



Superresolution Reconstruction of Severely Undersampled Point-spread Functions Using Point-source Stacking and Deconvolution

Teresa Symons¹ , Michael Zemcov^{1,2} , James Bock^{2,3}, Yun-Ting Cheng³ , Brendan Crill² , Christopher Hirata⁴, and Stephanie Venuto¹

¹ Center for Detectors, Rochester Institute of Technology, 1 Lomb Memorial Drive, Rochester, NY 14623, USA; tas4514@rit.edu

² Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA

³ California Institute of Technology, 1200 East California Boulevard, Pasadena, CA 91125, USA

⁴ Center for Cosmology and AstroParticle Physics, The Ohio State University, 191 West Woodruff Avenue, Columbus, OH 43210, USA

Received 2020 July 16; revised 2020 November 3; accepted 2020 November 13; published 2021 February 1

Abstract

Point-spread function (PSF) estimation in spatially undersampled images is challenging because large pixels average fine-scale spatial information. This is problematic when fine-resolution details are necessary, as in optimal photometry where knowledge of the illumination pattern beyond the native spatial resolution of the image may be required. Here, we introduce a method of PSF reconstruction where point sources are artificially sampled beyond the native resolution of an image and combined together via stacking to return a finely sampled estimate of the PSF. This estimate is then deconvolved from the pixel-gridding function to return a superresolution kernel that can be used for optimally weighted photometry. We benchmark against the <1% photometric error requirement of the upcoming SPHEREx mission to assess performance in a concrete example. We find that standard methods like Richardson–Lucy deconvolution are not sufficient to achieve this stringent requirement. We investigate a more advanced method with significant heritage in image analysis called iterative back-projection (IBP) and demonstrate it using idealized Gaussian cases and simulated SPHEREx images. In testing this method on real images recorded by the LORRI instrument on New Horizons, we are able to identify systematic pointing drift. Our IBP-derived PSF kernels allow photometric accuracy significantly better than the requirement in individual SPHEREx exposures. This PSF reconstruction method is broadly applicable to a variety of problems and combines computationally simple techniques in a way that is robust to complicating factors such as severe undersampling, spatially complex PSFs, noise, crowded fields, or limited source numbers.

Unified Astronomy Thesaurus concepts: [Astrostatistics techniques \(1886\)](#); [Deconvolution \(1910\)](#); [Astronomy data analysis \(1858\)](#); [Computational methods \(1965\)](#)

1. Introduction

Whether explicitly or implicitly, astrophysicists require knowledge of how the brightness in an image relates to true brightness on the sky to interpret the shapes and intensities of observed sources of emission. The instrument response function of a telescope is the transfer function that maps between quantities measured at the focal plane to the physical intensity on the sky. For natively imaging array detectors, or images constructed from time-domain scans of the sky, this information is often called the point-spread function (PSF), which is defined to be the response of a focused imaging system to a point source (see Gai & Cancelliere 2007 for a review). PSF estimation is a foundational problem in astronomical image analysis and interpretation.

Measuring the PSF of an instrument can be challenging. For some instruments, it is possible to measure the PSF using collimated images of unresolved sources in the laboratory or on the sky. It is often possible to use geometric or physical optics to model the system and propagate an image of an unresolved source to the detector. More recently, forward-modeling approaches incorporating measurements of the PSF have been used to model complex electrical effects in detectors (e.g., Donlon et al. 2018). In general, astrophysicists have a preference for using images of unresolved⁵ objects to assess the PSF in an image. This method has the advantage of

measuring the as-built optical system, which not only captures the design performance of an optical system, but also any effects from scattered light, mechanical structures, or other nonidealities in the telescope or detector.

The intrinsic PSF of the optical system and the pixelization of the displayed image are not immutable quantities, but rather independent parameters fixed at either the instrument design, or for the case of time-domain maps, during the map-making process. This is illustrated in Figure 1, which shows the effects of changing the relative sizes of the pixelization and PSF. In practice, the choice of image pixelization depends on the sources under investigation and the desired properties of the final image. When possible, the pixel size is chosen to be slightly smaller than the intrinsic PSF, but not to “super-resolve” it. This is because of both the noise penalties associated with spreading a fixed flux into an increased number of pixels, and the diminishing fidelity of superresolved astronomical images. A resolution of two pixels per FWHM is often quoted as representative of Nyquist sampling, with a smaller number of pixels per FWHM indicating undersampling. If the optical PSF is significantly narrower than the width of a pixel, then the PSF is severely undersampled and the precise location of the source within the pixel may be difficult to recover (Puetter et al. 2005; Robertson 2017). It can be shown that the optimum signal-to-noise ratio (S/N) on unresolved sources is achieved by either concentrating as much of the optical PSF into a single detector element as possible, or by performing optimal photometry that uses

⁵ For our purposes, “unresolved” is the condition $\theta_{\text{source}} \ll \theta_{\text{PSF}}$, where θ_{source} is the width of the source and θ_{PSF} is the width of the PSF.




Figure 1. Examples of the relationship between pixelization and PSF. The top row shows the effect of changes in the pixel gridding of the input image, shown in panel A, while the bottom row shows the effect of changing the width of the PSF. In panel A, the pixelization is matched to the optical PSF, so that the $\text{FWHM} \sim 1$ pixel. In this case, the spatial resolution of the telescope dominates the spatial resolution of the image. In panel B, we show the case where $\theta_{\text{pix}}/\text{FWHM} \sim 5$, and the image spatial resolution is dominated by the pixel grid. Panel C shows the $\theta_{\text{pix}}/\text{FWHM} \sim 20$ case where the image spatial resolution is heavily gridding dominated. The method described here takes advantage of the fact that the PSF is sampled in many different ways with respect to the pixel grid to allow the reconstruction of the subpixel PSF shape. In the bottom row, we show examples of $\text{FWHM}_{\text{PSF}}/\theta_{\text{pix}} = \{2, 5, 20\}$ in panels D, E, and F, respectively. In these cases, the subpixel PSF can be easily measured from point-like sources, and the method described here offers no improvement.

knowledge of the PSF to weight pixels according to their flux contributions (e.g., Horne 1986; Naylor 1998). The need for accurate photometry of sources is a strong motivation to improve our knowledge of the spatial properties of the PSF, sometimes beyond that offered by the native image pixelization.

Though the optical PSF may not be resolved by the image pixelization, information about its shape must be retained in any ensemble of images of unresolved sources that are placed at random locations with respect to the pixel grid. In this situation, the images of the PSF are sampled by the pixel grid slightly differently for each member of the set, so with a sufficient number of draws it should be possible to recombine the information to reconstruct the underlying “unpixelized” optical PSF. This concept has been successfully used to reconstruct superresolved PSFs in the past using a method commonly referred to as “stacking” (Dole et al. 2006; Béthermin et al. 2012; Zemcov et al. 2014). In short, the stacking method is a measurement of the covariance between a catalog of sources and their corresponding images in a map. Superresolution⁶ methods such as stacking can offer a way to recover Nyquist sampling of the PSF (Starck et al. 2002). The stacking method (Cady & Bates 1980) was first used in this context to perform statistical studies on low-resolution, confused images of the cosmic infrared background (e.g., Dole et al. 2006; Marsden et al. 2009, and others), but has since been extended to a wide range of wavelengths and scientific topics (see, e.g., Viero et al. 2013).

⁶ We define superresolution to mean any resolution higher than an image’s native resolution.

In this paper, we study PSF reconstruction and modeling using the stacking technique. We demonstrate that, in the limit where the input catalog is composed of point sources, the image resulting from a stacking process is that of the intrinsic PSF convolved with the pixelization grid (the effective PSF), and that the pixel grid can be reliably deconvolved to retrieve the instrument’s intrinsic optical PSF. As a case study, we examine this PSF reconstruction algorithm as it may apply to the Spectro-Photometer for the History of the universe, Epoch of Reionization, and ices Explorer (SPHEREx), an upcoming spectroscopic survey mission designed to image every 6.2×6.2 arcsec² pixel over the entire sky for $0.75 < \lambda < 4.8 \mu\text{m}$. Accurate photometry with methodological uncertainties $< 1\%$ is required to help attain SPHEREx’s ambitious scientific goals, but the PSF is intentionally designed to be underresolved by the pixel grid with an optical FWHM of $2''.01$ at $0.75 \mu\text{m}$, which is consistent with Korngut et al. (2018) when manufacturing tolerance and aberrations are included. This gives a sampling rate of ~ 0.3 pixels per FWHM, well below the Nyquist rate. It is therefore necessary to construct an algorithm that provides accurate PSF kernels for the precision optimal photometry. Because of SPHEREx’s severe undersampling, PSF reconstruction techniques used in other cases where accurate spatial knowledge of the PSF is equally important, such as for the study of weak lensing, cannot be used here (Bertin 2011; Rowe et al. 2011; Schmitz et al. 2020). These methods typically require applying subpixel shifts and thus some form of interpolation, but in our method, all shifts are in the form of integer values due to the large upsampling factor. The proposed approach also offers the benefit of only requiring a single exposure instead of combining multiple exposures to reconstruct the PSF and is less computationally expensive (Seshadri et al. 2013). The

Table 1
Important Variables Used in the Text

Variable	Description
<i>Image quantities, native pixelization</i>	
M	sky image
Δ	noise in M
P_{true}	underlying optical PSF of system
P_{grid}	pixel grid function
P	the PSF of sources in M , convolution of P_{true} and P_{grid}
<i>Image quantities, super pixelization</i>	
r	pixel grid scaling factor
\tilde{M}	sky image M scaled up by r
\tilde{P}	PSF of sources in M scaled up by r
<i>Input source catalog quantities</i>	
α	list of sources in image M
F	known flux of sources in M
(x, y)	known positions of sources in M
N	number of sources in M
<i>PSF reconstruction quantities</i>	
T	thumbnail cutout of a source centered on its known position
P_S	stacked PSF
P_{RLD}	deconvolution of P_S via RLD
P_{IBP}	deconvolution of P_S via IBP, best estimate for P_{true}
I_{RLD}	number of RLD iterations
I_{IBP}	number of IBP iterations
<i>Optimal photometry quantities</i>	
P_P	recentered, rescaled version of P_{IBP} used in optimal photometry
W_P	weight function used for optimal photometry
F_P	flux of a source calculated from optimal photometry
$\langle \delta F \rangle$	average deviation of flux from its known value
$\sigma_{\delta F}$	PSF reconstruction figure of merit

SPHEREx case illustrates the utility of our method in a situation where the intrinsic PSF is unresolved in individual images of point sources but can be reconstructed from a large ensemble of measurements contained within a single exposure. In Section 2, we introduce the PSF stacking concept in detail, various methods of deconvolution necessary for reconstructing the PSF in complex cases, as well as the method by which the PSF reconstruction can be used as a weight function in calculating photometry. In Section 3, we apply these methods to example cases including simulated SPHEREx PSFs as well as real data from the Long Range Reconnaissance Imager (LORRI) on New Horizons (Cheng et al. 2008). In Section 4, we put our work in the context of generalized PSF reconstruction techniques. Table 1 provides a glossary of mathematical variables used throughout the text.

2. Methods

2.1. Stacking

A simple way to understand the relationship between the PSF and the pixel grid function is through the convolution theorem. The act of signal detection with a pixelized system is equivalent to a convolution of the true underlying optical PSF,

P_{true} , with the pixel response P_{grid} :

$$P = P_{\text{true}} * P_{\text{grid}} = \mathcal{F}^{-1}[\mathcal{F}[P_{\text{true}}] \cdot \mathcal{F}[P_{\text{grid}}]], \quad (1)$$

where \mathcal{F} represents the Fourier transform and P is the PSF that can be observed in the image itself (Starck et al. 2002; Puetter et al. 2005), sometimes called the effective PSF (Lauer 1999). Figure 2 illustrates the relationship between these quantities.

