm. It is the portion of the companion light that has been considered as starlight by the LOCI algorithm and thus subtracted. This bias can occur in all implementations of LOCI but is larger when using chromatic differential imaging, compared with ADI, because of the potential poor conditioning of the correlation matrix $\mathbf{M}_O^{(R,R)}$. The relative size of the subtraction and optimization zones, which directly scales with N_A , impacts $I_{p_0}^{\text{Bias}}$. This number ought to be carefully calibrated using synthetic point sources when measuring the photometry within any LOCI processed image.
2. The bias I_{Star} corresponds to the residual starlight after LOCI in the pixels of interest. If we assume that this residual background is constant around the companion PSF, it can be estimated and subtracted. As a consequence, we neglect it here.

3. I_{Cross} corresponds to the contamination of the measured flux at a given wavelength by the companion flux present in the composite reference at adjacent wavelengths. Indeed, the measurement at λ_{p_0} is contaminated by the spectral content of the companion at nearby wavelengths I_p^C multiplied by the corresponding LOCI coefficient, c_p , and a quantity, $\gamma_{p_0,p}$, that measures the spatial extent of the companion PSF between wavelength channels. \mathcal{R} is chosen such that for all p and q , $\gamma_{p,q} \ll 1$. However, in the absence of damping the coefficients $|c_p|$ can become large enough to make the spectral contamination non-negligible, $|c_p \gamma_{p,q}| \simeq 1$. This spectral cross-talk is specific to the use of LOCI with IFS images. It can be somewhat mitigated by optimizing the value of N_δ using synthetic companions. However, such an optimization relies on knowing a priori the spectrum of the companion to characterize and thus can only be carried out when this companion is bright enough to obtain a rough estimate of the spectrum in the absence of any speckle calibration.

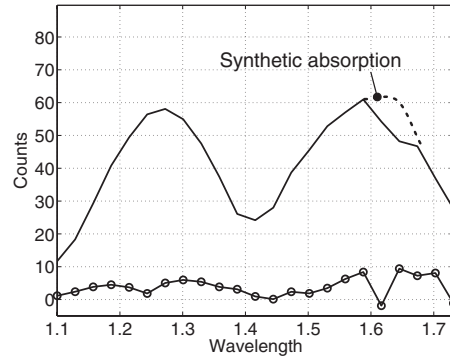

The final error on the retrieved spectrum is the sum of the algorithmic error, $I_{p_0}^C - I_{p_0}^{C,\text{LOCI}}$, and the spectroscopic calibration uncertainties. Calibrating the P1640 spectrograph involves a careful measurement of the instrumental spectral response function, whose determination procedure is detailed in Zimmerman et al. (2011), Zimmerman et al. (2010), and Hinkley et al. (2010a). The uncertainty on this response is not expected to be larger than 5%. As we will see below, the quantity $I_{p_0}^C - I_{p_0}^{C,\text{LOCI}}$ when using LOCI reduced data is much larger than this expected telluric calibration uncertainty and is thus very often the dominant source of errors. As a consequence, we carried out our spectral fidelity computations in ‘‘camera space,’’ using wavelength series that have not been normalized by the instrumental spectral response function. By doing so, we chose to leave aside discussions about telluric calibration and only focus on issues related to the effects of the PSF subtraction algorithms.

3.2. PSF Binary Weighting



In this section we consider the case of a faint companion that has been detected and needs to be characterized. We call \mathcal{P} the image plane location of the detected off-axis source and assume the geometry described on the left-hand side of Figures 1 and 2, where \mathcal{P} and \mathcal{S} have the same area but do not fully overlap. While the case of using a subtraction zone that is smaller than \mathcal{P} is of particular interest for the detection problem, this configuration leads to large signal depletions and is thus not discussed in the present communication. The top panel of Figure 1 illustrates a test case carried out using aperture photometry on an IFS data set reduced using the LOCI algorithm described in Lafrenière et al. (2007). We introduced a fake companion located at $1''$ in four Alcor data cubes. This corresponds to about 10 minutes of exposure, each exposure being composed of $N_\lambda = 23$ frames (around 80 references). We ran our reductions with constant LOCI parameters $dr = 2$, $N_\delta = 1$, $g = 3$, and $N_A = 100$. This choice of N_A leads to solving inverse problems that are slightly overconstrained and thus favorable to the regular LOCI in terms of detectability. In this figure the spectrum of the fake companion is similar to the one of the host star, except at $1.65 \mu\text{m}$, where a synthetic 10% flux depletion was inserted, and its brightness was chosen such that its integrated delta magnitude over the J and H band was $\Delta m_{J+H} = 6.5$. The retrieved spectrum is very different from the injected one. Because the observed discrepancy cannot be described as a ‘‘gray gain,’’ it is a combination of I_{bias} and I_{cross} , and it cannot be calibrated using a set of synthetic

(a) Regular LOCI:



Minimize least-squares fit in an optimization zone that includes the subtraction zone.


(b) S-zone masking LOCI:

Minimize least-squares fit in an optimization zone that does not includes the subtraction zone.


(c) Companion masking LOCI:

Minimize least-squares fit of the target frame in an optimization zone that includes the subtraction zone but not the pixels of the companion.


Legend:

- Spectrum of the star, similar to the spectrum of the injected synthetic companion except for synthetic absorption at 1.65 microns
- Spectrum of the injected companion
- ○ ○ Extracted Spectrum after LOCI

Figure 1. Illustration of three PSF binary weighting approaches, with their respective performances in the case of a companion slightly fainter than the P1640 speckles, $\Delta m = 6$. Note that the scale of the bottom panel has been compressed to accommodate for the large bias. None of these solutions provide robust spectrophotometric accuracy: they all lead to the combination of I_{bias} and I_{cross} . \mathcal{O} is the optimizations zone, used to find the best least-squares fit of the speckle realization. \mathcal{S} is the subtraction zone. \mathcal{P} is the position of the planet. \mathcal{R} is the ensemble of reference images.

companions. The 10% flux depletion at $1.65 \mu\text{m}$ allows us to illustrate the impact of I_{cross} . This weak absorption is strongly amplified via wavelength cross-talk in the reduction method, as shown in the top panel of Figure 1. This aspect is difficult to calibrate using synthetic sources since depletion depends on the depth of the putative absorption, which is unknown a priori. We first sought to address these issues by exploring several binary weighting strategies of the LOCI cost function.

1. *Subtraction zone masking* consists of reducing I_{bias} by not including the pixels in the \mathcal{S} zone in the least-squares fit. It excludes the companion flux from the inverse problem and

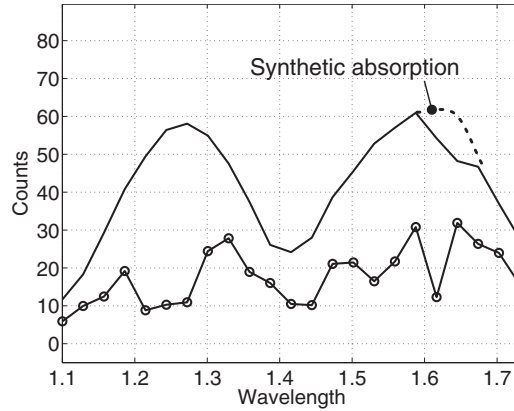

thus should preserve better the flux of potential companions. In practice, it is implemented using a weighting function $w(x, y)$ that is 1 in the optimization zone and 0 in the subtraction zone and corresponds to solving the following optimization problem:

$$\text{Find } c_k \text{ st } \min \left\{ \left\| T - \sum_k^{N_{\mathcal{R}}} c_k R_k \right\|_{\mathcal{O}-\mathcal{S}} \right\}. \quad (19)$$

The second panel of Figure 1 illustrates the efficiency of this approach in our test case. I_{bias} is indeed reduced; however,



(d) d-LOCI

Minimize least-squares fit of the target frame in an optimization zone while maximizing the residual in the subtraction zone.



