OPTIMIZED PRINCIPAL COMPONENT ANALYSIS ON CORONAGRAPHIC IMAGES OF THE FOMALHAUT SYSTEM[*]

Tiffany Meshkat[1], Matthew A. Kenworthy[1], Sascha P. Quanz[2], and Adam Amara[2]
[1] Sterrewacht Leiden, P.O. Box 9513, Niels Bohrweg 2, 2300-RA Leiden, The Netherlands
[2] Institute for Astronomy, ETH Zurich, Wolfgang-Pauli-Strasse 27, 8093-CH Zurich, Switzerland

ABSTRACT

We present the results of a study to optimize the principal component analysis (PCA) algorithm for planet detection, a new algorithm complementing angular differential imaging and locally optimized combination of images (LOCI) for increasing the contrast achievable next to a bright star. The stellar point spread function (PSF) is constructed by removing linear combinations of principal components, allowing the flux from an extrasolar planet to shine through. The number of principal components used determines how well the stellar PSF is globally modeled. Using more principal components may decrease the number of speckles in the final image, but also increases the background noise. We apply PCA to Fomalhaut Very Large Telescope NaCo images acquired at $4.05\,\mu$m with an apodized phase plate. We do not detect any companions, with a model dependent upper mass limit of 13–18 $M_{\rm Jup}$ from 4–10 AU. PCA achieves greater sensitivity than the LOCI algorithm for the Fomalhaut coronagraphic data by up to 1 mag. We make several adaptations to the PCA code and determine which of these prove the most effective at maximizing the signal-to-noise from a planet very close to its parent star. We demonstrate that optimizing the number of principal components used in PCA proves most effective for pulling out a planet signal.

Key words: methods: data analysis – planets and satellites: detection – stars: individual (Fomalhaut) – techniques: image processing

Online-only material: color figures

1. INTRODUCTION

The detection and characterization of extrasolar planets has grown dramatically as a field since the first detection in 1992 (Wolszczan & Frail 1992). The most successful detection techniques thus far are radial velocity (RV) and transit detection. Using ground and space based surveys (HARPS, *Kepler*, COROT, etc.), these indirect techniques have discovered over 800 planets (exoplanet.eu) as well as thousands more planet candidates.

The direct detection of planets provides a unique opportunity to study exoplanets in the context of their formation and evolution. It complements the underlying semi-major axis exoplanet distribution from RV surveys (from 100 AU down to a few AUs) and enables the characterization of the planet itself with an examination of its emergent flux as a function of wavelength. The detection of the planets HR8799 bcde (Marois et al. 2008), Fomalhaut b (Kalas et al. 2008), β Pic b (Lagrange et al. 2009), 2MASS1207 (Chauvin et al. 2004), 1RXS J1609−2105 b (Lafrenière et al. 2008), HD 95086 b (Rameau et al. 2013b), KOI-94 (Takahashi et al. 2013) as well as discoveries of protoplanetary candidates LkCa 15 b (Kraus & Ireland 2012) and HD100546 b (Quanz et al. 2013), demonstrate the potential breakthroughs of the technique. However, thus far, most dedicated high contrast imaging surveys have yielded null results (e.g., Rameau et al. 2013a; Vigan et al. 2012; Chauvin et al. 2010; Biller et al. 2007; Heinze et al. 2008). These null results are due to the lack of contrast at small orbital separations, where most planets are expected to be found. Since planets are concluded to be rare at large orbital separations (Chauvin et al. 2010; Lafrenière et al. 2007), high contrast imaging must probe close to the parent star to detect a planet.

High contrast imaging is limited by the diffraction limit, set by the telescope optics, which determines the minimum angular separation achievable under ideal conditions. Since planets are low mass, cold, and red compared to their parent star (Spiegel & Burrows 2012; Baraffe et al. 2003), the contrast ratio of their magnitudes is an additional constraint on their detectability. New instruments and techniques have been developed to combat these constraints at the acquisition and image processing stage.

Coronagraphs have been developed to reduce the light scattered in the telescope optics from diffraction during acquisition, but at a cost of throughput and angular resolution (Guyon et al. 2005). Coronagraphic optics allow us to probe smaller inner working angles, but are limited by the stellar "speckles" which can dominate the flux from a planet (Hinkley et al. 2009).

By turning off the telescope derotator on an alt-az telescope, the planet is able to "rotate" around the star, while the stellar PSF stays relatively stable and the speckles vary randomly in time. This technique is used in angular differential imaging (ADI; Marois et al. 2008). It takes advantage of this rotation to identify and subtract (in post-processing) the contribution from the stellar PSF and speckles. There are a number of image processing techniques aimed at modeling and subtracting the stellar PSF from every image, allowing the sky fixed planet signal to shine through. Locally optimized combination of images (LOCI; Lafrenière et al. 2007) is an extension of ADI, which models the local stellar PSF structure in every image. Principal component analysis (PCA; Amara & Quanz 2012; Soummer et al. 2012; Brandt et al. 2013) models how the PSF varies in time by identifying the main linear components of the variation. Application of these image processing techniques has been demonstrated to increase the limiting magnitude achievable by up to a factor of five (Lafrenière et al. 2007; Amara & Quanz 2012).