In our formulation, the quantity P_{true} is the combination of any effects that contribute to the overall PSF measured in an image. These include the underlying optical PSF, how the optical PSF is played around pixels due to pointing jitter, and any electrical effects like interpixel capacitance or the “brighter-fatter” effect (Hirata & Choi 2020). For a linear detector, the quantity P_{true} incorporates all of these effects as a convolution over the history of the exposure. It is commonly noted that the observer only ever has access to P , so it is expedient to perform analysis with this quantity. But in the presence of nonlinear effects that are known to exist in modern hybridized HgCdTe NIR detectors (Plazas et al. 2018; Hirata & Choi 2020), it may in fact be necessary to deconvolve P_{grid} or other effects from P . There are potentially many reasons and ways to do this; here, we investigate one that uses the fact that the covariance between a well-understood catalog of stars and sources in the image allows for superresolution reconstruction, which when combined with deconvolution methods can return the underlying optical PSF. This is particularly motivated by the fact that, in the presence of these complicating factors, P_{true} is the kernel required to perform optimal photometry.

The stacking method takes advantage of the fact that an “ideal” gridded image of the astronomical sky can be represented as

$$M_j = \Delta_j + \sum_i^{N_j} F_i, \quad (2)$$

where M_j is the brightness of the image M in pixel j , Δ_j is the noise in that pixel, N_j is the number of sources falling into pixel j , and F_i is the flux of source i in a list α of N total sources of emission (Marsden et al. 2009; Viero et al. 2013). To make progress with PSF estimation, we note that, by the argument above, instruments with non-point PSFs contribute flux into more than one pixel. We define $P_j(x_i, y_i)$ to be the beam-convolved and mean-subtracted shape of the PSF centered at some source position (x_i, y_i) , and then write Equation (2) as

$$M_j = \Delta_j + \sum_i^{N_j} F_i P_j(x_i, y_i). \quad (3)$$

This expression accounts for the fact that all sources can contribute to the intensity of pixel j , as the PSF spreads flux to neighboring pixels.

Our goal is to estimate the shape of the PSF P from the measured sky M . Because P has the same amplitude for each source, we can invert Equation (3) and solve

$$P = \sum_i^{N_j} \frac{M(x_i, y_i)}{F_i}, \quad (4)$$

where $M(x_i, y_i)$ is the image centered on the source position (x_i, y_i) , and we assume that the noise obeys $\langle \Delta \rangle = 0$ over the sum. Furthermore, as with simple stacking, we require the source positions to be uncorrelated so that contributions from sources

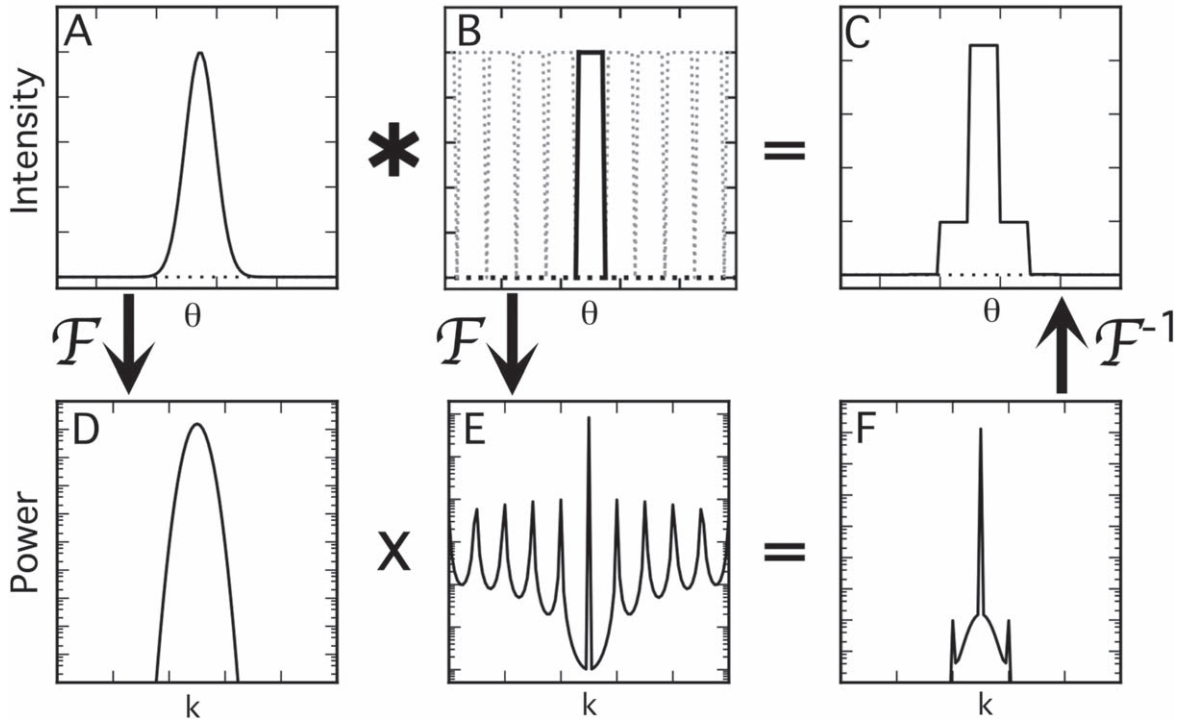


Figure 2. Relationship of the PSF and pixel-gridding function in both real space and Fourier space. Panel A shows a hypothetical PSF, which has a similar size to the pixel-gridding function (dotted line) shown in panel B. The spatial resolution of the image that results is computed as the convolution of the PSF and a single pixel (solid line in panel B), as shown in panel C. Flux from a point source is spread across several pixels according to the brightness in the PSF as sampled by the pixel-gridding function. As an alternative visualization, panel D shows the power spectrum (absolute square of the Fourier transform) of the PSF, and panel E shows the power spectrum of the pixel gridding function. By the convolution theorem, the inverse Fourier transform of their product, shown as a power spectrum in panel F, is also the final spatial quality of the image of a point source. These two visualizations are useful for understanding interactions between the PSF and pixel-gridding function.

in the image but not being stacked on do not add coherently (Marsden et al. 2009).

In this formalism, the superresolution recovery arises due to the nature of pixels, which have the property of averaging photons. We can always create larger pixels which contain more photons, but spread over a larger area, in such a way that the measured surface brightness is conserved. The corollary also holds when going from larger to smaller pixels. The cost of this regridding operation is changes in any fixed-amplitude noise: regridding larger pixels to smaller ones increases the noise per pixel, while smaller to larger decreases the noise in the larger pixel.

This averaging property of pixels allows us to write the relation

$$P_j = \frac{1}{r} \sum_i^r \tilde{P}_i^k, \quad (5)$$

where \tilde{P}^k is the image of the PSF sampled on a finer grid, and the scale factor between the areas of pixels k and j is r . Because the pixel size does not appear in Equation (4), we can write

$$\tilde{P} = \sum_i \frac{N_j}{F_i} \tilde{M}(x_i, y_i), \quad (6)$$

where \tilde{M} is a version of the image on the finer grid and is related to M through $\tilde{M}_j = (1/r) \sum_i^r \tilde{M}_i^k$. Any r such that the r -times finer image has significantly higher resolution than the native image resolution will increase the sampling rate and facilitate the stacking method; increasing r increases

the small-scale structure it would be possible to capture in the reconstructed PSF. However, increasing r can also lower the S/N or introduce greater numerical instability due to the larger number of subpixels per pixel. The fundamental cost of this superresolution stacking is that the noise per pixel is increased over the native pixel resolution. This can easily be addressed by making the list α large to compensate. As regridding conserves S/N in an area, the total S/N on the PSF will remain fixed for a fixed noise and number of sources in the stack.

To implement the proposed method in an unbiased manner, we require several conditions to be true:

1. The source positions are uncorrelated and sources are uncrowded, preventing combined or overlapping sources from coherently adding in the stack.
2. The source list we stack on comprises unresolved sources so that we are reconstructing an image of the PSF.
3. The source images do not suffer from detection artifacts like saturation or nonlinearity.
4. $\langle \Delta \rangle = 0$, meaning that stacking over many sources averages down the noise.

In practice, we can meet these requirements by making suitable choices for the list of sources to stack on and imposing some weak assumptions about the size of the PSF relative to the pixels. Requirements 1 and 2 are met by using a catalog of stars on which to stack. Stars are unresolved by all but very specialized telescopes and (at middle and high galactic latitudes where source density is low) have uncorrelated positions (Zemcov et al. 2014). Requirements 3 and 4 can be met by




Figure 3. Regridding of a single source to alter resolution. Panel A shows a single Gaussian point source generated in the fine grid. In panel B, the source resolution is sampled down by a factor of 10 to represent a simulated undersampled image (see Section 2.4 for details), and in panel C, the resolution is sampled back up by the same factor. Although the source contains the same number of coarse pixels in panels B and C, the source in panel C contains 100 subpixels in each of these. In panel C, the center of the source appears to have shifted from the coarse grid lines due to the subpixel center now being located with more accuracy than is possible in the coarse grid. Stacking on each source’s subpixel center makes superresolution PSF reconstruction possible by allowing the stacked subpixels to generate high-resolution structure.

selecting sources with fluxes faint enough to keep in the linear regime of the detector and restricting the catalog to a relatively narrow range of fluxes so the denominator F_i in Equation (6) does not strongly overweight noisy sources. However, the noise in the final PSF measurement is a function of the number of sources in the list α , so the tradeoff between the flux range and the number of sources depends on the details of both the instrument and the corresponding survey.

In order to approach this problem in a computationally efficient manner, we stack a tractable number of small “thumbnails” around each source centered on (x_i, y_i) , which focuses attention on the regions of interest and removes the necessity of generating many shifted versions of M . The size of the thumbnail image is determined by including a region large enough that a large fraction of the response from the PSF is included, without including unnecessary background noise.

With these requirements in mind, the stacking PSF estimation algorithm can be broken into five steps:

1. Resample the sky image M into an image \tilde{M} on a pixel gridding r times finer. The image will not appear different because a single pixel value from M will fall into multiple pixels in \tilde{M} , but \tilde{M} will have r times more pixels on a side than M (see Figure 3).
2. For each star in the list α , cut out a thumbnail T centered on the known position of a source (x, y) . The size of the thumbnail depends on the desired angular extent of the final PSF estimate. Due to the nature of resampling, the source’s subpixel center will not necessarily be centered in a coarse-grid pixel. The thumbnail is cut such that the subpixel center will be aligned with the center of the stack.
3. Suppress the overall constant offset value of the thumbnail. For equally bright sources, subtracting the mean of the thumbnail image is acceptable, but for sources of varying brightness, an estimate of the sky brightness away from the star is a good choice.
4. Add the thumbnail centered on the known source position into the stack.
5. Repeat for all stars in the catalog. All of the source thumbnails are combined by taking the mean of all corresponding pixels (i.e., the i th pixel in the stack contains the mean of all the i th pixels in each thumbnail), and this becomes the stacked PSF, P_S . This is a

Table 2
Source Crowding as a Function of Galactic Latitude

(ℓ, b)	Total Available Sources with $11 \leq m_{AB} \leq 15$	Isolated Sources	Sources from Masking	Total Sources in Stack
$(0^\circ, 90^\circ)$	1653	540	808	1348
$(0^\circ, 60^\circ)$	2078	438	1169	1607
$(0^\circ, 30^\circ)$	4710	48	2292	2340
$(0^\circ, 15^\circ)$	6794	5	2185	2190

superresolution image of the underlying optical PSF, P_{true} , convolved with the pixel grid function, P_{grid} . Deconvolution is still necessary to return P_{true} (Guillard et al. 2010).