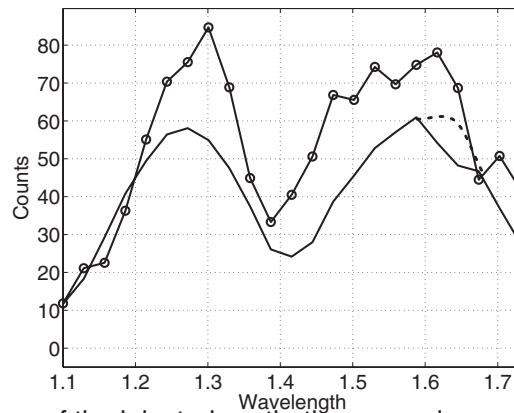

(e) Companion and S-zone masking LOCI:

Minimize least-squares fit of the target frame in an optimization zone that excludes the subtraction zone and the pixels of the companion.



(f) Companion masking d-LOCI:

Minimize least-squares fit of the target frame in an optimization zone that excludes the pixels of the companion, while maximizing the residual in the subtraction zone.



Legend:

- - - Spectrum of the star, similar to the spectrum of the injected synthetic companion
- - - except for synthetic absorption at 1.65 microns
- Spectrum of the injected companion
- ○ ○ Extracted Spectrum after LOCI

Figure 2. Illustration of three PSF binary weighting approaches combined with d-LOCI, with their respective performances in the case of a companion slightly fainter than the P1640 speckles, $\Delta m = 6$. Solutions (e) and (f) provide the best two spectral estimates, and their biases are carefully quantified in the present manuscript. \mathcal{O} is the optimizations zone, used to find the best least-squares fit of the speckle realization. \mathcal{S} is the subtraction zone. \mathcal{P} is the position of the planet. \mathcal{R} is the ensemble of reference images. The upward arrows illustrate the flux preservation constraint in d-LOCI.

it is achieved at the cost of an increase in I_{cross} : the flux of the 18th spectral channel is strongly depleted because of the large coefficients at neighboring wavelengths.

2. *Companion masking in the target image.* One way to reduce I_{cross} is minimizing the magnitude of the c_k 's at the location of the companion. Since the off-axis source has already been detected, this can be achieved by masking its pixels in the target image and thus solving the following least-squares problem:

$$\text{Find } c_k \text{ st min } \left\{ \left\| w_{\mathcal{P}}(x, y)T(x, y) - \sum_k^{N_{\mathcal{R}}} c_k R_k(x, y) \right\|_{\mathcal{O}} \right\}, \quad (20)$$

where $w_{\mathcal{P}}(x, y)$ is equal to 0 at the location of the detected companion and 1 everywhere else. This is equivalent to finding a composite reference that minimizes both the residual in the \mathcal{O} zone and the impact of the composite reference at the location of the companion:

$$\text{Find } c_k \text{ st min } \left\{ \left\| T - \sum_k^{N_{\mathcal{R}}} c_k R_k \right\|_{\mathcal{O}-\mathcal{P}} + \left\| \sum_k^{N_{\mathcal{R}}} c_k R_k \right\|_{\mathcal{P}} \right\}. \quad (21)$$

This is illustrated in the third panel of Figure 1. This approach can be interpreted as a simplified version of d-LOCI, without the coefficient positivity constraint and under the implicit assumption that $\Lambda = 1$. This approach considerably reduces I_{bias} to the point that the flux from the companion is overestimated.

3. *Subtraction zone masking and companion masking in the target frame* combine the two approaches and are illustrated in the second panel of Figure 2. They do alleviate spectral contamination but overestimates the companion flux, albeit at a much lower level than solely masking the companion in the target image.

The analysis using synthetic companions presented in Figures 1 and 2 was carried out at several angular separations and delta magnitudes, and consistent results were obtained. While using subtraction zone masking in conjunction with companion masking in the target frame yields good spectral fidelity at low-contrast levels, it degrades quickly when the synthetic companions get fainter. Replacing the \mathcal{S} zone masking by d-LOCI provides a robust alternative to this problem.

3.3. Results

3.3.1. Absolute and Relative Photometry

IFS observations of faint companions with instruments such as P1640 yield two astronomical observables: band-averaged photometry and low-resolution spectroscopy. Estimating both quantities from aggressively LOCI-subtracted data cubes requires two fundamentally different reduction strategies.

1. Band-averaged photometry (J and H bands in the case of P1640): accurate broadband photometry is crucial for deriving effective temperatures and determining masses based on evolutionary models. The band-averaged flux is insensitive to the intrinsic spectral shape; this is regardless of whether or not the companion has absorption or emission bands in the spectrum (main source of errors in I_{bias}). Therefore, the band-averaged LOCI response is relatively insensitive to the underlying spectrum and can be calibrated

using synthetic companions with a flat spectrum. This is equivalent to work done by several authors (Marois et al. 2008; Thalmann et al. 2009; Currie et al. 2011; Janson et al. 2011). LOCI parameters (N_A , g , dr , N_{δ}) can be tuned to optimize band-averaged performances. For instance, the bias of the integrated flux over both J and H bands, corresponding to the reductions illustrated in Figure 3, can be minimized using an LOCI parameter optimization based on synthetic companions. On this figure, regular LOCI reductions (middle column of Figure 3) exhibit more accurate band-averaged photometric estimates in the H band than in the J band: this suggests that the parameters can be further optimized in J band.

2. Low-resolution spectroscopy: accurate spectroscopy allows further constraints on physical parameters such as effective temperature, surface gravity, and the detection of molecular absorption bands. Here the color response of the algorithm is sensitive to the intrinsic spectral shape (main source of errors in I_{cross}). Calibrating this response without foreknowledge of the companion's spectrum is difficult. This is the main problem addressed in this paper: we show here that, even at the detection limit, d-LOCI provides an SED estimate much more accurate than classical LOCI. Importantly, this does not require any calibration of the color response of the algorithm.

We first discuss qualitatively the performance of four possibilities, with masking/no-masking and with regular/d-LOCI, with respect to band-averaged photometry and low-resolution spectroscopy. To do so, we introduced a synthetic companion located at $\simeq 1''$ separation in four data cubes on the star Alcor (Zimmerman et al. 2010). This is about 10 minutes of total exposure, with each three-dimensional cube composed of $N_{\lambda} = 23$ wavelength frames. The four data cubes allow a reference frame set of 80 frames. Unlike the A5V host Alcor, the 1.2–1.8 μm SED of the synthetic companion has H -band excess; this lowers the companion S/N at short wavelengths (J band) where the surrounding speckles are brighter. We then ran all reductions with constant LOCI parameters $dr = 2$, $N_{\delta} = 1$, $g = 3$, and $N_A = 100$; this choice of N_A leads to solving a set of inverse problems that are slightly overconstrained, with 80 degrees of freedom and 100 constraints. The results are shown in Figure 3, for three different companion magnitudes ranging a factor of 10 in flux. The leftmost column shows the before and after LOCI images used for these analyses. The “post-reduction” inset illustrates the companion's detection level in a single spectral channel. These images were obtained with classical LOCI and no priors on the companion's location. The behavior of the four different reductions illustrated in this figure allows insights into the algorithmic photometric and spectrophotometric accuracy of both unmasked and masked approaches. We now refer the reader to Figures 1–3.