In this work, we present a detailed study of LOCI and PCA image processing techniques in order to optimize the
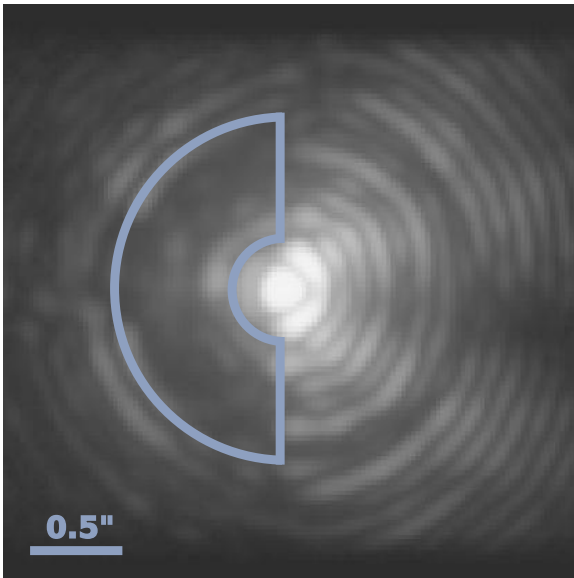
Figure 1. Image demonstrating the APP airy diffraction pattern with the diffraction suppressed region outlined in blue. This is the only region that is used in the data reduction.

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

signal-to-noise ratio (S/N) of a planet at small angular separations with the apodizing phase plate (APP) coronagraph (Kenworthy et al. 2010, 2007). We compare our results to the previous result of Kenworthy et al. (2013).

The Fomalhaut dataset that is used in the following analyses are a deep but typical observing sequence and will act as an example for the rest of our surveys.

2. DATA

Data were obtained of the star Fomalhaut at the Very Large Telescope (VLT)/UT4 with NaCo (Lenzen et al. 2003; Rousset et al. 2003) in 2011 July and August (087.C–0701(B)) and were analyzed and published in Kenworthy et al. (2013). Fomalhaut was used as the natural guide star with the visible band wavefront sensor. The L27 camera on NaCo was used with the NB4.05 filter ($\lambda = 4.051\,\mu$m and $\Delta\lambda = 0.02\,\mu$m) and the APP coronagraph (Kenworthy et al. 2010; Quanz et al. 2010) to provide additional diffraction suppression. We used pupil tracking mode to perform ADI (Marois et al. 2006). The PSF core is intentionally saturated to increase the signal from any potential companions.

The APP provides diffraction suppression over a $180°$ wedge on one side of the target (Figure 1). Additional observations are required with a different position angle (P.A.) to cover the full $360°$ around the star. For these observations, we have three different datasets with different P.A.s ensuring full P.A. coverage around the target star. Each dataset has a large amount of field rotation: $119°$, $117°$, $120°$.

Data were acquired in cube mode. Each data cube contains 200 frames, each with an integration time of 0.23 s. Approximately 70 cubes were obtained for each hemisphere dataset, totaling in an integration time of 160 minutes. A three point dither pattern was used to allow subtraction of the sky background and detector systematics as detailed in Kenworthy et al. (2013). Unsaturated short exposure data with the neutral density filter were also taken for photometry.

3. CREATING THE SIMULATED DATASETS

Data cubes at each dither position were pairwise subtracted to remove the sky background and detector systematics. The cubes were shifted to move the core PSF into the middle of a square image and bad frames (open loop and poor adaptive optics (AO)) were discarded (5% of hemisphere 1, 16% of hemisphere 2, and 7% of hemisphere 3). The three different APP P.A. datasets were processed separately. Each hemisphere dataset has its own corresponding unsaturated data for photometry.

Fake planets are subsequently used to determine the limiting contrast after image processing. The unsaturated Fomalhaut data is used to add a fake planet in each saturated frame. One fake planet is added at a time, between $0{.}''2$ and $1{.}''0$ in steps of $0{.}''1$ and δ magnitudes in steps of 1 mag from $dM = 7$–13. Due to the asymmetric nature of the APP PSF, it is also necessary to determine the S/N of a planet at different P.A.s. For our analysis we placed a planet on opposite sides of the star (P.A. $= 45°$ and $225°$ relative to the sky) to take into account the asymmetric PSF of the APP. These two P.A. orientations ensure that the planet is on the dark side of the APP in at least two of the hemispheres at once. The mean of the limiting contrast at each P.A. is stored.

The final science frames are processed with several different algorithms to recover the fake planet signal. All of the algorithms take advantage of the fake planet's rotation in the sky around the star to model and subtract the stellar PSF from each image (ADI; Marois et al. 2006). Before each algorithm is applied, the innermost region is masked out ($r < 0{.}''15$) where the star has saturated the image and no planet could be detected. The method of modeling the stellar PSF differs between the algorithms, detailed in the following subsections.

