



Coma Off It: Regularizing Variable Point-spread Functions

J. Marcus Hughes , Craig E. DeForest , and Daniel B. Seaton Southwest Research Institute, 1050 Walnut Street, Suite 300, Boulder, CO 80302, USA; deforest@boulder.swri.edu*Received 2022 December 5; revised 2023 March 16; accepted 2023 March 16; published 2023 April 18*

Abstract

We describe a rapid and direct method for regularizing, post facto, the point-spread function (PSF) of a telescope or other imaging instrument across its entire field of view (FOV). Imaging instruments in general blur point sources of light by local convolution with a PSF that varies slowly across the FOV, due to coma, spherical aberration, and similar effects. It is possible to regularize the PSF in post-processing, producing data with a homogeneous “effective PSF” across the entire FOV. In turn, the method enables seamless wide-field astronomical mosaics at higher resolution than would otherwise be achievable, and potentially changes the design trade space for telescopes, lenses, and other optical systems where data uniformity is important. For many kinds of optical aberration, simple and rapid convolution with a locally optimized “transfer PSF” produces extremely uniform imaging properties at low computational cost. PSF regularization does not require access to the instrument that obtained the data, and can be bootstrapped from existing data sets that include starfield images or other means of estimating the PSF across the field.

Unified Astronomy Thesaurus concepts: [Astronomy data analysis \(1858\)](#); [Photometry \(1234\)](#); [CCD observation \(207\)](#); [Optical telescopes \(1174\)](#); [Telescopes \(1689\)](#)

1. Introduction

Imaging instruments focus light onto a focal plane. Focusing is generally not perfect: in an ideal instrument, light from each specific direction would be steered to an equally specific location on the focal plane; but in real instruments, even perfectly collimated incident light is dispersed across a small pattern, the point-spread function (PSF), that varies slowly with respect to location in the focal plane. Each region of the resulting image is therefore a convolution of the actual data being measured, with the local PSF for that region of the instrument focal plane. The PSF is itself a convolution of image patterns from several effects, including diffraction at the aperture, scattered light, and aberrations caused by the imperfect nature of geometric optics.

Reducing optical aberration has been the dominant challenge of telescope optical design for some time (Kepler et al. 1611; King 2003). Even in the modern era, essentially all imaging instruments, from commercial camera lenses to the Hubble Space Telescope, demonstrate aberration of various kinds (e.g., Smith 2008; Khorin et al. 2022), including coma, spherical aberration, and position-dependent astigmatism among higher-order corrections. This introduces a slowly varying PSF that can be quite complicated in shape. Typical telescopes may be diffraction-limited near the center of the image plane, but have wide, and/or strongly anisotropic, PSFs near the edges of the field of view (FOV). This effect makes telescope resolution a slippery quantity: following the Rayleigh criterion of separability of point sources, an instrument may (and most do) have different spatial resolutions at different focal plane locations or in different directions at the same focal plane location.

Fortunately, the local interaction between a PSF and the underlying image data is well understood mathematically. The convolution theorem (e.g., Bracewell 2000) establishes a

correspondence between convolution in the spatial domain and multiplication in the Fourier domain; a corollary is that the set of all images approximates a commutative group under convolution: just as multiplication by nearly any number can be reversed through division, the image effects of convolution by nearly any PSF may be reversed through convolution by an inverse PSF, which is essentially division in the Fourier domain. Using this insight, one may reverse many PSF-widening effects, improving post facto the effective stray-light qualities of a telescope through simple transformations of its images; the technique has been used to improve the stray-light characteristics of solar telescopes in particular, improving photometry of dark features embedded in the bright solar corona (DeForest et al. 2009; Shearer et al. 2012). Although linear known-PSF deconvolution cannot fully compensate for effects, such as diffraction, that literally zero out particular Fourier components of the image, for some such cases approximate inverses may be constructed that improve image quality significantly, for example by partially suppressing diffraction lines from a mesh in front of the instrument (DeForest et al. 2009; Poduval et al. 2013). Further, many common optical aberration and stray-light “PSF halo” effects do not fully or even deeply suppress the corresponding Fourier components, and therefore can be removed through direct known-PSF deconvolution with minimal impact on the image signal-to-noise ratio (S/N)—even if the effects are strongly obtrusive in raw images (e.g., DeForest et al. 2009).

In addition to direct deconvolution, many techniques have been developed to transfer images from raw PSF to a desired, more homogeneous PSF that simplifies or even enables further analysis. PSF adjustment, with various methods, has been used on a whole-image basis to improve slitless imaging spectroscopy (Atwood & Kankelborg 2018), to better match starfield images to models (Alard & Lupton 1998), to improve data comparison across instruments with different PSFs (e.g., Boucaud et al. 2016), or simply to homogenize atmospheric seeing effects in astronomical surveys (e.g., Desai et al. 2012).