Flux preservation properties of unmasked LOCI. The extracted spectral profiles represented by the lines labeled “Companion non-masked” in Figure 3 correspond to the implementation of non-masked reductions with classical LOCI (center column), solution (a), top panel of Figure 1 and with d-LOCI (right column), solution (d), top panel of Figure 2. d-LOCI exhibits an improvement in the algorithmic throughput of a factor of two. As there is no prior about the companion's location, and as the S/Ns of both solutions are similar (Section 3), we find that d-LOCI preserves a larger portion of the companion's flux. Moreover, the top set of panels of Figure 3 demonstrates that classical LOCI without binary weighting can lead to gray

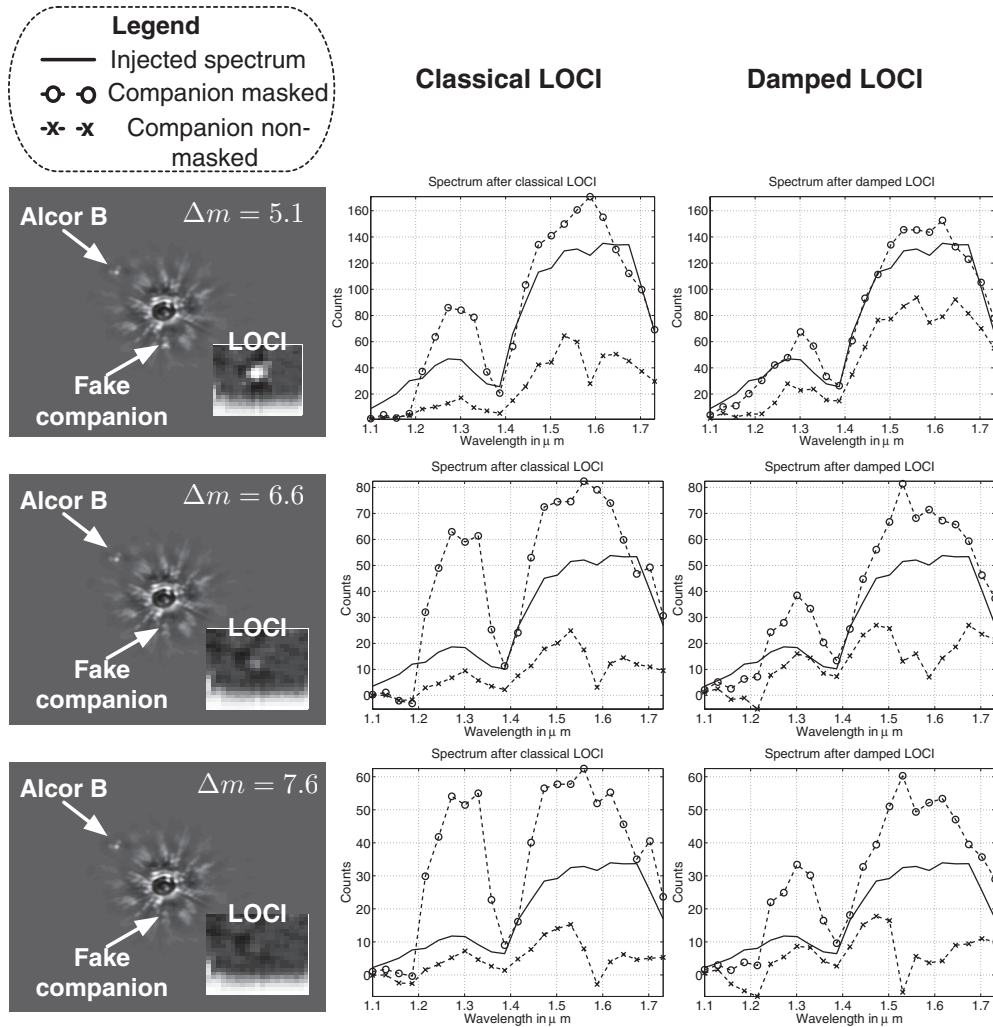


Figure 3. Qualitative comparison of the spectrophotometric accuracy of different LOCI reductions. We used synthetic companions at 1 arcsec at three different levels of contrast. The injected spectrum is represented by a solid line in all panels. Left panel: images before and after LOCI. Middle panel: classical LOCI, solution (a) in Figures 1 and 2, in crossed-dashed line; classical LOCI with dual masking, solution (e), in circled dashed line. Right panel: d-LOCI, solution (d), crossed-dashed line, d-LOCI with target frame masking, solution (f), in circled dashed line. When using d-LOCI, solution (f), most of the residual error corresponds to a gray gain, or absolute error. The relative SED error is smaller in this d-LOCI than in the classical masked case, solution (d). This figure illustrates approaches that do not include any iterations on the underlying companion’s spectrum or LOCI parameters.

spectral transmissivities. This holds when the companion is as bright as the speckles, as is seen in Bowler et al. (2010), where the authors used a set of synthetic companions to derive a gray algorithmic spectral response for HR 8799 b. However, this property quickly degrades with contrast, and more robust methods are needed when the companion is fainter than the speckles.

Spectrophotometric accuracy in masked or damped approaches. In the middle column of Figure 3, the spectral profiles represented by the lines labeled “companion masked” correspond to the implementation of LOCI with both subtraction zone masking and companion masking in the target frames. This refers to solution (e) in Figure 2. In the right panel the same spectra are generated by d-LOCI with companion masking in the target frames. This refers to solution (f) in Figure 2. The performance of both methods degrades as the companion gets fainter. However, the former clearly results in larger spectrophotometric errors than the latter (with damping).

These findings are summarized in Table 1, which describes the various levels of refinement of LOCI reductions necessary to carry out accurate measurement of band-averaged photometry and low-resolution spectroscopy. We identified three levels

of synthetic companion brightness: brighter than speckles (can be detected in a broadband image but its photometry is contaminated by the speckles), as bright as the speckles (can only be identified by its immobility in a movie of sequential wavelengths), and fainter than the speckles. The level of refinement of LOCI reduction was qualitatively inferred from the series of reductions carried out to generate Figure 3 (15 different companion brightnesses ranging from $\Delta m = 5$ to $\Delta m = 8$). Note that in this paper, no effort was made to optimize the sensitivity to faint off-axis sources. For details pertaining to the behavior of the detectability of P1640 as a function of observing conditions (e.g., seeing, host star magnitude, integration) and reduction procedure (e.g., optimization of LOCI parameters), we direct the reader to Crepp et al. (2011). Next we quantify the absolute and relative spectral accuracies that are obtained using the two most promising approaches identified in Table 1: dual PSF masking with classical LOCI and single PSF masking with d-LOCI.

3.3.2. Band-averaged Photometry Absolute Error

Here the absolute spectrophotometric error is the root-mean-squared error in the measurement of the companion flux,

Table 1
Summary of Our Findings

Preferred ^a LOCI method for given companion brightness and science goals			
Companion Brightness	Photometry w/o Iterations ^b	Photometry with Iterations ^c	SED Estimates ^d
Brighter than speckles	Dual masked LOCI	LOCI	Dual masked LOCI
As bright as speckles	d-LOCI	Dual masked LOCI	d-LOCI
Fainter than speckles	None	d-LOCI	d-LOCI

Notes.

- ^a For example, LOCI reduction with the smallest degree of sophistication that delivers a given science goal for a given companion brightness.
^b The iterations can be carried out on the LOCI parameters or the brightness of the object. In any case a sequence of synthetic companion injections and LOCI reductions is necessary in this scheme.
^c For example, one single LOCI reduction is sufficient to provide a robust photometric estimate.
^d Here we assume that LOCI biases on the SED depend on the underlying spectrum and thus are very difficult to calibrate using iterations without any prior on the faint companion.

integrated over the spectral channels spreading the J and H band

$$\epsilon_{\text{Abs}} = \frac{\sqrt{\sum_{p=1}^{N_\lambda} (s_p^{\text{telluric}})^2 (I_p^C - I_p^{C,\text{LOCI}})^2}}{\sum_{p=1}^{N_\lambda} s_p^{\text{telluric}} I_p^C}, \quad (22)$$

where s_p^{telluric} is the spectral response function, derived using reference A stars according to the method detailed in Zimmerman et al. (2011). The left column of Figure 4 shows how this metric varies as a function of the brightness and angular separation of a potential companion. The band-averaged photometric accuracy scales with the absolute spectrophotometric error. For comparison, the average level of photometric uncertainty in current ADI broadband discoveries of faint companions ($\delta m \simeq 0.2$; e.g., Marois et al. 2008; Thalmann et al. 2009; Currie et al. 2011; Janson et al. 2011) is indicated in Figure 4. d-LOCI consistently leads to smaller absolute spectrophotometric errors for the reduction parameters (N_A , g , N_δ , dr) tested here. The improvement in ϵ_{Abs} between dual masked LOCI and d-LOCI varies as a function of angular separation. This suggests, as discussed above, that the LOCI parameterization can be further optimized in order to reduce the systematic error bar on band-averaged photometry. Note that most of the spectrophotometric error when using d-LOCI in the faintest case of Figure 3 (bottom-right panel) consists of a band-average bias. Such an error can be mitigated using a set of synthetic companions to optimize LOCI parameters. In the faintest case, the actual improvement in relative SED (up to a scaling constant) provided by d-LOCI, when compared with classical LOCI (bottom-middle panel of Figure 3), is actually significant. We quantify this improvement in the next subsection.