One metric for detectability of planets is S/N. It is a measure of the detectability of a point source, assuming the noise is Gaussian and decorrelated between diffraction limited elements at that radius. The equation below is similar to those in the literature, describing local S/N:

$$\left(\frac{\text{S}}{\text{N}}\right)_{\text{planet}} = \frac{F_{\text{planet}}}{\sigma(r)\sqrt{\pi r_{\text{ap}}^2}},$$

where F_{planet} is the sum of the planet flux in an aperture with radius $r_{\text{ap}} = 3$ pixels and σ is the root mean square of the pixels in a $180°$, 6 pixel wide arc at the same radius, surrounding the star.

The equation above assumes statistically independent pixels, which in the case of speckle noise limited regimes is typically not be the case. For the sake of consistency with other papers in the literature, we use one of the most common definitions of S/N calculation to facilitate comparison with other methods. This is a widely acknowledged issue in this research field, so while the S/N quoted may be off by a scaling factor, the conclusions in this paper do not rely on the absolute scaling as we are comparing analysis techniques.

4. DATA ANALYSIS

Coadding the frames in a data cube is a common practice but the best number of frames to coadd has not yet been thoroughly studied. We experimented with different numbers of coadded frames using fake planets. We ran the data through our PCA pipeline (detailed in Section 4.2) with different numbers of frames coadded (Figure 2). For example, 100 frames coadded means that twice as many images are passed to our pipeline as in the 200 coadded frames case. Figure 2 shows four S/N
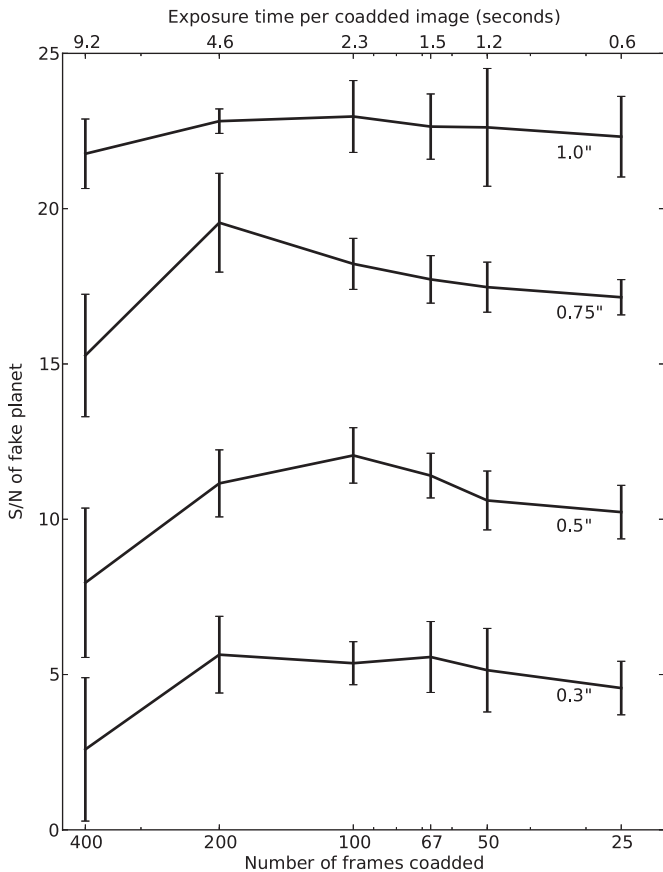
Figure 2. Signal-to-noise ratio curves for a fake planet at four angular separations, with different amounts of frames coadded. At each angular separation, planets were added at four position angles and averaged. The error bars are 1σ. The number of PCs is fixed at 20. Each contrast curve is offset from a S/N of 10 for clarity.

curves for planets injected a different angular separations, each with a S/N of approximately 10. The curves are offset from 10 for clarity. Coadding 200 or less frames yields a higher S/N. However, the S/N varies by less than a factor of two over all coadds, making this a relatively small effect. For the following analysis, we keep the coadds fixed at 200 frames, which yields S/N as good as less coadds, but is computationally much faster. This corresponds to ~70 coadded images in each hemisphere which are passed to our pipeline. Since there is little field rotation between individual frames in a data cube, the smearing effect within a cube is negligible.

4.1. LOCI

Locally optimized combination of images (Lafrenière et al. 2007) is a widely used planet detection algorithm which spatially models the stellar PSF to remove speckles. An image is divided into rings, which are subdivided into wedges. An optimal, linear combination of images subtracts speckles within that region. The least squares fit succeeds at minimizing speckles, but also reduces the planet flux through the subtraction for small angular separations.

Each hemisphere dataset is processed with LOCI independently and the final three hemisphere sky aligned cubes are collapsed. Since we are using the APP, we only perform LOCI on the "dark side" of the image frames. This $180°$ D shaped region (inner $= 2\lambda/D$, outer $= 7\lambda/D$) is the only part of each frame that is coadded in the final image.