Source crowding is a concern when stacking sources. Sources appearing in a thumbnail that have brightness similar to the stacking target will contribute flux to the stack and broaden the estimate of P . In the limit that source positions are uncorrelated, interloper sources have random positions and so act as an extra source of noise in P_S . However, with a finite number of sources, there may be significant sample variance from stack to stack. To mitigate this problem, we mask interloper sources as part of the stacking procedure by measuring the distance between the target and any interlopers. Interlopers with center position ≤ 11 coarse-grid pixels from the target source’s center are masked by excluding the thumbnail pixels within a radius of 5 pixels from the interloper position from the sum. If the two sources are closer than 8 coarse-grid pixels apart and the target source is not at least an order of magnitude brighter than the interloper, that source is rejected from the stack entirely. Table 2 demonstrates the effects of our crowding cut on representative 3.5×3.5 deg² fields at various galactic latitudes with fixed longitude. As an example, at $(\ell, b) = (0^\circ, 90^\circ)$, 18% of sources are rejected, while in a crowded field at $(\ell, b) = (0^\circ, 15^\circ)$, 68% of sources are rejected. A benefit of the algorithm discussed here is that, even in such crowded fields, we are able to reconstruct a useful estimate for P_{true} (see Section 3.4).

2.2. Reconstructed PSF Deconvolution

Due to the oversampling inherent to our superresolution stacking method, we require a deconvolution step in order to




Figure 4. Deconvolution of the pixel grid function. In the top left, P_{true} for a Gaussian source is shown, followed by P_S (top right) for a stack of 1000 Gaussian sources at a scaling factor of $r = 10$, P_{grid} (bottom left) that is computed using Equation (7) and deconvolved from P_S via RLD, and the resulting P_{RLD} (bottom right). P_{RLD} exactly matches P_{true} , further demonstrated in Figure 7.

remove P_{grid} and return P_{true} from P_S (Lauer 1999; Starck et al. 2002; Park et al. 2003). This type of issue has also arisen in Planck data, for which the optical beam cannot be recovered without the deconvolution of the effect of sampling time (Planck Collaboration et al. 2016; Tauber et al. 2019). Guillard et al. (2010) follow a similar procedure of superresolution PSF reconstruction followed by deconvolution for the Mid-Infrared Instrument on the James Webb Space Telescope, although their superresolution method combines multiple images followed by deconvolution via a maximum a posteriori method. We choose the Richardson–Lucy deconvolution (RLD), a common algorithm used in this type of problem (Richardson 1972; Lucy 1974). This algorithm is also referred to as the expectation-maximization method and is a form of maximum likelihood estimation (Starck et al. 2002). We have implemented RLD on P_S according to the following prescription. The known blurring factor P_{grid} (shown in Figure 4) is a matrix where each pixel contains the fraction of light detected from a point source:

$$P_{\text{grid}}(x, y) = [r - (x - x_0)][r - (y - y_0)], \quad (7)$$

where r is the upscaling factor, and x_0 and y_0 are the center of the stacking area. Starting from P_S and P_{grid} , each iteration of

the RLD proceeds as follows:

$$P_{\text{RLD}}^{i+1} = P_{\text{RLD}}^i \left(\frac{P_S}{P_{\text{RLD}}^i * P_{\text{grid}}} * P_{\text{grid}}^{\text{ref}} \right), \quad (8)$$

where the initial P_{RLD}^i is simply P_S , and i increases on every iteration for I_{RLD} total iterations. $P_{\text{grid}}^{\text{ref}}$ is the reflection of P_{grid} across both axes. The final result is P_{RLD} , the deconvolution of P_S , and a more accurate reconstruction of P_{true} . This process is demonstrated for ideal Gaussian sources in Figure 4. However, we find that there is a limit to the resulting improvement in reconstruction quality available from increasing the number of RLD iterations due to RLD’s tendency to amplify artificial structure after too many iterations (Hanisch et al. 1997). This is problematic for applications that require fine features of the PSF, for example, optimally weighted photometry (see Section 2.3).

In order to improve the quality of the reconstructed PSF beyond that which can be achieved by RLD, we implement a more advanced method of deconvolution with the iterative back-projection (IBP) algorithm, in which the error or difference between simulated and observed low-resolution images is iteratively reduced (Irani & Peleg 1991). This method is also based on maximum likelihood estimation but was

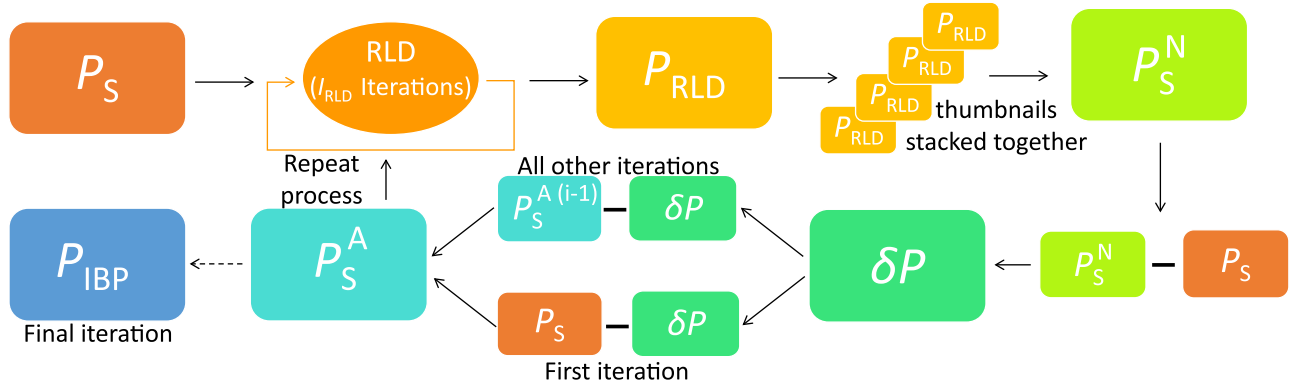


Figure 5. Steps of the iterative back-projection algorithm. On the first iteration P_S is deconvolved using the regular RLD method, while on all subsequent iterations P_S^A is deconvolved. Deconvolution is followed by a repeat of the stacking procedure using P_{RLD} as P_{true} , creating P_S^N . The error is calculated to be the difference between P_S and P_S^N , and that error is subtracted from either P_S (if first iteration) or the previous P_S^A , creating the new P_S^A . The cycle then repeats with deconvolution via RLD again, resulting in P_{IBP} on the final iteration.

developed for superresolution image reconstruction. The approach is similar to back-projection used in tomography and has previously been used in reconstruction of blurred or degraded images. It is similar to RLD in that it offers no unique solution and has the potential to falsely amplify noise after too many iterations (Park et al. 2003). This method has been applied in remote sensing and planetary science (Bell et al. 2006), and other versions of this method have been used to analyze images of solar flares (Schwartz et al. 2014) and in the testing of superresolution image reconstruction methods on simulated Euclid data (Castellano et al. 2015).

After P_S is obtained, the IBP iterations proceed as follows:

1. An RLD is performed as previously described, with I_{RLD} iterations. This produces P_{RLD} .
2. The entire stacking procedure is repeated using P_{RLD} as P_{true} . P_{RLD} is placed into the fine grid (using the same number of sources as were used previously), and the fine grid is downsampled and upsampled by the previously used scale factors. These sources are stacked as before to produce the new stacked PSF, P_S^N .
3. An error term δP is calculated as

$$\delta P = P_S^N - P_S. \quad (9)$$

4. δP is used to produce an error-adjusted stack, P_S^A . On the first iteration, P_S^A is created via

$$P_S^A = P_S - \delta P. \quad (10)$$

On every subsequent iteration, P_S^A is

$$P_S^{A(i+1)} = P_S^{A(i)} - \delta P. \quad (11)$$

5. The cycle repeats with the RLD of the newly formed $P_S^{A(i+1)}$ (which is resampled and stacked to form P_S^N) until a specified completion criterion is reached. The result of the final iteration is P_{IBP} , which is the updated and more accurate reconstruction.

The overall flow of this algorithm is illustrated in Figure 5.

Standard RLD fails in the presence of noise for several reasons, including the amplification of noise over many iterations and, depending on the choice of image zero point, the presence of negative values. This arises because the RLD algorithm is based on the maximum likelihood for Poisson statistics, so its solution requires positive input data. However, stacking requires a choice about the image zero point assumed

in each thumbnail that may produce negative values into P_S . Though it is possible to design algorithms that do not have this feature, the most general case is that we have to account for mildly negative values in P_S .

We handle this problem by implementing a more advanced version of the RLD algorithm (Hanisch et al. 1997) that adds an initial estimate of the background and read-noise value to every pixel and suppresses the contribution from pixels with a value less than the damping factor in each iteration. This damping prevents the amplification of noisy pixels that do not contain structure related to the PSF. Because our PSF is concentrated in the center of the stacked image, we use a damping matrix that performs no damping in the center but damps heavily beyond a circular radius of 13 fine-grid pixels from the center, preventing each iteration from drastically changing the values of damped pixels. Beyond a 13 pixel radius, the PSF is largely noise dominated and does not contribute much signal to the optimal photometry. As presented below, this version of the IBP algorithm successfully deconvolves P_{grid} from P_S even in the presence of noise.

2.3. Using a Reconstructed PSF for Optimal Photometry

Optimal photometry (Naylor 1998) is an example of an application in which detailed knowledge of the PSF is required to reach the maximum possible S/N on point-source fluxes. The best estimate for P_{true} (here P_{IBP} unless explicitly noted) is shifted via interpolation to account for the difference between the source's known coordinates via catalog reference and those same coordinates scaled by the regridding factor, r . It is then downsampled by r to match the resolution of the source. This shifted and resampled reconstructed PSF is normalized and used to give each pixel the weight of its individual flux contribution. This weight is defined as

$$W_P = \frac{P_P}{\sum(P_P^2)}, \quad (12)$$

where P_P is the shifted and resampled reconstructed PSF. Each source's flux is then computed as

$$F_P = \sum(T \cdot W_P), \quad (13)$$

where T is a thumbnail cutout of each source, and W_P is the previously calculated weight function. For an image with many point sources, we define the average deviation of all source fluxes determined via optimal photometry from known

(catalog) values to be

$$\langle \delta F \rangle = \frac{\sum_i^N [(F_{P,i} - F_i) / F_i]}{N}, \quad (14)$$

where $\langle \delta F \rangle$ is expressed in percent, N is the number of sources in the image, F is a source's known flux, and F_P is that computed by optimal photometry.

2.4. Image Simulation

In order to test and characterize these algorithms, simulated images of the sky are necessary. We simulate images of undersampled point sources consistent with the SPHEREx expectation (Korngut et al. 2018) according to the following prescription. First, we generate a grid that represents the r -times finer grid or pixelized image. The native detector size for SPHEREx will be 2048×2048 pixels covering a 3.5×3.5 deg² field of view (FoV) in each of the six main wavelength bands. We select r to be 10 for Gaussian P_{true} , so that our r -times finer image contains $20,480 \times 20,480$ pixels, or r to be 20 for SPHEREx P_{true} , so that the fine-grid image contains $40,960 \times 40,960$ pixels. This increases the sampling rate for SPHEREx to ~ 6 pixels per FWHM, which allows for superresolution recovery of the PSF. The parameter r that we use to simulate the high-resolution image does not need to have the same value as that used to increase resolution during stacking. We choose to set them equal for simplicity, but for real data, the stacking r can be empirically determined. We populate the generated image with point sources using one of two methods:

1. One thousand point sources are randomly placed over the FoV, and for simplicity, all sources are given uniform flux F_p as a dimensionless quantity such that the S/N is chosen to be 20. An S/N of 20 corresponds to $m_{\text{AB}} = 17.87$. See the Appendix for more details on the units we use to define flux and how we calculate S/N.
2. Realistic star fields are generated for any desired sky position with an all-sky catalog of 332 million sources, derived by selecting stars in the Gaia DR2 (Gaia Collaboration et al. 2016) catalog with close counterparts (within a $1''$ angular separation) in the AllWISE catalog (Wright et al. 2010). The use of the g and r_p photometry from Gaia in combination with the W1, W2, and W3 photometry from WISE gives a rough SED for each source from which realistic fluxes for the SPHEREx bands can be estimated. For the special case of single-field tests presented in Section 3, we use the field centered at the north galactic pole (NGP; $(\ell, b) = (0^\circ, 90^\circ)$), where star coordinates are uncorrelated and fields are the least crowded. This “minimal” field contains $\sim 20,000$ sources ranging from $3 < m_{\text{AB}} < 21$.