Note that the two reduction methods considered in Figure 4 assume that a companion has already been detected and has been masked in the target frames. As a consequence, they can only be carried out for sources already identified using non-masked LOCI reductions. We separately evaluated the detectability of the synthetic companions in Figure 4 and did not include sources below the band-averaged $5 - \sigma$ level limit. We find that in the case of the P1640 data set used for these calculations, this regime occurs at $\Delta m \simeq 7.3$ with slight variations depending on the angular separation. While this number is relatively modest, we emphasize that this is mostly due to our conservative choices of data sets (average observing conditions, no wave front calibration before observations) of fake spectrum (faint in the J band), short exposure times (9 minutes), and reduction parameters (N_A , g , N_δ , dr). Further

optimizations of these parameters can surely lead to deeper detectability and absolute error that is consistent with the band-averaged photometric error bars published in previous ADI discoveries, all the way to the $5 - \sigma$ detection limit.

3.3.3. Low-resolution Spectroscopy: Relative Error

When seeking to characterize the spectral profile of a discovered companion within its near-infrared bands, calibrating the color response of the LOCI reduction without foreknowledge of the companion's spectrum is difficult. When it is sufficiently bright with respect to the speckles, the special case of a gray spectral response for the reduction arises. Bowler et al. (2010) provide a thorough illustration of this regime in the case of HR 8799b. When the algorithmic color response is gray, it is safe to assume that I_{bias} is responsible for most of the flux depletion. Indeed, since I_{cross} mostly amplifies local absorption features, it is unlikely to be constant over the full spectrum. In this particular case both the photometry and the low-resolution spectral features of the object can be estimated. However, as the companion gets fainter, the reduction's spectral response starts varying with wavelength. Without prior knowledge about the spectral profile of the companion, it is very difficult to disentangle whether such variations are the sole result of LOCI (e.g., I_{bias}) or stem from the spectrum of the faint source propagated through LOCI (e.g., I_{cross}). One solution to address this issue would be to carry out a series of iterative reductions: (1) extract a biased spectrum from a preliminary LOCI reduction, (2) subtract this spectrum from the non-reduced data, (3) run a secondary LOCI reduction, and (4) iterate. Since the biases depend on the set of LOCI parameters chosen for the reduction (N_A , g , N_δ , dr), this process should be repeated over a wide set of parameters in order to ensure that the biases on the SED are minimized. In this paper, we answer an orthogonal question, namely, *how to minimize the relative spectrophotometric bias in one iteration, for a well-chosen but not optimized set of LOCI parameters.*

We quantify the relative spectrophotometric error as a measure of the uncertainty in estimating the SED of the discovered object:

$$\epsilon_{\text{Rel}} = \sqrt{N_\lambda \sum_{p=1}^{N_\lambda} (s_p^{\text{telluric}})^2 \left(\frac{I_p^{C,\text{LOCI}}}{\sum_{k=1}^{N_\lambda} s_p^{\text{telluric}} I_k^C} - \frac{I_p^C}{\sum_{k=1}^{N_\lambda} s_p^{\text{telluric}} I_k^C} \right)^2}. \quad (23)$$

In other words, ϵ_{Rel} captures how well the overall shape of the low-resolution spectrum can be retrieved, up to a scaling

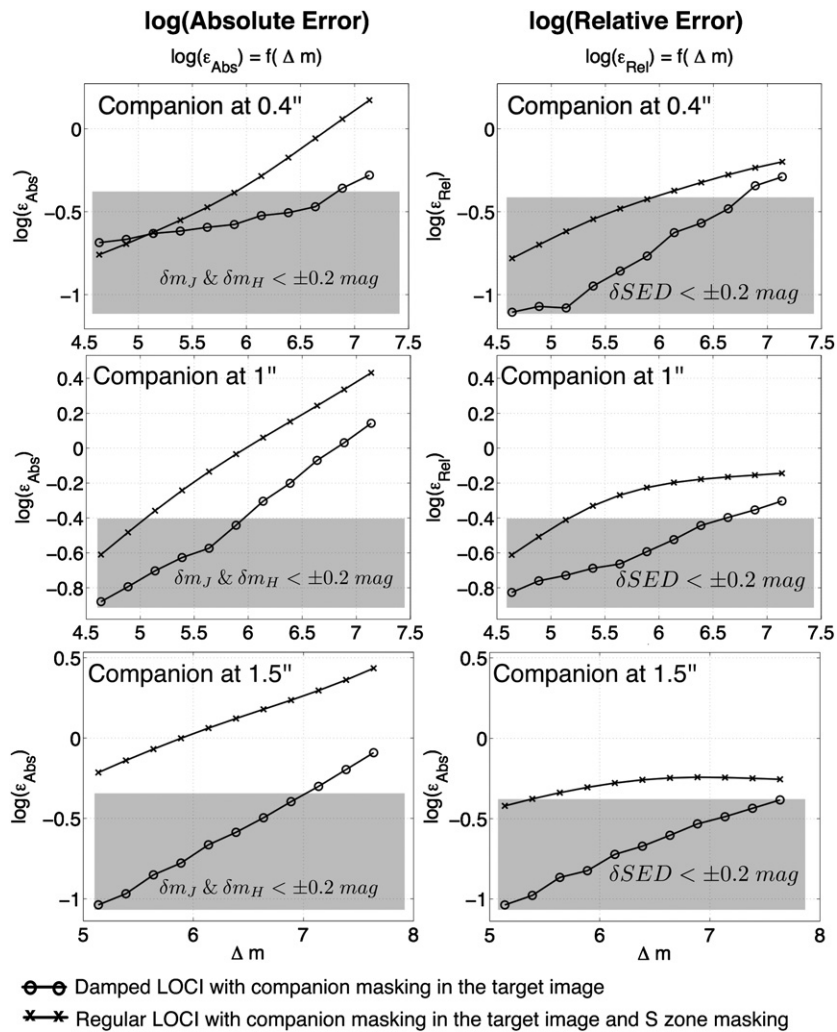


Figure 4. Quantitative comparison of the spectrophotometric accuracy for LOCI reductions (d) and (e) as a function of contrast at three different angular separations. Left: band-averaged photometric error. Right: bias in the SED estimate. The gray area corresponds to published band-averaged photometric error bars on recently discovered faint companions using ADI: δm_J & $\delta m_H < \pm 0.2$ mag. Even in the challenging “IFS + no field rotation” configuration, d-LOCI delivers systematic errors on the SED of the order of state-of-the-art uncertainties derived from band-averaged ADI observations, all the way to the detection limit of the P1640 instrument.

constant. Assuming that the band-averaged photometry of the object has already been estimated, this SED provides the quantitative basis to further characterize the discovered companion. In particular, it provides a remarkable level arm to estimate surface gravity, atmospheric chemistry (Burgasser et al. 2006; Oppenheimer et al. 1995), and, to a lesser extent, the atmospheric cloud content of the faint object (Burgasser et al. 2010). Therefore, the spectrophotometric error will have a direct impact on the estimation of the physical parameters of the faint exoplanets discovered by campaigns such as GPI, SPHERE, and P3000-P1640. Without any priors about the spectral content of such objects, calibration of the biases introduced by the reduction using synthetic sources becomes a very delicate exercise. A first avenue that can be explored when seeking to address this issue of faint companion spectral characterization is to devise a new class of algorithms that rely on priors about the speckle field (e.g., Burke & Devaney 2010; Mesa et al. 2011) or the spectral type of the companion (e.g., spectral deconvolution as in Sparks & Ford 2002). The solution proposed in the present manuscript, d-LOCI, focuses on class of solution, which assumes no priors about the quasi-static errors or the companion.