Kenworthy et al. (2013) analysis of these data used the LOCI algorithm. Monte Carlo simulations exploring LOCI parameters ensured that this is the best sensitivity LOCI could produce.

4.2. Principal Component Analysis

Principal component analysis is a mathematical technique that relies on the assumption that every image in a stack can be represented as a linear combination of its principal orthogonal components, selecting structures that are present in most of the images. Its recent application to high contrast exoplanet imaging (Amara & Quanz 2012; Soummer et al. 2012) has been shown to be very effective. Unlike LOCI (Lafrenière et al. 2007) which models the local stellar PSF structure, PCA models the global PSF structure.

The full stack of images with sky rotation is used for PCA. However, since we are using the APP, only the "dark side" of each image is used in the fit. The S/N from a fake planet is lower if we include the "bright side." We follow the description of PCA outlined in Amara & Quanz (2012) for the following analysis.

The number of principal components (PCs) used determines how well the stellar PSF is fit. The first few components are the most stable, have less noise, and contain the most common structure in all the images. For our default analysis, we used 20 PCs to model the stellar PSF. PCA is run on each hemisphere dataset independently, as the PCs are correlated with time. The final de-rotated frames are coadded into one final image covering the full $360°$ around the star.

The following subsections discuss self-subtraction due to the PCA algorithm as well as a series of modifications we performed on PCA to optimize the detection of a planet at small λ/D.

4.2.1. Self-subtraction

Self-subtraction due the LOCI algorithm has been well documented by previous authors (Lafrenière et al. 2007; Marois et al. 2010), but its impact on PCA is not yet well studied. The LOCI algorithm requires that the frames nearest in time to the current frame are not considered in the least squares fit, thus limiting the self-subtraction a potential planet. However, this frame rejection technique does not completely account for flux loss from a planet.

For our PCA analysis, we draw a distinction between two types of modes: detection and characterization. Characterization mode requires fully accounting for flux loss of the planet as a function of number of PCs as we map between the measured flux and the calibrated estimate of "true" flux. However, in detection mode, since we only care about our ability to separate the planet signal from the background noise, the main issue is the flux loss relative to the separation of the background noise. In this paper, we address simply the detection mode.

Figure 3 shows three plots with the flux ratio and S/N versus the number of PCs it was processed with. The top figure is for a planet injected at $1''.0$, middle is at $0''.5$ and bottom is at $0''.3$. The "flux ratio" is the ratio of the injected planet flux to the PCA processed flux in a 4 pixel aperture. For each angular separation, the L' contrast which yields a S/N of approximately 10 is plotted. This figure demonstrates that the PCA method is more efficient at capturing the patterns associated with the background fluctuations of the field than capturing information associated with the planet translation. This differential effect means that in detection mode, it is acceptable for the flux ratio to decrease as long as the noise is decreasing as or more rapidly.
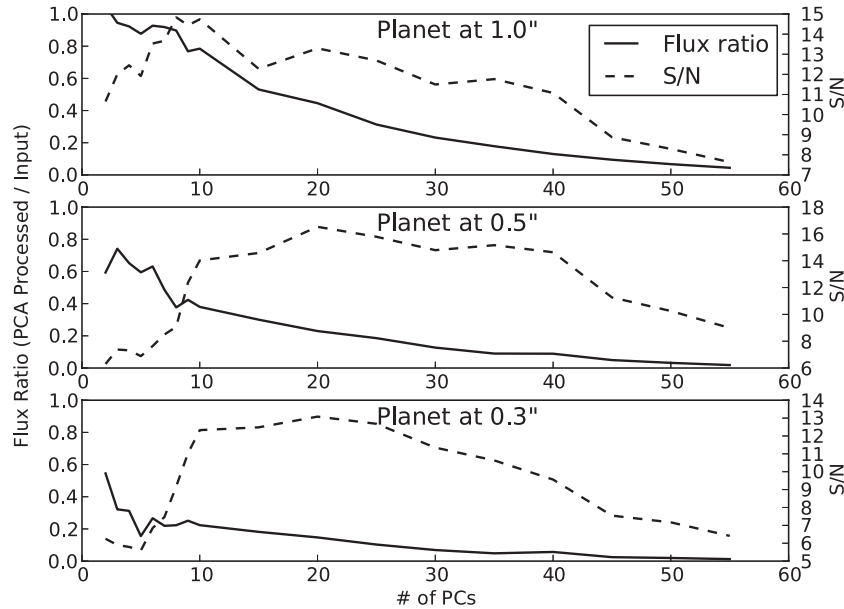
Figure 3. Comparison of the flux ratio and S/N based on number of PCs at different radii. The top panel is for a fake planet at $1''0$, the middle panel is for $0''5$ and the bottom panel is for $0''3$. These panels demonstrate that, while the flux ratio does decrease with PCs, the S/N follows a different curve.