After the image of point sources gridded according to P_{grid} is created, the map is convolved with the P_{true} (either Gaussian or SPHEREx⁷) under study. This image is then sampled down by $r = 10$ or $r = 20$ as appropriate to the native image resolution.

⁷ The SPHEREx P_{true} is derived from optical simulations performed by L3-SGG as part of the SPHEREx Phase A study using an end-to-end telescope design.

Table 3
Summary of Simulations

PSF Type	Flux Type	RLD/IBP	Noise	Figures	Sections
Gaussian	Uniform	RLD	No	7	3.1
Gaussian	Catalog	RLD	No	8	3.2
Gaussian	Catalog	IBP	No	9	3.2
Gaussian	Catalog	IBP	Yes	10	3.3
SPHEREx	Catalog	RLD	No	11, 12	3.4
SPHEREx	Catalog	IBP	No	11, 12, 13	3.4
SPHEREx	Catalog	IBP	Yes	14, 15, 16, 17, 18	3.4

Figure 3 demonstrates how the regridding process works for a single source.

3. Results

We now have all the pieces required to test the described method against various cases, including nonideal PSF shapes, noise, and crowding, and to assess its overall performance. In Sections 3.1–3.3, we apply the IBP method to Gaussian PSFs in various scenarios to develop intuition. In Section 3.4, we introduce the SPHEREx P_{true} and assess the effects of noise, field position, and other quantities of interest. Table 3 gives a summary of the various tests and the figures and sections where the corresponding results can be found. Finally, in Section 3.5, we apply the IBP PSF reconstruction to data from the LORRI instrument on New Horizons and are able to identify additional complicating factors present in real data such as pointing instability.

3.1. Gaussian Point Sources with Uniform Flux

We first explore the fundamental properties of the PSF estimation method using a simple Gaussian model without noise. This allows us to perform an idealized test of the various stacking and deconvolution methods without complicating factors that may introduce their own sources of error. We construct a simulated image as specified in Section 2.4 with Gaussian point sources and uniform flux. We begin with the stacking procedure described in Section 2.1, which is demonstrated for Gaussian point sources in Figure 6. To evaluate the quality of P_S as a reconstruction of P_{true} , we compare sections through P_S and P_{true} in Figure 7. As expected, we find that P_S is significantly wider than P_{true} as P_S is P_{true} convolved with P_{grid} .

We next evaluate P_S 's performance when used as the kernel for the weighted optimal photometry described in Section 2.3 and find that the measured fluxes using P_S have a wide spread and are $>20\%$ larger than their expected values from P_{true} . To improve the accuracy of the photometry, we implement $I_{\text{RLD}} = 10$ iterations of the RLD algorithm. Figure 7 shows the horizontal and diagonal profiles of P_{RLD} , as well as the difference between P_{RLD} and P_{true} , which is negligible compared to the amplitudes of P_S and P_{true} .

In order to quantify the accuracy of P_{RLD} as a pixel-weighting kernel for photometry, we simulate 50 realizations of noiseless, constant-flux sources with randomized source coordinates. We find $\langle \delta F \rangle = 0.379\% \pm 0.007\%$, which is evidence for a significant output flux bias in the most idealized constant-flux case.

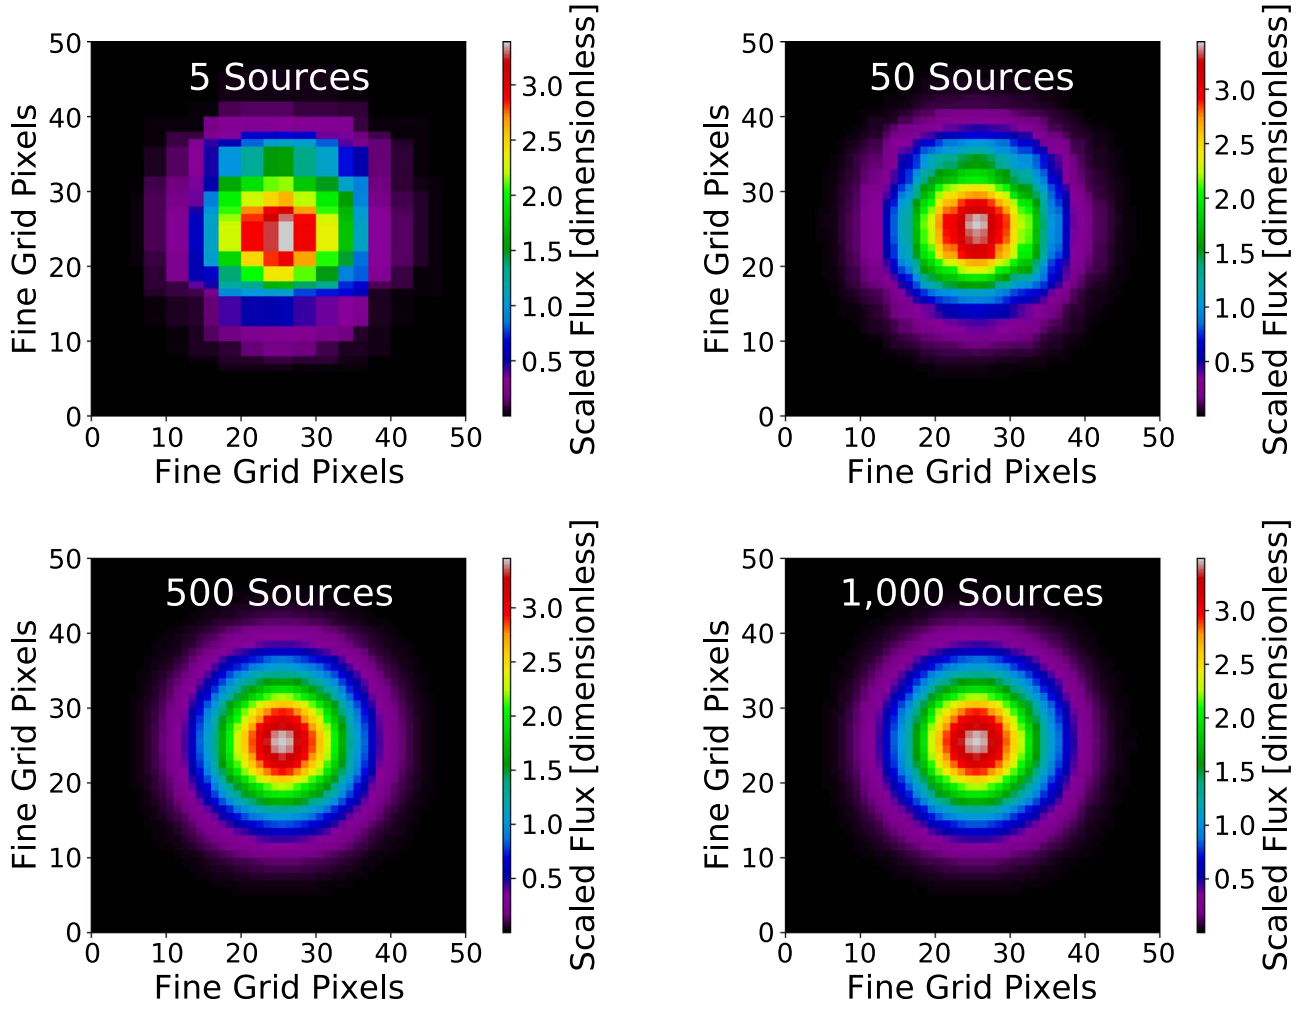


Figure 6. Demonstration of stacking an increasing number of sources. On the top left, five source thumbnails are combined together via their mean. Each new thumbnail is centered on the corresponding source’s known coordinates. The number of sources increases on the top right to 50 sources, on the bottom left to 500 sources, and on the bottom right to 1000 sources. Although each individual thumbnail bears little resemblance to P_{true} , the more thumbnails that are stacked together, the more closely the stack begins to match P_{true} .




Figure 7. Comparison of P_{RLD} to P_{S} and P_{true} for a stack of Gaussian point sources. The left panel shows a horizontal cut comparison of the various PSFs and the difference between the input PSF, P_{true} , and the output PSF, P_{RLD} . The horizontal cut (left panel) is a horizontal section with respect to the pixel grid through the center of the PSF, while the diagonal cut (right panel) is a section from the top-left to the bottom-right corner. While P_{S} is much broader than P_{true} , the deconvolved P_{RLD} very closely matches the width and shape of the optical PSF.




Figure 8. The left panel shows the expected vs. measured flux for noiseless simulations of Gaussian PSFs in the field located at the NGP. Gray lines mark our photometric benchmark within which all fluxes have less than a 1% difference from their known values. Photometry is only performed on sources with $m_{AB} > 16$ as this marks the (approximate) upper limit of unresolved galaxy brightness for SPHEREx. Though the variation is within the requirement, a clear positive bias of the output flux can be seen. The right panel shows the cumulative distribution function of $\langle \delta F \rangle$ for 50 trials of fields with $b > 70^\circ$. The mean deviation, $0.410\% \pm 0.003\%$, is marked with the dotted line, with the gray shaded region indicating the standard error of the mean. The same positive bias as seen on the left is confirmed in an ensemble of fields and appears to be due to subtle biases in the RLD reconstruction of the PSF, indicating the need for a more accurate reconstruction of P_{true} .




Figure 9. The same test as shown in Figure 8 is repeated with the addition of $10 I_{IBP}$. On the left, the comparison of the difference between measured and expected flux to expected flux for a single field at the NGP reveals a much tighter correlation with bias at the 10^{-4} level. On the right, the CDF of $\langle \delta F \rangle$ from 50 trials with varying galactic longitude at $b > 70^\circ$ has a mean of $0.0123\% \pm 0.0002\%$, marked with the dotted line and the gray region. In this noiseless, idealized simulation IBP reduces the bias from RLD alone by an order of magnitude.

3.2. Gaussian Point Sources with Catalog-based Flux

To check the effect of a more realistic distribution of input fluxes on the RLD reconstruction, we simulate an image with Gaussian point sources and catalog fluxes for a field at the NGP constructed as described in Section 2.4, again without noise in the image. Performing the same stacking and deconvolution procedure yields the results shown in Figure 8. All fluxes fall within 1% of their known values, but a positive bias at the level of $\sim 0.4\%$ remains. To study the variance in an ensemble of simulated images, instead of randomizing individual source coordinates, we calculate results for 50 randomly centered and independent fields with $b > 70^\circ$. The $\langle \delta F \rangle$ from this ensemble is $0.410\% \pm 0.003\%$, which verifies the positive bias seen in the single trial at the NGP and indicates that the use of

catalog-based flux instead of uniform flux has not had a negative impact. This bias is due to P_{RLD} still being slightly broader than P_{true} at the level of $\sim 1:10^4$. We conclude that in order to remove the bias, a more accurate reconstruction of P_{true} is necessary.