The right panel of Figure 4 demonstrates that d-LOCI is a robust approach when seeking to characterize the spectral features of a faint companion without any foreknowledge of its spectral type. Indeed, the relative error using d-LOCI in conjunction with target frame masking is, for most separations and levels of contrast, half an order of magnitude smaller than the one obtained using dual masked classical LOCI. More importantly, it yields SED uncertainties below the error-bar level of state-of-the-art band-averaged photometric estimates (Marois et al. 2008; Thalmann et al. 2009; Metchev et al. 2009). This feature is quite remarkable since it reduces the spectrophotometric uncertainties associated with observations with an IFS *without* field rotations to the level of observations carried out in a more favorable mode: ADI ($I_{\text{cross}} = 0$) and images integrated over several P1640 channels (more photons from the companion). For the brightest synthetic companions tested here, the SED bias obtained using d-LOCI is of the order of the P1640 telluric calibration uncertainties. We also emphasize that such results are obtained without any iterations on LOCI parameter or underlying companion spectrum, and thus d-LOCI reduced images are a good first guess when seeking to

obtain even higher spectrophotometric accuracy using iterative schemes.

d-LOCI is thus an efficient method when one seeks to estimate the SED of a detected point source in P1640 data, without any prior knowledge about the spectral profile of the companion or about the chromaticity of the post-coronagraphic speckles. This result ought to be carefully extrapolated to other instruments, in particular using similar tests based on synthetic companions. We expect, however, that such conclusions will stand when the raw contrast of direct imaging campaigns improves, using either a real-time calibration system (Wallace et al. 2010) or advanced pre-observing calibration methods (Serabyn et al. 2010). We thus envision that the approach suggested herein will be critical to the estimation of the physical parameters and atmospheric content of exoplanets detected via future direct imaging campaigns (Macintosh et al. 2006; Beuzit et al. 2008; Dekany et al. 2006; Hinkley et al. 2011).

4. CONCLUSION

In this paper we discuss an algorithm for PSF subtraction in data acquired with an IFS that is efficient and accurate for retrieving the spectral content of faint companions located within several diffraction widths of a target star. A previous communication by Crepp et al. (2011) used the newly functional Project 1640 spectrograph (Hinkley et al. 2011) as a laboratory to devise a robust method purely for off-axis faint source detection. This method relies on incorporating the spectral diversity inherent in IFS data into an LOCI PSF calibration scheme. In that paper, we reported H -band 5σ detection limits of $\simeq 10^{-5}$ in contrast, at 1 arcsec separation, in 20 minutes of exposure on an $H = 3$ star. In the present paper, we expand on the general LOCI scheme and discuss the spectral characterization of already-detected faint companions. We show that when LOCI is tuned to generate an *aggressive* composite calibration PSF, the spectrophotometric content of the faint object is biased. It is affected by both signal depletion, which is common to all LOCI implementations, and spectral cross-talk with neighboring wavelengths, which is specific to IFS imaging spectroscopy. Herein, we introduced a new algorithm “d-LOCI” that alleviates many of these shortcomings.

We modified the formalism underlying LOCI PSF calibration by introducing an augmented cost function that seeks to preserve the flux from the companion in question. We compared the efficiency of LOCI and d-LOCI for “detectability” and found that both approaches performed similarly in the various test cases. We identified second-order trends in the behavior of S/N as a function of exposure time and LOCI parameters. Such trends can be explained by the conditioning of the inverse problem that is at the heart of LOCI reduction and have only marginal bearing on the conclusions pertaining to “detectability” discussed in Crepp et al. (2011). While d-LOCI clearly increases the signal of faint companions, it results in a noise floor that is shallower than that obtained in LOCI, and thus the detectability metric of both approaches is similar.

We finally delved into the problem of spectral extraction. We first used synthetic companions to explore the behavior of classical LOCI reductions, implemented with a variety of binary weightings. Our tests showed that a classical implementation of LOCI, with well-chosen parameters as in Lafrenière et al. (2007), can be applied to companions slightly brighter than the speckles and yields an “almost gray” algorithmic throughput that yields a 10% error in SED estimates. In this regime we find that the spectral fidelity can be improved by the use of

a dual binary weighting when posing the PSF reconstruction problem. Since we expect telluric calibration errors below 10%, the main source of errors in the spectrophotometric characterization of the discovered companions resides in bias within the reduction algorithm. Moreover, these biases increase rapidly as a function of contrast, and when the companion is fainter than the speckles, the error in the SED estimate can be as large as 50%. Using d-LOCI, in conjunction with a binary weighting of the target frames at the location of the companion, consistently reduces this error by a factor of five: it lowers the algorithmic errors in the SED close to the levels imposed by the telluric calibration. Translating these uncertainties in error bars on the estimates of the physical parameters of potentially discovered companions is beyond the scope of this paper and will be illustrated in a future communication. All results are obtained using the Project 1640 IFS data, but we fully expect that these results can be generalized to future high-contrast imaging instruments operating in conjunction with advanced wave front calibration techniques. These next-generation instruments (Macintosh et al. 2006; Beuzit et al. 2008; Dekany et al. 2006; Hinkley et al. 2011) will operate at contrasts significantly higher than those possible with the current PALAO+P1640 configuration, and they will use IFS as science cameras. Accurate in situ spectral characterization of faint exoplanets buried in the residual speckle fields of the long exposures will thus be a crucial component of upcoming campaigns. It will lead to far tighter constraints on planetary bulk physical properties and atmospheric chemistry and would be a direct application of the d-LOCI methods discussed herein.

The research described in this publication was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Project 1640 is funded by National Science Foundation grants AST-0520822, AST-0804417, and AST-0908484. This work was partially funded through the NASA ROSES Origins of Solar Systems Grant NMO710830/102190, the NSF AST-0908497 Grant. The adaptive optics program at Palomar is supported by NSF grants AST-0619922 and AST-1007046. L.P. was supported by an appointment to the NASA Postdoctoral Program at the JPL, Caltech, administered by Oak Ridge Associated Universities through a contract with NASA. L.P. and S.H. acknowledge support from the Carl Sagan Fellowship Program.

APPENDIX A

FRAME SELECTION AND GEOMETRY

In this section we consider the case of a series of IFS cubes that have already been centered and spatially rescaled so that the optical artifacts are stationary within the image space and the astronomical image has been stretched radially (Crepp et al. 2011; Sparks & Ford 2002). Proper scaling and centering of the images are crucial to subsequent data analysis. We carry out this stage by finding the centering and scaling parameters that maximize the autocorrelation between two images. We solve this nonlinear search using the sub-pixel image registration algorithm described by Guizar-Sicairos et al. (2008), which has been modified using the discrete Fourier transform algorithm described in Soummer et al. (2007b), in order to accommodate for the scaling parameter. The series of images resulting from such a calibration corresponds to a virtual four-dimensional cube, with two spatial dimensions, one spectral dimension,