4.2.2. PCA Modifications

1. *Frame rejection.* For our PCA code detailed above, all the frames are used in the fit and none are rejected. This was done under the assumption that self-subtraction of the planet happens less rapidly than the noise subtraction when using PCA. As discussed in Section 4.2.1, while we are in "detection mode" of direct exoplanet detection, the important factor to consider is the S/N rather than planet flux. To test this, we used only a subset of the frames to determine the PCs. The frames nearest in time to the frame being fitted were rejected. These are frames where a potential planet would overlap by 0.5 FWHM or more. The number of frames to reject depends on the separation of the planet from the star. The total rotation of the planet is limited by the amount of sky rotation achieved during each dataset. A planet very far from its parent star would appear to rotate faster between frames, thus less frames need to be rejected. This test allows us to compare the S/N of a fake planet processed with standard PCA and "0.5 FWHM rejection," where we mimic the routine in LOCI to reject the frames closest in time.

2. *Radius limited.* Next, we modified the PCA basis set by only using the image out to a certain radius. The outer radius (R_{out}) passed to the PCA code determines the amount of information provided to the singular value decomposition (SVD) algorithm. Extra information does not necessarily provide a better fit. Our previous applications of PCA kept R_{out} fixed. The information passed to the SVD algorithm should be directly related to the stellar PSF. We modified our PCA code to vary R_{out} based on the location of the fake planet. The new R_{out} is $1 \lambda/D$ greater than the radius of the fake planet (see Figure 4), thus performing PCA on a smaller region. This experiment was performed to test how significant the stellar PSF fit was affected by radii greater than the planet location.

3. *Number of PCs.* The main parameter which can be manipulated in PCA is the number of PCs used in the SVD fit. The first principal value (the highest singular value in the
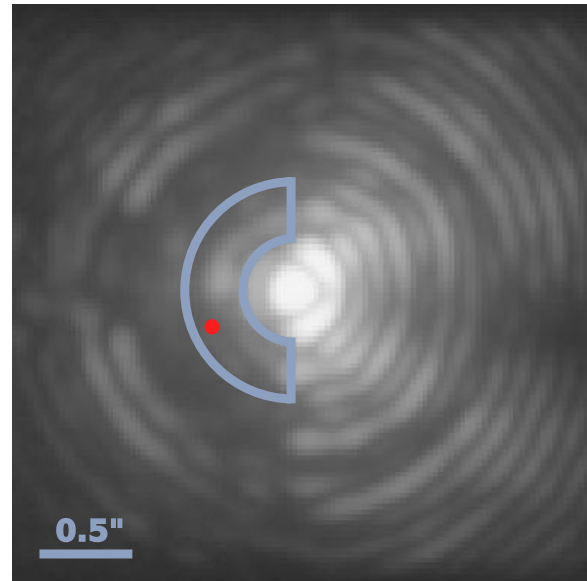

Figure 4. Image demonstrating the APP airy diffraction pattern with the radius limited region outlined in blue. This is the only region that is used in the data reduction.

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

diagonal matrix) is the "variance" of the image stack from the mean, in the direction of the first PC. The same is true about the second principal value and so on.

Figure 5 shows the PCA coefficient values in descending order for one of our datasets. The first few PCA coefficient values are significantly greater than the later values, implying that those PCs contain the most dominant features. Increasing the number of PCs in the stellar PSF fit can help bring out the planet signal by removing structure, however it also can add noise. Determining the optimal number of PCs for a certain stellar PSF fit is an essential but expensive task. The optimal number of PCs depends on the time variability of complex speckles.
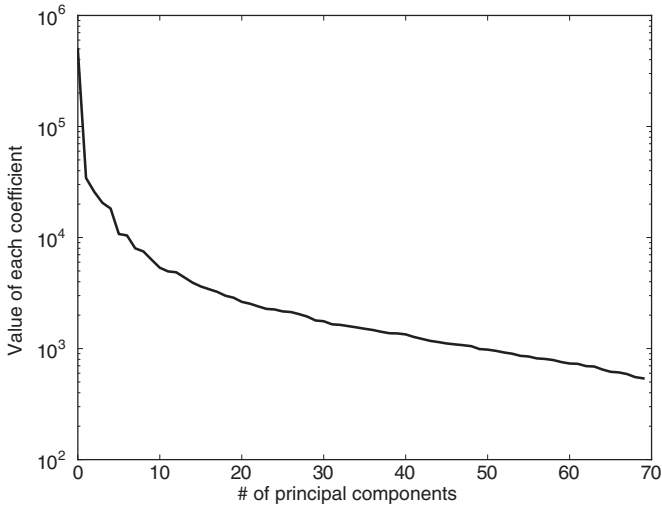
Figure 5. Plot of the PCA coefficient values. The highest PCA coefficient value corresponds to the most significant PC.
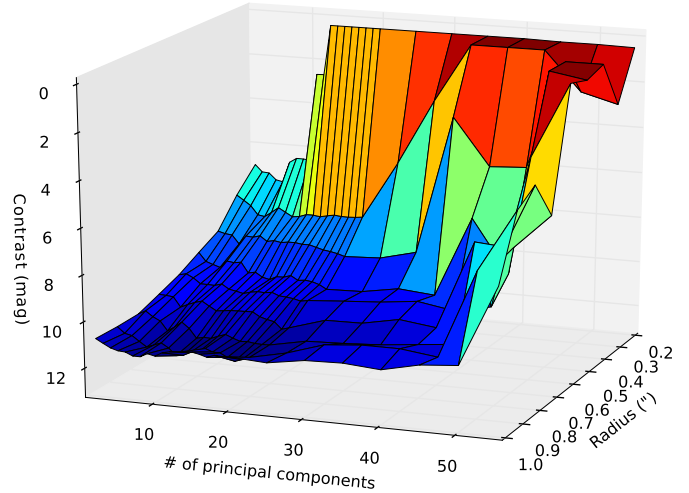