In order to obtain a more accurate reconstruction, we use the IBP algorithm outlined in Section 2.2. Similar to I_{RLD} , I_{IBP} is likewise determined by minimizing $\langle \delta F \rangle$'s total spread and bias until no further improvements can be achieved, which we empirically find converges by $I_{IBP} = 10$. Results from the more advanced algorithm for the 50 independent fields with $b > 70^\circ$ are shown in Figure 9 and show significant improvements in both the mean and variance of δF . This is good evidence that P_{IBP} is a more accurate reconstruction of P_{true} than P_{RLD} . Performing the same test on sources with uniform flux also




Figure 10. For a single trial using a Gaussian P_{true} and catalog-based sources, varying photometric accuracy results depending on whether we add noise to the sources, the kernel, or both. Panel A shows the photometry using a noisy P_{IBP} as the kernel for noisy sources. Results are similar in panel B, which shows flux values obtained using the noiseless P_{true} applied to noisy sources. For both panels A and B, instrumental and photon noise dominate the photometric results. Using the noisy P_{IBP} as the kernel for photometry of noiseless sources (panel C) gives all flux values within the desired 1% requirement, demonstrating that our reconstruction methods are not introducing photometric error beyond the SPHEREx requirement.

results in the previously seen bias in $\langle \delta F \rangle$ being reduced by an order of magnitude, which further demonstrates that the bias was due to reconstruction quality and not the type of source fluxes used.

3.3. Noise

We simulate noise from the instrument consistent with SPHEREx Band 1 (centered at $0.93 \mu\text{m}$; see Doré et al. 2018 for full specifications) with pixel rms of $46 \text{ nW m}^{-2} \text{ sr}^{-1}$ by adding a white noise component to the native resolution source image. We also model photon noise from the sources themselves as an additional source of noise. At this noise amplitude, a source detected at 5σ has $m_{\text{AB}} = 19.4$. We restrict the sources included in P_S to the range $10 \leq m_{\text{AB}} \leq 21.1$. The bright limit is determined by sources that would require $>50\%$ correction due to nonlinearity in the SPHEREx detector, and the lower limit by sources with $S/N = 1$. The latter choice is not motivated by any known limitation of the algorithm, but rather by including only those sources not dominated by noise. In the NGP field, using P_{IBP} constructed from noisy sources as the kernel for optimal photometry returns photometric fluxes well within the desired 1% benchmark, as illustrated in Figure 10. Performing a more realistic simulation including noise in both the photometered source fluxes and the PSF stack sources results in photometry that exceeds the 1% requirement, but that is dominated by the noisy fluxes rather than the intrinsic error from the photometry kernel.

3.4. Spatially Structured PSF

In order to provide a concrete test case and assess our IBP reconstruction of the SPHEREx PSF, we simulate an image matching SPHEREx’s specifications using the sky catalog for the NGP as described in Section 2.4. We convolve the point sources with the SPHEREx P_{true} derived from optical simulations of the instrument, shown in Figure 11, and optionally add instrument and photon noise. Next, we perform the stacking procedure, pixel grid deconvolution, and optimal photometry as described in Section 2.

First, we perform a noiseless test. A comparison of P_{true} to the resulting P_S , P_{RLD} , and P_{IBP} calculated in the absence of noise is displayed in Figure 11. As expected from the Gaussian simulations, P_S is not similar to P_{true} , but P_{RLD} with 200 I_{RLD} offers a significant improvement in reconstruction quality. Using the full IBP algorithm allows an even more accurate

reconstruction of P_{true} . A comparison of the squared Fourier representations of P_{true} and P_{IBP} is shown in Figure 12, along with the corresponding differences between P_{true} , P_{RLD} , and P_{IBP} . Our choice of r dictates the frequency of information we are able to recover in the reconstruction. We find that using $r = 10$ for the SPHEREx P_{true} does not return enough high-frequency information in the reconstruction, so we use $r = 20$ to obtain accurate results. This is an empirical choice based on reconstruction quality as measured by the output photometric accuracy. Applying P_{IBP} as the photometry kernel results in no detectable bias within uncertainties, as demonstrated in Figure 13. With this procedure, all measured fluxes are within 1% of their known values, and 50 trials of fields with randomized galactic longitude at $b > 70^\circ$ produce $\langle \delta F \rangle$ consistent with zero.

Next, we test the reconstruction of the SPHEREx PSF in the presence of noise in the stacked sources. We find that setting pixel values beyond an exclusion radius of 31 fine-grid pixels to zero after both the initial RLD and the IBP sequence yields the best results in this case. This is because noisy pixels with little or no PSF signal can amplify noise artifacts. The damping radius controls the weight given to pixels between the damping and exclusion radii during the RLD, while the exclusion radius provides a hard cutoff for pixels in the P_S image that have no significant effect on the PSF. Radii for damping and exclusion are determined through minimization of the average flux deviation derived from using P_{IBP} as a kernel for optimally weighted photometry, $\langle \delta F \rangle$, and its bias from the known input F . Because the average deviation from the expected flux is desired to be zero, any significant trend from zero indicates a bias being introduced during reconstruction or photometry. The value and bias of $\langle \delta F \rangle$ are minimized until the number of iterations, I_{IBP} , and the exclusion and damping radii both converge. The radii and number of iterations have converged when no further improvements occur. Figure 14 demonstrates the deconvolution process for the SPHEREx PSF in the presence of noise.

3.4.1. Evaluating Reconstruction Quality

In order to evaluate the quality of the reconstruction, as well as its effectiveness when used as a weight kernel for optimal photometry, we develop a figure of merit (FOM) based on $\langle \delta F \rangle$.

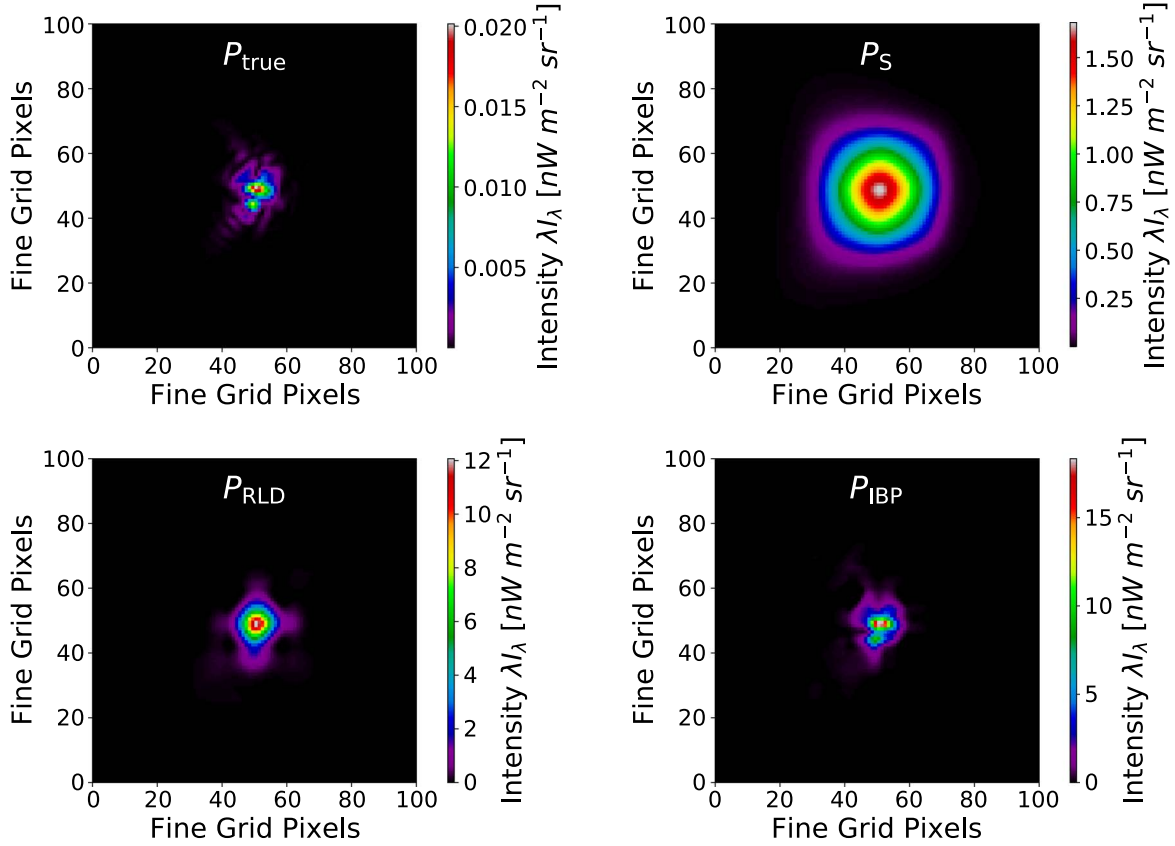


Figure 11. Images of the SPHEREx P_{true} (top left), P_S (top right), P_{RLD} after 200 RLD iterations (bottom left), and P_{IBP} after 300 RLD iterations combined with the adaptive IBP iterations (bottom right) on $0''.31 \text{ pixel}^{-1}$ gridding. Because the IBP offers a more accurate reconstruction of P_{true} , it also provides more accurate photometry when used as pixel weighting.

Our FOM is the rms of δF :

$$\sigma_{\delta F} = \sqrt{\frac{\sum_i^N [(\delta F_i)^2]}{N}}, \quad (15)$$

where δF_i is the difference between measured and expected flux for any individual source, and N is the total number of sources. This gives a measure of the total dispersion in δF such that $\sigma_{\delta F} < 1$ indicates photometric results that meet the SPHEREx requirement.

An important issue is the optimal range of sources to use for constructing P_S . Bright sources allow measurement of P_S in the coarse gridding with high fidelity, but as there are relatively few sources, this comes at the cost of the spatial resolution of the reconstruction. Faint sources are numerous, but the effect of noise is larger in the stack, and this has a cascading effect on the fidelity of P_{IBP} . In order to optimize the range of fluxes to use, we calculate P_{IBP} over narrow magnitude ranges and from these calculate $\sigma_{\delta F}$ to determine where it is optimal. Figure 15 demonstrates how the average $\sigma_{\delta F}$ of 50 trials with randomized galactic longitude for $b > 70^\circ$ changes when selecting only sources within a single AB magnitude bin for use in P_S . We find that $\sigma_{\delta F}$ is minimized for $11 \leq m_{\text{AB}} \leq 15$, which are the statistically optimal source magnitudes for optimal PSF reconstruction at the SPHEREx noise amplitude.

Next, we test how using a much smaller subset of the total available sources affects the reconstruction quality. We restrict the number of sources used in P_S by an increasing percentage via a representative sample for the optimal range of

$11 \leq m_{\text{AB}} \leq 15$. Figure 16 shows how the average $\sigma_{\delta F}$ varies with this restriction for a series of 50 trials of fields with varying galactic longitude and $b > 70^\circ$ at each fraction of the total number of sources. As the fraction of sources being used increases, $\sigma_{\delta F}$ continues to decrease, indicating better reconstruction quality. Even restricting to only 6% of the available sources yields $\sigma_{\delta F} < 1$ on the mean. We conclude that it should be possible to produce accurate kernel reconstructions over regions as small as 10% of the full SPHEREx FoV. Because the SPHEREx PSF is expected to change somewhat over the FoV as the effective bandpass at the detector changes, this capability is invaluable to produce P_S appropriate to a region where the gradient of the PSF properties will be smaller.