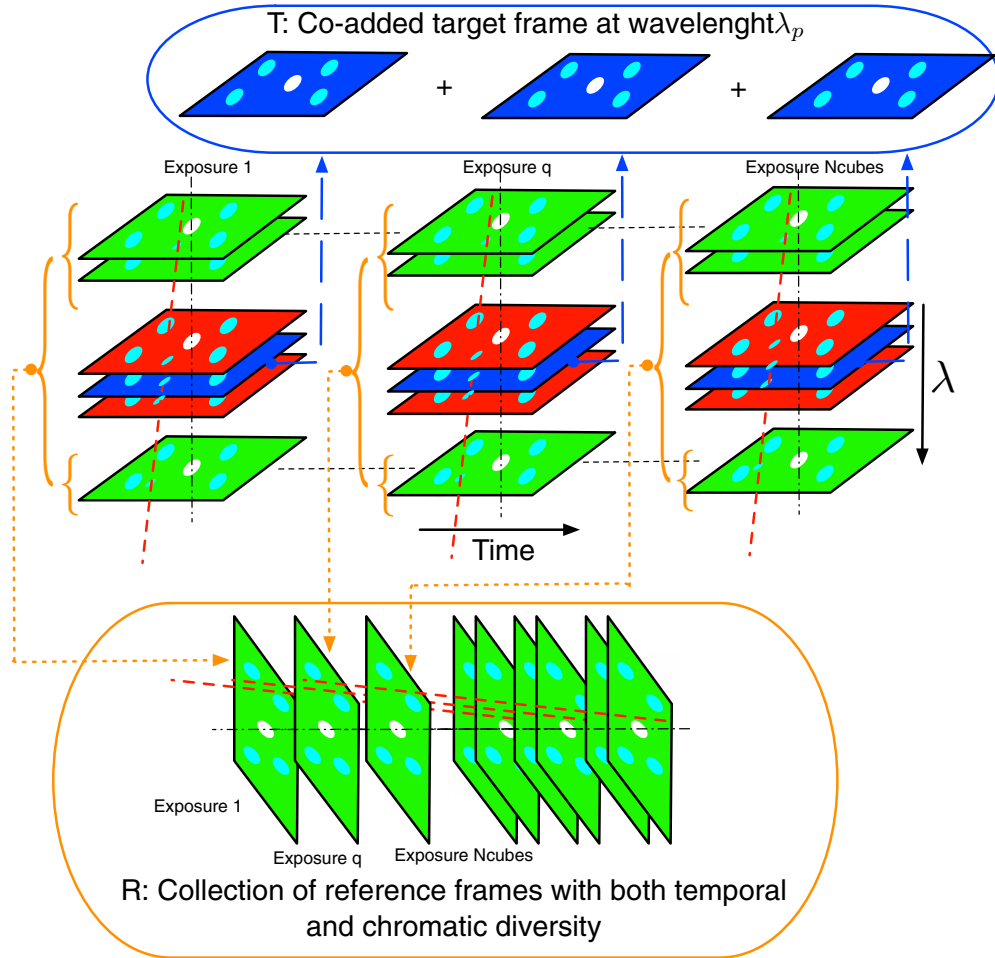


Figure 5. Illustration of the frame classification method. The data consist of a collection of cubes, represented in the middle panel. Each cube corresponds to one exposure, and each slice within a cube is an image at a given wavelength: here, the temporal dimension is represented horizontally and the wavelength vertically. We co-add the frames at the wavelength λ_{p_0} into a target frame T , top panel. We create a library of reference frames that will be used to build the composite calibrator $\{R_k\}$. These references are chosen so that the flux of a potential companion at neighboring wavelengths is minimal at the radial location of interest (e.g., minimize self-subtraction). The reference ensemble presents both chromatic and temporal diversity.

(A color version of this figure is available in the online journal.)

and a temporal dimension. We seek to devise a reduction algorithm that will stack the signal of a potential companion for a given wavelength, while subtracting the artifacts due to optical imperfections. We write each frame in our four-dimensional cube as $I_p^{(q)}$, where the subscript $p = 1, \dots, N_\lambda$ stands for the spectral channel and the superscript $q = 1, \dots, N_{\text{cubes}}$ stands for the cube number, or the temporal dimension. Associated with these slices is a collection of scaling parameters addressed similarly, $\gamma_p^{(q)}$. We first build an integrated target frame T , with an associated scaling parameter γ_{p_0} , that corresponds to the sum of the frames at the wavelength of interest λ_{p_0} :

$$T = \sum_{q=1}^{N_{\text{cubes}}} I_{p_0}^{(q)} \quad (\text{A1})$$

$$\gamma_{p_0} = \frac{1}{N_{\text{cubes}}} \sum_{q=1}^{N_{\text{cubes}}} \gamma_{p_0}^{(q)}. \quad (\text{A2})$$

This is illustrated in the top panel of Figure 5, where slices of the same wavelength are co-added. We seek to create a composite reference that locally subtracts the quasi-static speckles without

removing the flux from a putative companion. Thus, at a given radial location within the target frame, written as r , we need to sort the remaining frames into two subsets: $\{R_k\}$, the subset of slices that will be used to create the composite reference, and $\{J_k\}$, a subset of “useless” frames whose wavelength is too close to λ_p so that the image of a potential companion would overlap with the region of interest. Since the several wavelengths in the data cube provide radial diversity, the image of a potential companion moves radially in the rescaled space, and we can rewrite the self-subtraction condition established by Lafrenière et al. (2007) as

$$\mathcal{R} = \{R_k\} = \{I_p^{(q)} : |r(\gamma_p^{(q)} - \gamma_{p_0})| \geq N_\delta W\} \quad (\text{A3})$$

$$\mathcal{J} = \{J_k\} = \{I_p^{(q)} : |r(\gamma_p^{(q)} - \gamma_{p_0})| < N_\delta W\}, \quad (\text{A4})$$

where W is the FWHM of a PSF at wavelength λ_{p_0} and N_δ is a parameter that determines the aggressiveness of the LOCI reduction algorithm. This sorting process involves rejecting the slices whose wavelength is too close to λ_{p_0} from the collection of references, so that there is minimal companion flux in the composite references. We actually do not integrate over time the slices in the *reference ensemble* and let LOCI choose




Figure 6. Geometry associated with a polar implementation of LOCI. The reference coefficients are chosen in order to obtain the best least-squares fit of the target speckle pattern in the *large* optimization zone, \mathcal{O} . The corresponding composite reference PSF is then subtracted from the target frame only in a *smaller* area, the subtraction zone \mathcal{S} , that is included in \mathcal{O} . A zoom on the \mathcal{S} after LOCI subtraction is displayed in the rightmost panel of the figure.

(A color version of this figure is available in the online journal.)

the best coefficients to form an optimal reference PSF over time and wavelength. The strategy summarized in Figure 5 was chosen in order to mitigate computational cost. Indeed, calculating an LOCI reference PSF for each slice of each exposure using PSFs at other wavelengths and slightly decorrelated in time would mean running $N_{\text{cubes}} \times N_{\lambda}$ LOCI reductions, with each reduction looking for coefficients over an ensemble of $N_{\text{cubes}} \times (1 - \epsilon(N_{\delta}))N_{\lambda}$ reference PSFs ($\epsilon(N_{\delta})$ is the number of slices excluded to avoid cross-talk). Instead, our reduction strategy only leads to N_{λ} LOCI reductions with each reduction looking for coefficients over an ensemble of $N_{\text{cubes}} \times (1 - \epsilon(N_{\delta}))N_{\lambda}$ reference PSFs. The performance of each reduction strategy in terms of detectability depends on the actual shot noise level of the speckle field (e.g., whether LOCI actually reconstructs the true realization of the speckle field in the target image or fits the photon noise associated with it). The true merit of a given frame selection strategy thus depends strongly on observational conditions such as stellar magnitude or integration time. Finding the optimal binning scheme for given observational conditions is beyond the scope of the present manuscript, and we did not delve into such trade-offs. Instead, we chose a hybrid frame selection approach that preserves the temporal information in the ensemble of reference PSFs while integrating the time variability of the target science image. The next step of the reduction algorithm is to find a series of optimal coefficients in order to build the composite reference. This is the stage that can lead to a loss of the spectral information of the detected companion. In order to formally describe the challenges of LOCI reductions, we first briefly review the notations involved in the framework introduced by Lafrenière

et al. (2007). As shown in Figure 2, the local composite reference PSF is constructed by finding a linear combination of reference frames, $\{R_k \in \mathcal{R}\}$, that yields the best least-squares fit of the speckle pattern in the target image, T . This locally optimal calibration image is then subtracted from the target image. However, in order for this reference to only describe the speckle field and not the flux of a potential companion, its coefficients ought to be determined in a zone of the image that is larger than the zone within which the subtraction will occur. Thus, the LOCI algorithm uses two regions of interest in the images, a subtraction zone, \mathcal{S} , where we will subtract the composite reference from the target, and a large optimization zone, \mathcal{O} , that surrounds \mathcal{S} , which we will use to find the weights of the composite reference. We follow the presentation of Lafrenière et al. (2007) and use an annular geometry, as illustrated in Figure 6. The optimization zone is defined by its area, expressed in number of PSF cores, N_A , and its aspect ratio g . The aspect ratio is a unitless quantity that scales with the radial elongation of the zone \mathcal{O} ,

$$g = \frac{\Delta r}{(r + \Delta r/2)\Delta\phi}, \quad (\text{A5})$$

where r is the radial location of the inner annulus of the optimization zone, Δr is its radial extent, and $\Delta\phi$ is its azimuthal extent. They can be written as a function of N_A , g , and W :

$$\Delta r = \sqrt{\frac{\pi N_A g W^2}{4}} \quad (\text{A6})$$




Figure 7. Data set used for S/N calculations. Left column: raw (top) and processed (bottom) images in the J band. Right column: raw (top) and processed (bottom) images in the H band. We used a spiral of 11 synthetic companions with a flat spectrum whose angular separation varies from 0.3 to 0.5 arcsec. Because of the better performance of the adaptive optics system at longer wavelengths, the raw contrast is more favorable in the H band than in J band. In this example the exposure time is 140 s. Since the average contrast is $\simeq 10^{-4}$ in J band and $\simeq 10^{-3}$ in H band, the synthetic companions are detected with a higher S/N in the latter case. This figure was generated using images reduced by d-LOCI.