Figure 6. Contrast curves for a 7σ detection of a point source in our Fomalhaut APP data processed with LOCI, ADI, and variations of ADI. The LOCI curve is adapted from Kenworthy et al. (2013) to a 7σ detection. The numbers on the yellow curve signify the number of PCs which yield the highest S/N at that radius. The dashed line is the background limit. The PCA contrast curves are the mean value for fake planets inserted at two P.A.s on opposite sides of the star (P.A. = $45°$ and $225°$).

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

For each dataset and fake planet angular separation, PCA was run with different numbers of PCs ranging from 5 to 60, in increments of 5.

5. RESULTS AND DISCUSSION

Figure 6 shows the results of each image processing method detailed in Section 4.2. Each technique was run with varying planet contrasts at a given radius. We extrapolated between planet contrasts to determine contrast that yields a S/N of 7. For the method with varying PCs detailed in Section 4.2.2, we noted which number of PCs yielded the highest S/N at which radius. These are the numbers listed on the yellow curve in Figure 6.

Our standard PCA technique yields a better contrast curve than LOCI for our coronagraphic data. Our modifications to PCA, in some cases, yield better sensitivity.

Figure 7. Three-dimensional surface of the contrast achieved in a 7σ detection with varied numbers of PCs. Varying the number of PCs at small angular separations affects the 7σ detection limit by up to 8 mag. Beyond $0\farcs6$, the number of PCs used is less significant.

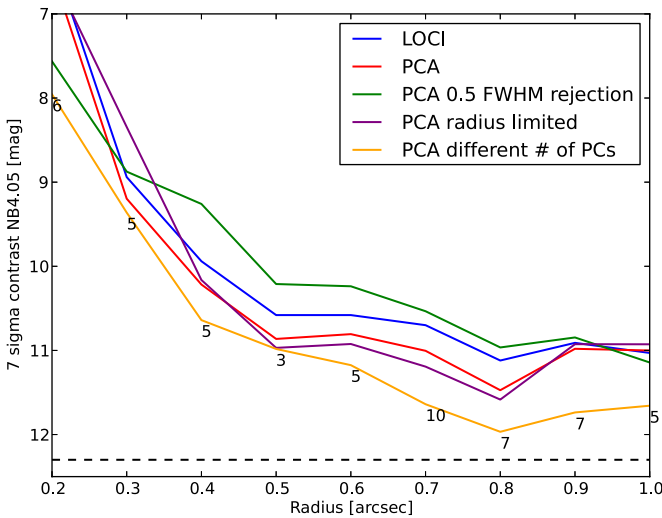(A color version of this figure is available in the online journal.)

Unlike the LOCI algorithm, rejecting the frames nearest in time (detailed in Section 4.2.2) yields a worse contrast curve than our standard PCA. This is likely due to the noise being more correlated in frames closer in time, thus providing important information to the SVD algorithm and increasing the S/N of the planet. We did not reject any frames in our final data analysis approach.

Limiting the outer radius passed to the SVD algorithm yielded a slightly better contrast ratio than standard PCA from $0\farcs5$ to $0\farcs8$. However, this contrast increase is not significant and is only beneficial because it is less computationally expensive.

Our standard PCA contrast curve was generated with 20 PCs. By varying the number of PCs we can increase the S/N from a companion. Our PC-varying result yields a consistently more sensitive contrast curve then all the other methods. We gain between 0.5 and 1 mag contrast over our LOCI analysis from $0\farcs2$ to $1\farcs0$. From Figure 6 we see that the number of PCs which yield the highest S/N for a planet varies based on its angular separation.

Figure 7 is a three-dimensional (3D) surface plot showing how the number of PCs at each radius affects the contrast at 7σ for a planet at a fixed P.A. Fake planets were added between 2 and 20 PCs in smaller steps to emphasize the structure. This figure demonstrates that at small angular separations ($<0\farcs6$), the S/N is sensitive to the number of PCs chosen. This is the region where the diffraction and speckles due to the star are more significant than the unstructured noise from thermal emission and the sky background. For example, at $0\farcs2$ choosing a small number of PCs yields an 8 mag gain in sensitivity than a large number of PCs. Increasing the number of PCs quickly leads to nearly complete self-subtraction. This can be seen in Figure 7 as a contrast of nearly zero. As we move to larger radii the optimal number of PCs remains in the 5–20 PC range. Beyond $0\farcs6$ where the number of PCs shows no significant preference below 45 PCs.