3.4.2. Overall Photometric Accuracy

To assess the overall reliability of the IBP algorithm, in Figure 17 we compare all six SPHEREx wavelength bands (centered at 0.930, 1.375, 2.030, 3.050, 4.065, and $4.730 \mu\text{m}$) using the $11 \leq m_{\text{AB}} \leq 15$ range for P_S . P_{true} and the noisy P_{IBP} are shown for each band, along with a cumulative distribution function (CDF) of the resulting $\langle \delta F \rangle$ for the noisy P_{IBP} applied to noiseless sources for 50 trials with randomized galactic longitude at $b > 70^\circ$. We use catalog-based fluxes from the Gaia catalog for all bands, resulting in a high level of accuracy with mean percent error consistent with zero for each band. This demonstrates the algorithm's robustness in successful PSF reconstruction of complex PSFs with different types of structure from images with varying levels of noise.




Figure 12. A comparison between the squared Fourier transform of the SPHEREx P_{true} (top left), P_{IBP} (top right), difference between P_{true} and P_{RLD} (bottom left), and difference between P_{true} and P_{IBP} (bottom right). We would expect a diffraction-limited system to go to zero response near $4\pi D/\lambda \sim 14 \text{ arcsec}^{-1}$; the measured value of 11.7 arcsec^{-1} is consistent with the beam being slightly larger than diffraction limited. The IBP method exhibits residuals at the 10^{-3} level, which is significantly less than the traditional RLD method that shows residuals at the $>10^{-2}$ level.

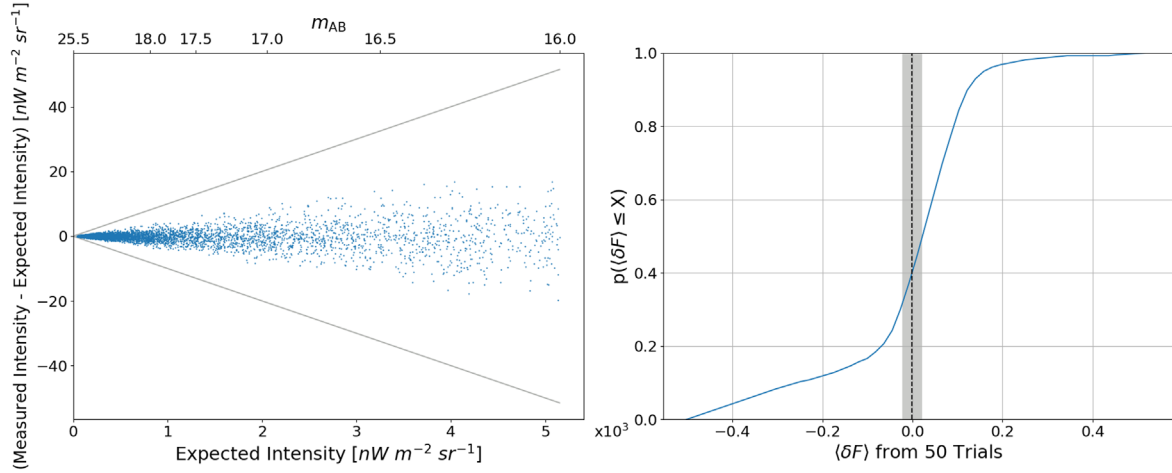


Figure 13. (Left) The difference between measured and expected flux as a function of expected value, with gray lines showing the 1% requirement for a single trial using the noiseless SPHEREx P_{true} and IBP algorithm. All fluxes lie within the 1% boundaries, and we detect no overall flux bias. (Right) $\langle \delta F \rangle$ for 50 trials of fields at $b > 70^\circ$ and randomized galactic longitude, which has a mean of $-0.001\% \pm 0.022\%$. There is no statistical bias in the output fluxes from this method in this test.

In previous tests, we have performed optimal photometry on intrinsically noiseless sources to assess the accuracy of the IBP reconstruction algorithm. It is also informative to investigate at which source fluxes the instrumental and shot noise present in real sources would dominate over the error arising from misestimation of P_{true} . In the top row of Figure 18, we show tests using combinations of noise in the stack and noise in the photometered sources (or lack thereof) for the NGP, compared

to using a noiseless P_{true} as the photometry kernel. We continue to restrict P_S to $11 \leq m_{\text{AB}} \leq 15$. As expected, only when the noisy P_{IBP} is used for photometry on sources with perfectly known flux is the resulting photometry within 1% variation at all F . Photometry on noisy sources is dominated by the noise in the source flux measurement. Noise in the PSF reconstruction itself appears to have a small effect compared with the intrinsic scatter of F from random noise. Additionally, the PSF

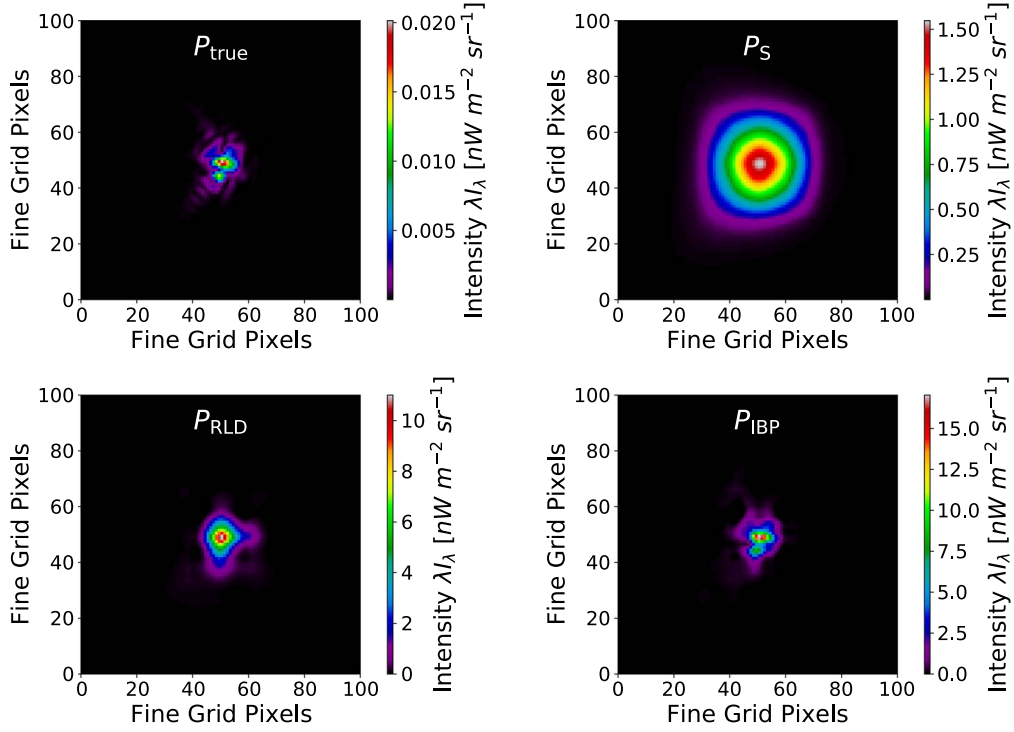


Figure 14. The deconvolution process in the presence of noise for the $0.93 \mu\text{m}$ SPHEREx wavelength band. In the upper-left panel, we show P_{true} for comparison with the noisy P_S on the top right, the initial P_{RLD} with the use of damping and background noise constant on the bottom left, and the final P_{IBP} after the IBP process on the bottom right. Due to the large number of high-S/N sources in the stack, instrument noise has minimal effect on reconstruction quality.




Figure 15. Comparison of $\sigma_{\delta F}$ when restricting the sources used in P_S to 1 mag wide bins (lower axis) centered between integer m_{AB} values. The blue points and line indicate the average $\sigma_{\delta F}$ (left axis) from 50 trials of each bin at varying galactic longitude and $b > 70^\circ$. The black dashed line shows the desired limit of $\sigma_{\delta F} = 1$, and the green dashed line indicates the average $\sigma_{\delta F}$ from 50 trials of a $11 \leq m_{\text{AB}} \leq 15$ bin, with the green region showing the 1σ error. This region provides the lowest $\sigma_{\delta F}$ compared to individual magnitude bins. The gray line shows the number of sources used in P_S in any given bin (right axis), and the upper axis shows the corresponding S/N for each magnitude bin. $\sigma_{\delta F}$ always decreases with an increasing number of sources or higher S/N and is minimized in the zone that balances both.

reconstruction is not correlated with details of the noise realization; the P_{IBP} derived from one noise realization, when used for photometry on sources from different noise realizations, returns an unbiased $\langle \delta F \rangle$.

To understand the relative effect of noise in the stack versus numerical inaccuracies in the algorithm, we perform 50 trials




Figure 16. Comparison of $\sigma_{\delta F}$ correlating to the fraction of the total number of available sources being used in P_S . The blue region corresponds to the 1σ error from 50 trials of randomized galactic longitude at $b > 70^\circ$. Increasing the number of sources used in P_S results in lower $\sigma_{\delta F}$ and thus better reconstruction quality, while using as few as 6% of sources still allows for $\sigma_{\delta F} < 1$, marked by the black dashed line.

with fields of varying galactic longitude at $b > 70^\circ$ for the same set of noise cases. We again find that the variation in the output flux is dominated by the intrinsic noise rather than the error in the IBP, demonstrated in the bottom row of Figure 18. Results are very similar for noisy sources with a noiseless P_{true} and noiseless sources with a noisy P_{IBP} , as both have mean $\langle \delta F \rangle$ consistent with zero and similar distribution widths. Again, the PSF kernel reconstruction is not introducing photometric error beyond that determined by the noise on the sources.