$$\Delta\phi = \left(\frac{g}{2} + \frac{2r}{W} \left(\frac{g}{\pi N_A} \right)^{1/2} \right). \quad (\text{A7})$$

In the geometry chosen here, illustrated in Figure 6, the inner radius of the subtraction and optimization zones is similar and equal to r . Their azimuthal span is likewise identical and equal to $\Delta\phi$. The radial extent of the S zone is dr , while the O zone stretches all the way to $r + \Delta r$. The efficiency of the LOCI reduction strongly depends on the choice of the parameters N_A , g , dr , and N_δ . Lafrenière et al. (2007) explored the parameter space for LOCI applied to ADI, and Crepp et al. (2011) presented similar work in the case of chromatic differential imaging. In this paper we show that a standard LOCI implementation applied to IFS data yields very good results in terms of detectability but perturbs the spectrum of the companions discovered. While the errors associated with the spectral estimation can be reduced by exploring the N_A , g , dr , N_δ parameter space, we propose a more direct solution that consists of constraining the problem associated with the LOCI coefficient search.

APPENDIX B

DETECTION AND SIGNAL-TO-NOISE RATIO

We first study the impact of d-LOCI on the detectability of faint point sources and show that the damping introduced herein only marginally affects the S/N of detected companions. To do

so, we use a data set obtained on the bright A5V star ($V = 4$) Alcor on 2009 March 16. We introduced in the data a spiral of 11 synthetic companions with a flat spectrum, at angular separations varying from 0.3 to 0.5 arcsec. The injection of the artificial sources occurs as high up as possible in the cube processing pipeline, as described in Crepp et al. (2011). Alcor bright visible magnitude enables good performances from the PALAO (Dekany et al. 2006). A faint M dwarf was recently detected using the P1640 spectrograph (Zimmerman et al. 2010), making this data set an ideal test case. The contrast of the astronomical companion is moderate, $1/30$ as estimated by Mamajek et al. (2010), and the secondary star is visible above the speckle halo: it provides a second reference for the photometry of the synthetic companions inserted. Figure 7 presents unprocessed images and the results of a data reduction, both integrated in J and H bands. Since the speckles are brighter at shorter wavelengths, and the spectrum of the synthetic sources injected is gray, the overall contrast associated with them varies between 10^{-2} and 10^{-3} in the H band and 10^{-3} and 10^{-4} in the J band. As a consequence, some companions are too faint to be detected in the J reduced images. These results were obtained using the following parameters: one exposure of 140 s, $N_\lambda = 23$ slices across the cube of interest, $N_A = 100$, $g = 1$, $dr = 2$, and $N_\delta = 1$. This choice of parameters, in particular $g = 1$, implies that two adjacent synthetic companions might be located in the same LOCI zones and thus might artificially hamper the S/N of reduced images. However, we included our “conservative,” $g = 1$, results in this manuscript, as its purpose was not to




Figure 8. Detectability of faint point sources in the case of an overconstrained LOCI reduction: S/N as a function of the area of the optimization zone for constant exposure time and number of references (here the exposure time is 140 s and $N_{\mathcal{R}} < 23$). The synthetic companions illustrated in Figure 7 were used for these calculations: the figure shows the average of the post-LOCI S/N over the 11 companions. Top curves: *H*-band results. This shows that in the overconstrained regime d-LOCI performs almost as efficiently as classical LOCI as far as the detectability of faint point sources is concerned. While these results were obtained using $g = 1$, reduction with $g \gg 1$ yields the same qualitative behavior, albeit at slightly higher S/N levels. Bottom curves: *J*-band results. Details about the second-order trends in this figure are discussed in the body of the present manuscript.

quantify the ultimate detection limits of P1640. We refer to Crepp et al. (2011) for an example of a P1640 reduction when introducing the companions one at a time. Note also that the exposure time can be lengthened and the other algorithm parameters can be finely tuned for better S/N. However, we do not delve into the exploration of the whole parameter space: the results of such an optimization are presented in Crepp et al. (2011). In this section we describe a couple of test cases that illustrate the different regimes of the regular and d-LOCI algorithms with regard to point-source detectability.

We first consider the case of an overconstrained problem for which the number of free parameters, which scales with the number of reference frames, is smaller than the number of constraints, which scales with N_A , the number of PSF cores in the zone \mathcal{O} . We only reduce one data cube. This corresponds to a short integration, 140 s, and a small number of references to choose from in order to build the calibration PSF: $N_{\mathcal{R}} < N_{\lambda} = 23$ (the actual value of $N_{\mathcal{R}}$ varies as a function of angular separation; at the location of the injected synthetic sources a typical value is $N_{\mathcal{R}} \sim 15$). We fix the geometry of the problem as follows: $g = 1$, $dr = 2$, and $\delta = 1$. We then vary the size of the optimization zone by changing the value of N_A from 20 to 300 and thus study the behavior of both LOCI and d-LOCI in the overconstrained case since $N_A > N_{\mathcal{R}}$ in these configurations. Figure 8 illustrates how the average S/N of the 11 injected companions, integrated over the wavelength channels in both *J* and *H* bands, behaves in this configuration. The signal is estimated using aperture photometry at the known location of each artificial companion. The noise is estimated by computing the standard variation of the flux of a set of 100 boxes surrounding the spiral of fake planets. Each photometric box has the area of a PSF core. Since the purpose of this

section is to illustrate general trends pertaining to detectability with LOCI and d-LOCI, we chose to only isolate two “raw contrast” levels: “faint” companions in *J* band and “bright” companions in *H* band, and only show how the averaged S/N over the 11 companions behaves for these two sub-groups. A detailed study, with synthetic sources of continuously decreasing brightness, was carried out in Crepp et al. (2011) in order to derive the sensitivity limit of P1640. Both LOCI and d-LOCI present the same overall level of performance with slight upward S/N trend in both bands. In spite of having a penalty term that worsens the least-squares residual in the zone of the image where the PSF is subtracted, d-LOCI performs nearly as well as regular LOCI in the overconstrained case. The small S/N difference can be understood as follows: the noise suppression using pseudo-inverse methods is always more aggressive than when a damping term is introduced, and, in the overconstrained case, the likelihood of fitting companion flux with starlight is low. The upward trend in S/N as a function of the area of the optimization zone means, for the data set studied here, that the higher companion throughput (e.g., larger signal) achieved by using a larger N_A slightly dominates the degradation of the noise owing to the fact that the composite calibration image is calculated over a less local area.

We then study the case of an underconstrained problem, for which the number of reference frames is larger than the number of PSF cores in the zone \mathcal{O} . To do so, we now fix the area of the optimization zone $N_A = 100$ and vary the number of cubes reduced from 1 to 14. According to our frame selection strategy, illustrated in Figure 5, using more data cubes should increase the signal, since co-adding cubes increases the exposure time, and should also reduce the noise after LOCI, since $N_{\mathcal{R}}$ increases with the number of cubes, thus making the subtraction more efficient. This means that we expect the S/N to be an increasing function of $N_{\mathcal{R}}$. Our results are shown in Figure 9. Classical LOCI *H*-band results exhibit this behavior, with S/N increasing from 8 to 11, and yield better performances than d-LOCI reductions, for which the S/N remains constant at 8. Note that the rise in detectability does not scale as $\sqrt{N_{\mathcal{R}}}$. This suggests that both our frame selection method (which was chosen for computational cost purposes) and our LOCI parameters are not fully optimized for the synthetic companions injected. In the *J* band, for which the contrast is less favorable, d-LOCI leads to constant detectability levels (S/N = 4.5), while the performances of classical LOCI actually degrade with $N_{\mathcal{R}}$. In this configuration the likelihood of fitting companion flux is high, since there are more degrees of freedom than there are constraints, and thus the algorithmic throughput decreases when $N_{\mathcal{R}}$ increases. When the companion is faint enough (in the *J* band), classical LOCI implementations that are underconstrained can reach a regime for which the companion’s flux suppression is greater than the speckle noise attenuation.