5.1. Comparison with Kenworthy et al. (2013)

Our PCA re-analysis of these data improves sensitivity at small inner working angles, from $0\farcs2$ to $1''$, in some cases
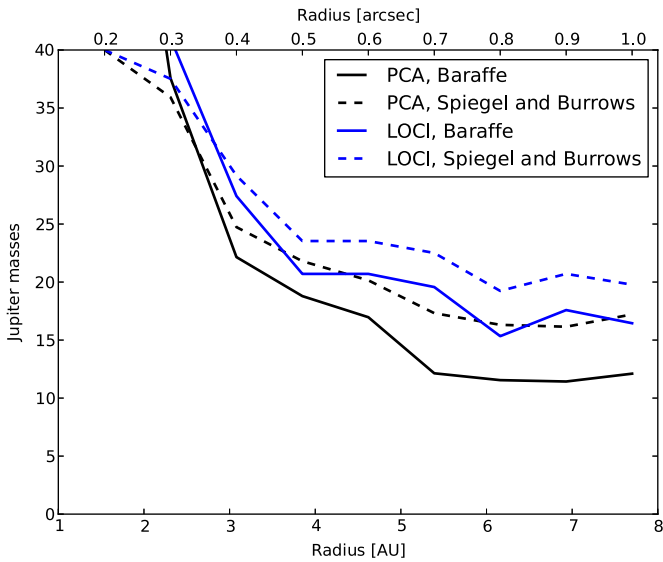
Figure 8. Detection limit for fake companions around Fomalhaut generated with PCA (black lines) and LOCI (blue lines, converted to 7σ detection from Kenworthy et al. 2013) using Baraffe et al. (2003; solid lines) and Spiegel & Burrows (2012; dashed lines).

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

by 1 mag (see blue and yellow curves, Figure 6). We convert the best 7σ detection contrast curve to an upper mass limit for planets using the Baraffe et al. (2003) and Spiegel & Burrows (2012) atmospheric models (Figure 8) assuming an age of 440 Myr (Mamajek et al. 2012). We confirm the non-detection of companions with model-dependent upper mass limit of 13–18 $M_{\rm Jup}$ from 4–10 AU. Our new upper mass limit is based on our more robust 7σ detection limit. The 1 mag increase in the contrast ratio at $0''.5$ translates to an increased sensitivity of $\Delta 7 M_{\rm Jup}$. The increase in sensitivity allows us to probe planetary masses ($<15 M_{\rm Jup}$) at small angular separations.

5.2. Fainter Fomalhaut

We have shown that the number of PCs which yield the highest S/N depends on the planet's distance from the parent star (yellow line, Figure 6). At small angular separations ($<0''.6$), the S/N is sensitive to the number of PCs chosen (Figure 7). This is the limit where the diffraction from the central star is equal to or less significant than the background noise.

We add Gaussian white noise to our data to test if this turnover point changes for a fainter target (Figure 9; 3D surface noise). Increasing the sky background noise makes Fomalhaut 1.5 mag fainter, while keeping the telescope conditions and Strehl identical. This is the ideal way to test how fainter targets will behave. Fake planets are once again injected and the best number of PCs at each angular separation is noted.

Changing the number of PCs used at each angular separation is still the best method for detecting companions. As expected, the regime of large numbers of PCs at small separations results in low contrast, which then improves down to a plateau at smaller PCs and larger radii. The turnover point remains near $0''.6$, beyond which the diffraction from the star is no longer significant and the optimal number of PCs is less clear. Beyond this separation, the background noise dominates the SVD fit and thus does not help subtract the stellar PSF.
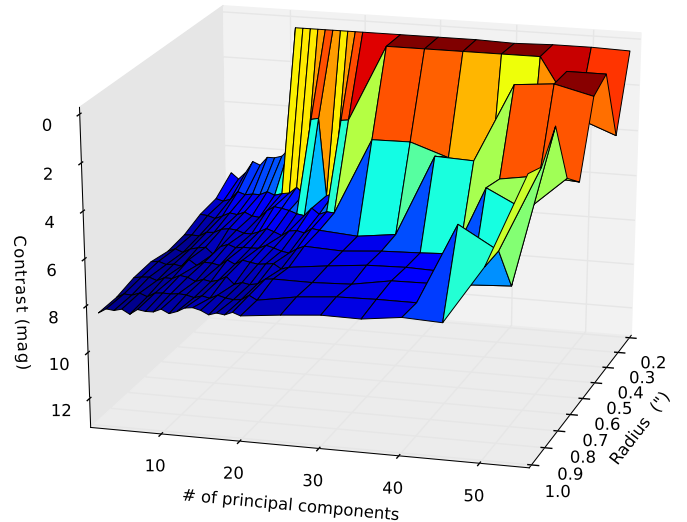

Figure 9. Similar to the 3D surface in Figure 7, but with Gaussian white noise added to the data. The resulting star is 1.5 mag fainter than Fomalhaut.