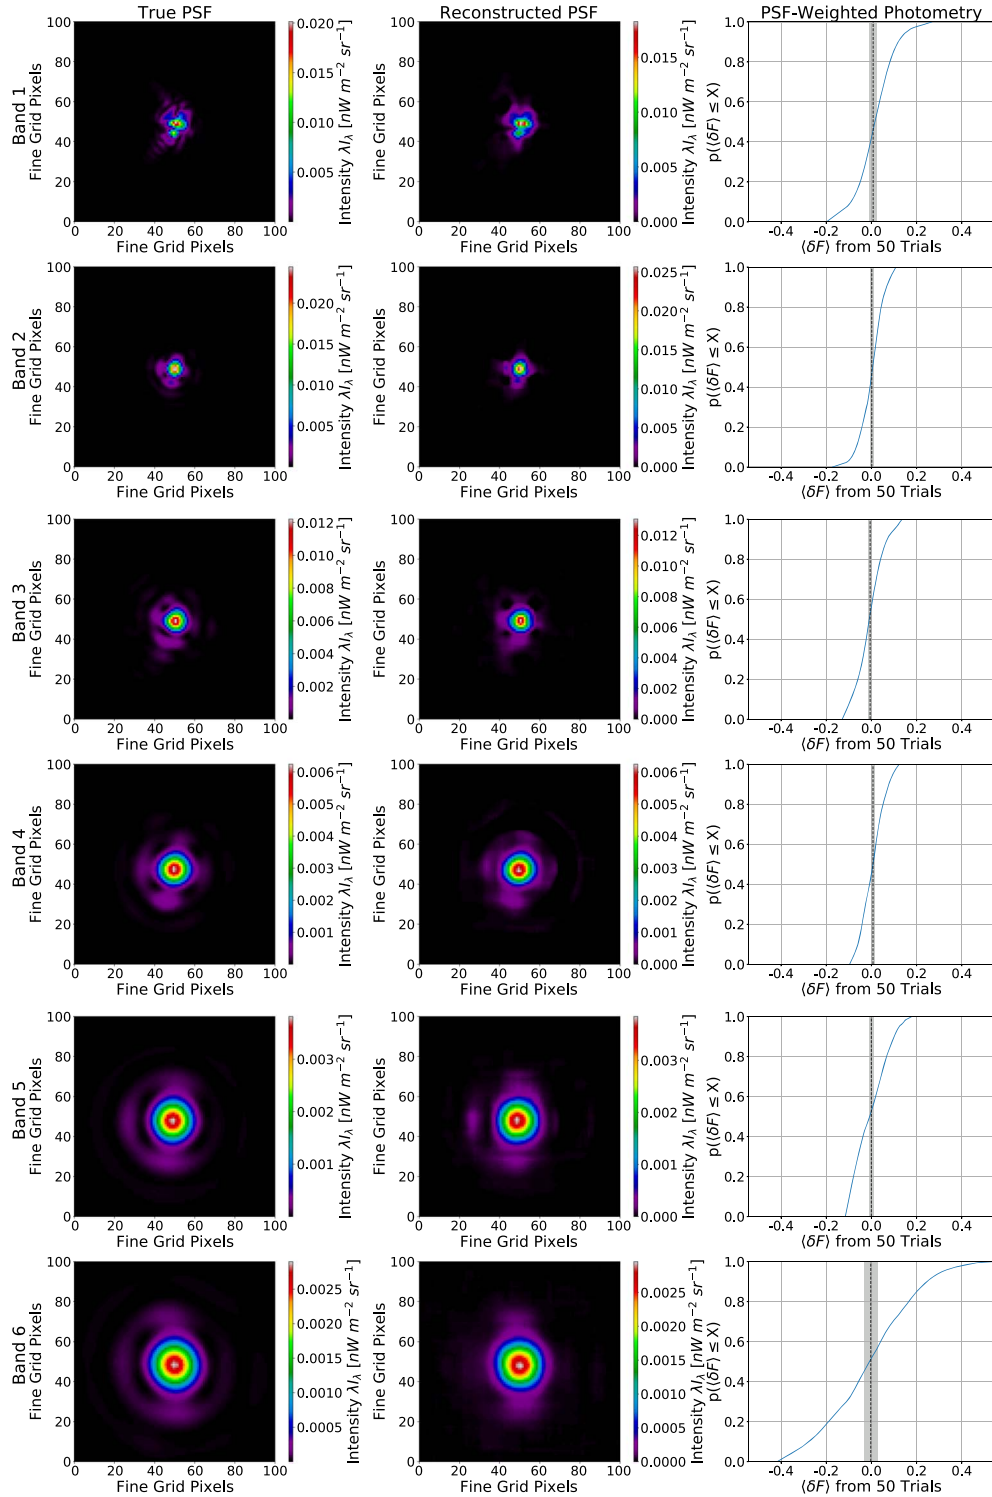


Figure 17. Comparison of the six SPHEREx wavelength bands, with each band’s P_{true} in the first column, final reconstructed P_{IBP} in the second column, and CDF of $\langle \delta F \rangle$ from 50 trials of randomized galactic longitude at $b > 70^\circ$ in the third column. Mean $\langle \delta F \rangle$ for all trials is marked with a dotted line, with a gray region showing the standard error of the mean. The mean $\langle \delta F \rangle$ is consistent with zero for each band.

3.4.3. Source Crowding

To understand the effects of source crowding (discussed in Section 2.1), we vary galactic coordinates in a single trial for each of $b = 15^\circ, 30^\circ, 60^\circ$, and $\ell = 0^\circ, 90^\circ$, and 180° with P_S restricted to $11 \leq m_{\text{AB}} \leq 15$. All trials use the SPHEREx P_{true} for the $0.93 \mu\text{m}$ wavelength band with catalog-based fluxes and

noisy P_{IBP} with noiseless sources during photometry. Table 4 gives a comparison of $\langle \delta F \rangle$ and $\sigma_{\delta F}$ for each set of coordinates. All fields have $\sigma_{\delta F} < 1$ with no bias in $\langle \delta F \rangle$ and no obvious increase in $\sigma_{\delta F}$. This indicates that accurate PSF reconstruction and photometry can be achieved with this method, even for highly crowded fields.

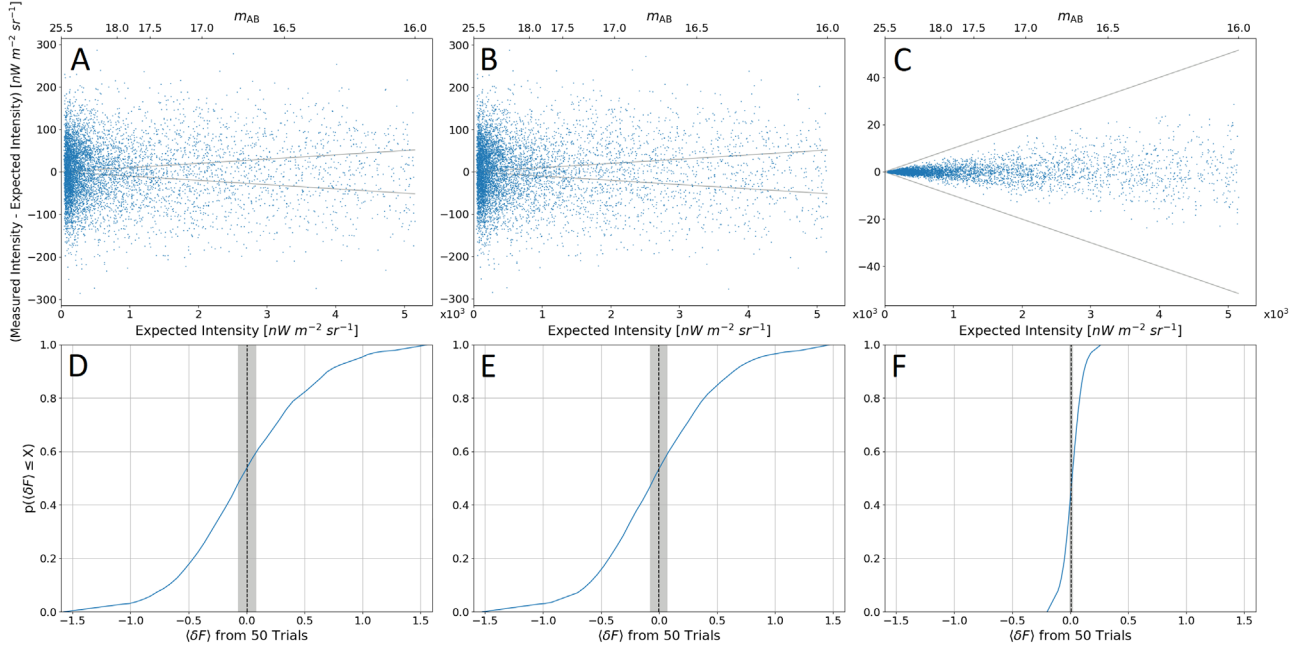


Figure 18. The top row shows single trials at the NGP using the SPHEREx $0.93 \mu\text{m}$ P_{true} with catalog-based flux values and instrument noise. Panels A and B show photometry resulting from noisy sources with a noisy P_{IBP} and noisy sources with a noiseless P_{true} , respectively. In both cases, flux values display similar scatter, indicating that reconstruction quality does not negatively affect photometric results. Panel C shows the results from noiseless sources and a noisy P_{IBP} , all of which are within 1% boundaries. To quantify the bias in these cases, the bottom row shows CDFs of $\langle \delta F \rangle$ corresponding to panels A–C for 50 trials with randomized field galactic longitude and $b > 70^\circ$, where the dotted line indicates the mean $\langle \delta F \rangle$ of all trials, and the gray region shows the standard error of the mean. Panel D shows a mean $\langle \delta F \rangle$ of $0.005\% \pm 0.080\%$, panel E has a mean $\langle \delta F \rangle$ of $-0.003\% \pm 0.075\%$, and panel F has a mean $\langle \delta F \rangle$ of $0.008\% \pm 0.016\%$.




Figure 19. On the left is an example LORRI image with $4''$ 1 pixel^{-1} . On the right is a comparison of the P_S and P_{IBP} measured from this image, shown in linear scale in the top row and logarithmic scale in the bottom row. The PSF measured in this image is significantly larger than the laboratory-determined PSF (Cheng et al. 2008), which is likely due to the pointing stability of the spacecraft (Noble et al. 2009).

Table 4
Photometric Results as a Function of Sky Position

	$\ell = 180^\circ$		$\ell = 90^\circ$		$\ell = 0^\circ$	
	$\langle \delta F \rangle$	$\sigma_{\delta F}$	$\langle \delta F \rangle$	$\sigma_{\delta F}$	$\langle \delta F \rangle$	$\sigma_{\delta F}$
$b = 60^\circ$	0.031%	0.193%	0.010%	0.163%	-0.017%	0.218%
$b = 30^\circ$	-0.046%	0.201%	0.004%	0.168%	-0.026%	0.207%
$b = 15^\circ$	-0.078%	0.205%	-0.066%	0.172%	-0.090%	0.210%

3.5. LORRI PSF Estimation

To this point, we have been working solely with simulated PSFs where we know the input and the noise is idealized. To

test this algorithm against real data, we operate on an image taken by the LORRI instrument on New Horizons (Cheng et al. 2008; Zemcov et al. 2017). This image was acquired by the

CCD chip on the LORRI instrument and, in LORRI's 4×4 binning mode, has size 256×256 pixels at $4''1 \text{ pixel}^{-1}$, as shown in Figure 19. In the PSF reconstruction, we use $r = 20$ and $I_{\text{IBP}} = 10$. The reconstructed PSF is shown in Figure 19. We find that the FWHM of the minor axis of P_{IBP} is $3''.73$, while the FWHM of the major axis is $5''.51$, both of which are significantly larger than the $1''.5$ PSF measured in the lab (Cheng et al. 2008).

What can account for this difference? Noble et al. (2009) demonstrate that the PSF width is a strong function of the integration time of the instrument, and the New Horizons spacecraft is known to exhibit pointing drift at the arcsec s^{-1} level. Performing the deconvolution on a set of several hundred 10 s LORRI exposures, we find the PSF is often extended with eccentricity $\epsilon > 0.8$ and a minimum pointing drift of $1''$. This is consistent with the expected drift in the spacecraft's pointing in relative control mode of $\pm 2''$ per exposure (Conard et al. 2017). In principle, if we understood the pointing history of the spacecraft, we could deconvolve this component out to isolate the underlying optical PSF using the same methods described above; this will be left to future work.

4. Discussion

4.1. Comparison to Similar Methods

The method of PSF reconstruction described here is robust to complicating factors such as severe undersampling, complex PSFs, noise, crowded fields, or a limited number of available point sources. Starck et al. (2002) and Puetter et al. (2005) provide thorough reviews of deconvolution methods commonly used in astronomy, which are numerous. New methods are still being actively developed for applications in astronomy (see, e.g., Sureau et al. 2020). In comparison to some algorithms, a major advantage of this method is that it works on a per-exposure basis. Many other superresolution methods that have been used in astronomy (e.g., Orieux et al. 2012; Li et al. 2018; Guo et al. 2019) rely on multiple exposures of the same field for information reconstruction. As an example, data from the Hubble Space Telescope prompted the development and application of a number of techniques driven by undersampling in the Wide Field and Planetary Camera 2 (WFPC2; Anderson & King 2000) and Wide Field Camera 3 (WFC3; Anderson 2016). These methods rely on dither-based solutions in which multiple images of the same field with some small shift in the position of the detector are recombined to generate a higher-resolution image of the same field (Lauer 1999; Fruchter & Hook 2002). Our method does not rely on similar constraints.

The SPHEREx application investigated here requires super-resolution knowledge of the PSF to optimally weight pixels for photometry and, as a result, requires the effect of the pixelization to be modeled on a per-source basis. Some reconstruction methods return the PSF convolved with the pixel-gridding function (the effective PSF or the PSF observed in the image) instead of the underlying optical PSF, such as the method introduced by Anderson & King (2000). While the effective PSF has important uses, it cannot be used directly for optimal photometry, pointing drift assessment, or other applications where the details of the underlying optical PSF are important. Other methods of superresolution image reconstruction can also deconvolve the pixel-gridding function and be applied to single exposures (Aujol et al. 2006), but have

applications significantly different from those discussed here (e.g., Castellano et al. 2015).