The cases illustrated in Figures 8 and 9 first confirm that when seeking faint companions using LOCI reduction methods, one should optimize the algorithm parameters beforehand using synthetic sources. In particular, they show how the conditioning of the correlation matrix $\mathbf{M}_{\mathcal{O}}^{(R,R)}$ is a critical metric in this parameter search and that N_A and $N_{\mathcal{R}}$ ought to be adjusted so that the inverse problem is not underconstrained. This was carried out in the case of P1640 by Crepp et al. (2011). Second, they show that d-LOCI detectability performances are close to the ones of classical LOCI. Overall, the damping makes d-LOCI less sensitive to the conditioning of the inverse problem, and its flat response to variations of the algorithm parameters might




Figure 9. Detectability of faint point sources in the case of an underconstrained LOCI reduction: S/N as a function of the number of cubes reduced (increasing exposure time and increasing number of references), for a constant area of the optimization zone, constant N_A . The synthetic companions illustrated in Figure 7 were used for these calculations: the figure shows the average of the post-LOCI S/N over the 11 companions. In the case of underconstrained problems, the relative performances of classical LOCI and d-LOCI with respect to detectability are a function of the contrast of the companion. Top curves: *H*-band results. For bright companions, *H*-band results, classical LOCI provides better noise suppression. Bottom curves: *J*-band results. For fainter companions, *J* band, the noise suppression with classical LOCI is associated with a large signal depletion (e.g., the companion's signal is fitted by one of the references). When using d-LOCI, this depletion is mitigated, and thus larger S/Ns are obtained. While these results were obtained using $g = 1$, reduction with $g \gg 1$ yields the same qualitative behavior, albeit at slightly higher S/N levels.

be used in order to discriminate astronomical companions from speckle residuals in the case of marginal detection. However, a careful study of this second-order trend is beyond the scope of the present paper.

REFERENCES

Aime, C., & Soummer, R. 2004, *ApJ*, 612, L85
 Beuzit, J.-L., Feldt, M., Dohlen, K., et al. 2008, *Proc. SPIE*, 7014, 701418
 Biller, B. A., Close, L. M., Masciadri, E., et al. 2007, *ApJS*, 173, 143
 Bordé, P. J., & Traub, W. A. 2006, *ApJ*, 638, 488
 Bowler, B. P., Liu, M. C., Dupuy, T. J., & Cushing, M. C. 2010, *ApJ*, 723, 850
 Burgasser, A. J., Burrows, A., & Kirkpatrick, J. D. 2006, *ApJ*, 639, 1095

Burgasser, A. J., Simcoe, R. A., Bochanski, J. J., et al. 2010, *ApJ*, 725, 1405
 Burke, D., & Devaney, N. 2010, *J. Opt. Soc. Am. A*, 27, A246
 Crepp, J. R., Pueyo, L., Brenner, D., et al. 2011, *ApJ*, 729, 132
 Currie, T., Burrows, A., Itoh, Y., et al. 2011, *ApJ*, 729, 128
 Dekany, R., Bouchez, A., Britton, M., et al. 2006, *Proc. SPIE*, 6272, 62720
 Fitzgerald, M. P., & Graham, J. R. 2006, *ApJ*, 637, 541
 Give'on, A., Belikov, R., Shaklan, S., & Kasdin, J. 2007, *Opt. Express*, 15, 12338
 Guizar-Sicairos, M., Thurman, S. T., & Fienup, J. R. 2008, *Opt. Lett.*, 33, 156
 Guyon, O., Pluzhnik, E., Martinache, F., et al. 2010, *PASP*, 122, 71
 Hinkley, S., Oppenheimer, B. R., Brenner, D., et al. 2010a, *ApJ*, 712, 421
 Hinkley, S., Oppenheimer, B. R., Soummer, R., et al. 2007, *ApJ*, 654, 633
 Hinkley, S., Oppenheimer, B. R., Zimmerman, N., et al. 2011, *PASP*, 123, 74
 Janson, M., Carson, J., Thalmann, C., et al. 2011, *ApJ*, 728, 85
 Kalas, P., Graham, J. R., Chiang, E., et al. 2008, *Science*, 322, 1345
 Krist, J. E., Stapelfeldt, K. R., Golimowski, D. A., et al. 2005, *AJ*, 130, 2778
 Lafrenière, D., Marois, C., Doyon, R., Nadeau, D., & Artigau, É. 2007, *ApJ*, 660, 770
 Leconte, J., Soummer, R., Oppenheimer, B. R., et al. 2010, *ApJ*, 716, 1551
 Macintosh, B., Graham, J., Palmer, D., et al. 2006, *Proc. SPIE*, 6272, 62720
 Mamajek, E. E., Kenworthy, M. A., Hinz, P. M., & Meyer, M. R. 2010, *AJ*, 139, 919
 Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, *ApJ*, 641, 556
 Marois, C., Macintosh, B., Barman, T., et al. 2008, *Science*, 322, 1348
 Marois, C., Macintosh, B., & Véran, J. 2010, *Proc. SPIE*, 7736, 77361
 Marois, C., Racine, R., Doyon, R., Lafrenière, D., & Nadeau, D. 2004, *ApJ*, 615, L61
 McElwain, M. W., Metchev, S. A., Larkin, J. E., et al. 2007, *ApJ*, 656, 505
 Mesa, D., Gratton, R., Berton, A., et al. 2011, *A&A*, 529, 131
 Metchev, S., Marois, C., & Zuckerman, B. 2009, *ApJ*, 705, L204
 Mugnier, L. M., Cornia, A., Sauvage, J., et al. 2009, *J. Opt. Soc. Am. A*, 26, 1326
 Oppenheimer, B. R., Kulkarni, S. R., Matthews, K., & Nakajima, T. 1995, *Science*, 270, 1478
 Patience, J., King, R. R., de Rosa, R. J., & Marois, C. 2010, *A&A*, 517, A76
 Perrin, M. D., Sivaramakrishnan, A., Makidon, R. B., Oppenheimer, B. R., & Graham, J. R. 2003, *ApJ*, 596, 702
 Pueyo, L., Kay, J., Kasdin, N. J., et al. 2009, *Appl. Opt.*, 48, 6296
 Racine, R., Walker, G. A. H., Nadeau, D., Doyon, R., & Marois, C. 1999, *PASP*, 111, 587
 Sauvage, J.-F., Fusco, T., Rousset, G., & Petit, C. 2007, *J. Opt. Soc. Am. A*, 24, 2334
 Serabyn, E., Mawet, D., & Burruss, R. 2010, *Nature*, 464, 1018
 Soummer, R. 2005, *ApJ*, 618, L161
 Soummer, R., & Ferrari, A. 2007, *ApJ*, 663, L49
 Soummer, R., Ferrari, A., Aime, C., & Jolissaint, L. 2007a, *ApJ*, 669, 642
 Soummer, R., Pueyo, L., Sivaramakrishnan, A., & Vanderbei, R. J. 2007b, *Opt. Express*, 15, 15935
 Sparks, W. B., & Ford, H. C. 2002, *ApJ*, 578, 543
 Thalmann, C., Carson, J., Janson, M., et al. 2009, *ApJ*, 707, L123
 Thatte, N., Abuter, R., Tecza, M., et al. 2007, *MNRAS*, 378, 1229
 Vigan, A., Langlois, M., Moutou, C., & Dohlen, K. 2008, *A&A*, 489, 1345
 Wallace, J. K., Burruss, R., Bartos, R. D., et al. 2010, *Proc. SPIE*, 7736, 77365
 Zimmerman, N., Brenner, D., Oppenheimer, B. R., et al. 2011, *PASP*, 123, 904
 Zimmerman, N., Oppenheimer, B. R., Hinkley, S., et al. 2010, *ApJ*, 709, 733