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

6. CONCLUSION

We re-analyze our Fomalhaut APP/NaCo/NB4.05 data using PCA and compare it with the LOCI algorithm. PCA yields a more sensitive contrast curve than the LOCI algorithm at small inner working angles. We tested several modifications to PCA and gain up to 1 mag of contrast over our LOCI analysis from $0''.2$ to $1''.0$. The most effective parameter which optimized PCA was varying the number of principal components. The number of principal components chosen is sensitive for planets at small inner working angles. The detection limit of a planet at small radii can vary by several magnitudes. Careful attention should be paid to determining the number of principal components used at radii where the speckles are more significant than the unstructured noise of thermal emission and the sky background. Running PCA for a range of principal components at each angular separation and generating a 3D surface is a useful way to visualize the optimal number of principal components needed to pull out a faint planet signal.

Further analysis is needed in other wavelengths, as differing Strehl ratios may affect the turnover point where the stellar diffraction is less significant than the background noise. These results have direct application for current and future planet imaging campaigns, which will likely use a combination of PCA, LOCI, and other image processing techniques.

REFERENCES

Amara, A., & Quanz, S. P. 2012, MNRAS, 427, 948
Baraffe, I., Chabrier, G., Barman, T. S., Allard, F., & Hauschildt, P. H. 2003, A&A, 402, 701
Biller, B. A., Close, L. M., Masciadri, E., et al. 2007, ApJS, 173, 143
Brandt, T. D., McElwain, M. W., Turner, E. L., et al. 2013, ApJ, 764, 183
Chauvin, G., Lagrange, A.-M., Bonavita, M., et al. 2010, A&A, 509, A52
Chauvin, G., Lagrange, A.-M., Dumas, C., et al. 2004, A&A, 425, L29
Guyon, O., Pluznhik, E. A., Galicher, R., et al. 2005, ApJ, 622, 744

Heinze, A. N., Hinz, P. M., Kenworthy, M., Miller, D., & Sivanandam, S. 2008, ApJ, 688, 583

Hinkley, S., Oppenheimer, B. R., Soummer, R., et al. 2009, ApJ, 701, 804

Kalas, P., Graham, J. R., Chiang, E., et al. 2008, Sci, 322, 1345

Kenworthy, M. A., Codona, J. L., Hinz, P. M., et al. 2007, ApJ, 660, 762

Kenworthy, M. A., Meshkat, T., Quanz, S. P., et al. 2013, ApJ, 764, 7

Kenworthy, M. A., Quanz, S. P., Meyer, M. R., et al. 2010, Proc. SPIE, 7735, 773532

Kraus, A. L., & Ireland, M. J. 2012, ApJ, 745, 5

Lafrenière, D., Doyon, R., Marois, C., et al. 2007, ApJ, 670, 1367

Lafrenière, D., Jayawardhana, R., & van Kerkwijk, M. H. 2008, ApJL, 689, L153

Lagrange, A.-M., Gratadour, D., Chauvin, G., et al. 2009, A&A, 493, L21

Lenzen, R., Hartung, M., Brandner, W., et al. 2003, Proc. SPIE, 4841, 944

Mamajek, E. E., Quillen, A. C., Pecaut, M. J., et al. 2012, AJ, 143, 72

Marois, C., Lafrenière, D., Doyon, R., Macintosh, B., & Nadeau, D. 2006, ApJ, 641, 556

Marois, C., Macintosh, B., Barman, T., et al. 2008, Sci, 322, 1348

Marois, C., Macintosh, B., & Véran, J.-P. 2010, Proc. SPIE, 7736, 77361J

Quanz, S. P., Amara, A., Meyer, M. R., et al. 2013, ApJL, 766, L1

Quanz, S. P., Meyer, M. R., Kenworthy, M. A., et al. 2010, ApJL, 722, L49

Rameau, J., Chauvin, G., Lagrange, A.-M., et al. 2013a, A&A, 553, A60

Rameau, J., Chauvin, G., Lagrange, A.-M., et al. 2013b, ApJL, 772, L15

Rousset, G., Lacombe, F., Puget, P., et al. 2003, Proc. SPIE, 4839, 140

Soummer, R., Pueyo, L., & Larkin, J. 2012, ApJL, 755, L28

Spiegel, D. S., & Burrows, A. 2012, ApJ, 745, 174

Takahashi, Y. H., Narita, N., Hirano, T., et al. 2013, arXiv:1309.2559

Vigan, A., Patience, J., Marois, C., et al. 2012, A&A, 544, A9

Wolszczan, A., & Frail, D. A. 1992, Natur, 355, 145