Computational speed is also an advantage of this method. More complicated methods such as blind deconvolution (Fish et al. 1995) or Bayesian methods (Shi et al. 2017) have also been used for analyses in astronomy, but these tend to be much more computationally expensive. Comparatively, our method of stacking and deconvolution requires ~ 1 minute to analyze a single exposure on a personal computer, which is scalable to large surveys.

4.2. Future Improvements

To summarize the primary result of this work, through simulations of the stacking method in all six SPHEREx wavelength bands including realistic noise, source catalogs from Gaia+AllWISE star catalogs, and realistic beam shapes, we find IBP-derived kernels allow photometry with accuracy to $\sim 0.2\%$ in a single SPHEREx exposure. We find that kernels derived from stars with $11 \leq m_{\text{AB}} \leq 15$ generate the best estimate of the underlying optical PSF across all six SPHEREx bands. At SPHEREx's noise level, the population of sources in this range balance the need for large S/N in the stack against the need for a large number of sources to maintain spatial fidelity. We find that stacking on subimages of the full SPHEREx FoV does not significantly degrade the performance of the algorithm up to fields with $\lesssim 10\%$ of the full FoV, at least at high and mid-galactic latitudes.

Though we have determined a method that meets the 1% accuracy requirement, further improvements and details could be investigated. The IBP reconstruction method is primarily tuned to returning a kernel for optimal photometry of point sources, rather than reconstruction over a wide range of spatial scales. "Stitching" together information from the brightest sources to probe the faint wings of the PSF with measurements of the central region of the PSF derived using IBP may offer an improved measurement of the PSF over a wide range of spatial scales (Zemcov et al. 2014). Further, more advanced algorithms that can separate effects like the optical PSF, pointing jitter and drift, and distortion over the array may be possible using larger volumes of data. Though such advancements might offer even more accurate reconstructions, the method presented here already illustrates how advanced statistical methods can offer a path to unlocking the full information content of astronomical images.

Thanks to Chi Nguyen for helpful comments and suggestions. This work was supported by NASA awards 80GSFC18C0011/S442557, NNN12AA01C/1594971, and 80NSSC18K1557. The research was partly carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA (80NM0018D0004).

This publication makes use of data products from the Wide-field Infrared Survey Explorer, which is a joint project of the University of California, Los Angeles, and the Jet Propulsion Laboratory/California Institute of Technology, and NEOWISE, which is a project of the Jet Propulsion Laboratory/California Institute of Technology. WISE and NEOWISE are funded by the National Aeronautics and Space Administration. This work has made use of data from the European Space Agency (ESA) mission Gaia (<https://www.cosmos.esa.int/gaia>), processed by the Gaia Data Processing and Analysis Consortium (DPAC, <https://www.cosmos.esa.int/web/gaia/dpac/consortium>). Funding for the

DPAC has been provided by national institutions, in particular the institutions participating in the Gaia Multilateral Agreement.

The authors acknowledge Research Computing at the Rochester Institute of Technology for providing computational resources and support that have contributed to the research results reported in this publication.

Software: Astropy (Astropy Collaboration et al. 2013, 2018), Matplotlib (Hunter 2007), NumPy (Van Der Walt et al. 2011), SciPy (Virtanen et al. 2020).

Appendix Point-source Flux and Units

For point sources that are assigned some magnitude and corresponding specific flux F_λ (measured in units of power per unit area per unit wavelength such as $\text{nW m}^{-2} \mu\text{m}^{-1}$), we calculate specific intensity (generally measured in units of power per unit area per unit solid angle per unit wavelength) I_λ as

$$I_\lambda = \frac{F_\lambda}{\Omega_{\text{beam}}} [\text{nW m}^{-2} \text{sr}^{-1} \mu\text{m}^{-1}], \quad (\text{A1})$$

where Ω_{beam} is the beam area on the sky for SPHEREx. We then multiply I_λ by a specific wavelength to get the intensity or diffuse surface brightness λI_λ , where

$$\lambda I_\lambda = \frac{\lambda F_\lambda}{\Omega_{\text{beam}}} [\text{nW m}^{-2} \text{sr}^{-1}]. \quad (\text{A2})$$

The quantity λI_λ is equivalent to the quantity νI_ν , where I_ν is measured in units of power per unit area per unit solid angle per unit frequency such as $\text{nW m}^{-2} \text{sr}^{-1} \text{Hz}^{-1}$. With the proper conversion, I_λ and I_ν can be used interchangeably. When we perform photometry of a point source, we are taking the integral

$$\int_{\text{beam}} I_\nu d\Omega = F_\nu [\text{Jy}] \quad (\text{A3})$$

to get some flux F_ν . For our point sources with catalog-based flux values, we choose to convert the photometric flux into intensity in units natural to SPHEREx, $\text{nW m}^{-2} \text{sr}^{-1}$, using Equation (A2). Given some source with intensity λI_λ and a pixel rms σ , the S/N can be defined as

$$\text{S/N} = \frac{\lambda I_\lambda}{\sigma}, \quad (\text{A4})$$

where λI_λ and σ have been integrated over a beam with area on the sky Ω_{beam} . For our simulated point sources with uniform flux, we assign λI_λ such that the S/N is a desired value based on the selected pixel rms, and quote the S/N as scaled flux of dimensionless units, easily converted into flux units via Equations (A4) and (A2).

ORCID iDs

Teresa Symons  <https://orcid.org/0000-0002-9554-1082>
Michael Zemcov  <https://orcid.org/0000-0001-8253-1451>
Yun-Ting Cheng  <https://orcid.org/0000-0002-5437-0504>
Brendan Crill  <https://orcid.org/0000-0002-4650-8518>

References

- Anderson, J. 2016, Empirical Models for the WFC3/IR PSF, Space Telescope WFC Instrument Science Report, [WFC3 2016-12](#)
- Anderson, J., & King, I. R. 2000, [PASP](#), **112**, 1360
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, [AJ](#), **156**, 123
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, [A&A](#), **558**, A33
- Aujol, J.-F., Gilboa, G., Chan, T., & Osher, S. 2006, [Int. J. Comput. Vis.](#), **67**, 111
- Bell, J. F., Joseph, J., Sohl-Dickstein, J. N., et al. 2006, [JGRE](#), **111**, E02S03
- Bertin, E. 2011, in ASP Conf. Ser. 442, Astronomical Data Analysis Software and Systems XX, ed. I. N. Evans et al. (San Francisco, CA: ASP), 435
- Béthermin, M., Le Floc'h, E., Ilbert, O., et al. 2012, [A&A](#), **542**, A58
- Cady, F. M., & Bates, R. H. T. 1980, [OptL](#), **5**, 438
- Castellano, M., Ottaviani, D., Fontana, A., et al. 2015, in ASP Conf. Ser. 495, Astronomical Data Analysis Software and Systems XXIV, ed. A. R. Taylor & E. Rosolowsky (San Francisco, CA: ASP), 257
- Cheng, A. F., Weaver, H. A., Conard, S. J., et al. 2008, [SSRv](#), **140**, 189
- Conard, S. J., Weaver, H. A., Núñez, J. I., et al. 2017, [Proc. SPIE](#), **10401**, 104010W
- Dole, H., Lagache, G., Puget, J.-L., et al. 2006, [A&A](#), **451**, 417
- Donlon, K., Ninkov, Z., & Baum, S. 2018, [PASP](#), **130**, 074503
- Doré, O., Werner, M. W., Ashby, M. L. N., et al. 2018, [arXiv:1805.05489](#)
- Fish, D. A., Brinicombe, A. M., Pike, E. R., & Walker, J. G. 1995, [JOSAA](#), **12**, 58
- Fruchter, A. S., & Hook, R. N. 2002, [PASP](#), **114**, 144
- Gai, M., & Cancelliere, R. 2007, [MNRAS](#), **377**, 1337
- Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, [A&A](#), **595**, A1
- Guillard, P., Rodet, T., Ronayette, S., et al. 2010, [Proc. SPIE](#), **7731**, 166
- Guo, R., Shi, X., & Wang, Z. 2019, [JEI](#), **28**, 023032
- Hanisch, R. J., White, R. L., & Gilliland, R. L. 1997, in Deconvolution of Images and Spectra, ed. P. A. Jansson (2nd ed.; San Diego, CA: Academic), 310
- Hirata, C. M., & Choi, A. 2020, [PASP](#), **132**, 014501
- Home, K. 1986, [PASP](#), **98**, 609
- Hunter, J. D. 2007, [CSE](#), **9**, 90
- Irani, M., & Peleg, S. 1991, [CVGIP: Graph. Models Image Process.](#), **53**, 231
- Komgut, P. M., Bock, J. J., Akeson, R., et al. 2018, [Proc. SPIE](#), **10698**, 106981U
- Lauer, T. R. 1999, [PASP](#), **111**, 1434
- Li, Z., Peng, Q., Bhanu, B., Zhang, Q., & He, H. 2018, [Ap&SS](#), **363**, 92
- Lucy, L. B. 1974, [AJ](#), **79**, 745
- Marsden, G., Ade, P. A. R., Bock, J. J., et al. 2009, [ApJ](#), **707**, 1729
- Naylor, T. 1998, [MNRAS](#), **296**, 339
- Noble, M. W., Conard, S. J., Weaver, H. A., Hayes, J. R., & Cheng, A. F. 2009, [Proc. SPIE](#), **7441**, 74410Y
- Orieux, F., Giovannelli, J. F., Rodet, T., et al. 2012, [A&A](#), **539**, A38
- Park, S., Park, M., & Kang, M. 2003, [ISPM](#), **20**, 21
- Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016, [A&A](#), **594**, A7
- Plazas, A. A., Shapiro, C., Smith, R., Huff, E., & Rhodes, J. 2018, [PASP](#), **130**, 065004
- Puetter, R., Gosnell, T., & Yahil, A. 2005, [ARA&A](#), **43**, 139
- Richardson, W. H. 1972, [JOSA](#), **62**, 55
- Robertson, J. G. 2017, [PASA](#), **34**, e035
- Rowe, B., Hirata, C., & Rhodes, J. 2011, [ApJ](#), **741**, 46
- Schmitz, M. A., Starck, J. L., Ngole Mboula, F., et al. 2020, [A&A](#), **636**, A78
- Schwartz, R. A., Torre, G., & Piana, M. 2014, [arXiv:1407.7343](#)
- Seshadri, S., Shapiro, C., Goodsall, T., et al. 2013, [PASP](#), **125**, 1065
- Shi, X., Guo, R., Zhu, Y., & Wang, Z. 2017, [J. Syst. Eng. Electron.](#), **28**, 1236
- Starck, J. L., Pantin, E., & Murtagh, F. 2002, [PASP](#), **114**, 1051
- Sureau, F., Lechat, A., & Starck, J.-L. 2020, [A&A](#), **641**, A67
- Tauber, J. A., Nielsen, P. H., Martín-Polegre, A., et al. 2019, [A&A](#), **622**, A55
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, [CSE](#), **13**, 22
- Viero, M. P., Moncelsi, L., Quadri, R. F., et al. 2013, [ApJ](#), **779**, 32
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, [NatMe](#), **17**, 261
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, [AJ](#), **140**, 1868
- Zemcov, M., Immel, P., Nguyen, C., et al. 2017, [NatCo](#), **8**, 15003
- Zemcov, M., Smidt, J., Arai, T., et al. 2014, [Sci](#), **346**, 732