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Because optical aberration generally varies across the FOV, its effects are not susceptible to full-image linear deconvolution. However, typical aberration PSF effects such as coma vary “slowly,” in that the length scale of variation is much longer than the length scale of the PSF itself. This means that a sufficiently small region of interest may be deconvolved using linear methods (e.g., Jarvis et al. 2008). In the context of conventional photography, with a wider region of interest, variable known-PSF deconvolution has been used for full-frame image reconstruction using several commercial compound lenses and even a spherical singlet camera (Schuler et al. 2011). The methods described by Schuler et al. come at a high computational cost: hours of processing time for each several-megapixel image processed, in addition to the computation required to estimate the PSF itself. More recently, Plowman et al. (2022) have developed sparse-matrix deconvolution methods that offer some advantages and flexibility over those of Schuler et al. (2011), including both self-consistent treatment of PSF variation and also a framework with extensibility to encompass more difficult effects such as instrument spectral response or multiple viewing angles and/or constraints (Plowman 2022). The computational cost remains high with sparse-matrix methods, because they must represent not merely a small number of kernels but a full PSF realization for every pixel in the original image plane.

A less rigorous but more practical approach to homogenizing PSF effects across an image is to treat the image as a collection of individual patches, each of which may be treated as having a separate uniform PSF. Methods of this sort have been used in a solar physics context to overcome aberration from atmospheric seeing (e.g., Wöger 2007; Scharmer et al. 2010) and, more recently, to address aberration in the IRIS solar UV spectrograph (Wuelser et al. 2020).

In developing the Polarimeter to UNify the Corona and Heliosphere (PUNCH) mission (DeForest et al. 2022), we have found that a straightforward method of homogenization, patched image-neighborhood known-PSF transfer deconvolution, can clean several-megapixel images in a few seconds; this enables routine use of PSF regularization as part of the PUNCH data-reduction pipeline, and also as a part of other astronomical or heliophysical image-analysis efforts. The technique is a three-step process: (i) generate a PSF model from bright stars in each subregion of the image plane, which may use a single image or a suite of images to ensure enough stars in each patch of image; (ii) perform analytic known-PSF deconvolution of each image subregion to a target PSF, using the collection of real PSFs determined in the first step; (iii) merge the image patches using a soft windowing function, such as the Hann window, to minimize boundary effects in the final reconstructed image. This sequence is directly applicable to astronomical imaging, sufficiently computationally inexpensive to be applied to survey data or time-dependent data sets, such as image or spectral-image sequences of the solar corona, and suitable for generating uniform-PSF mosaics from images acquired by either one camera, as in a sky survey, or many cameras with heterogeneous PSF characteristics, as for PUNCH. In particular, simple known-PSF transfer deconvolution is useful for scrubbing PSF “wings” and other nonideal artifacts from images affected by optical aberration, without excessively amplifying noise in any part of the images.

In Section 2, we briefly introduce the mathematical basis of known-PSF deconvolution and “transfer PSFs.” In Section 3,

we describe the numerical method and sample code to carry out the operation. In Section 4, we demonstrate deaberration of model and actual starfield data. In Section 5, we discuss the implications of routine post facto homogenization for design of new instruments and analysis of existing data. Finally, in Section 6, we recap and draw conclusions.

2. Theory of Point-spread Function Deconvolution and Windowing

The convolution theorem relates the operations of convolution and multiplication. If I and K are images (i.e., $I, K: \mathbb{Z}^2 \rightarrow \mathbb{Q}$ with pixel indices indicated by subscripts, so that $I_{i,j} \in \mathbb{Q}$ for some range of i and j), then

$$(I \otimes K) = \mathcal{F}^{-1}[\mathcal{F}(I) \mathcal{F}(K)], \quad (1)$$

where \mathcal{F} represents the discrete Fourier transform, \otimes represents convolution, and juxtaposition of terms represents element-wise multiplication (e.g., Bracewell 2000). Deconvolution may therefore be accomplished through element-wise division in the Fourier plane. Using \mathcal{I} for $\mathcal{F}(I)$, etc., given a measured $I_K \equiv (I \otimes K)$, we can immediately write

$$I = \mathcal{F}^{-1} \left[\frac{\mathcal{I}_K}{\mathcal{K}} \right] = \mathcal{F}^{-1} \left[\mathcal{I}_K \frac{\mathcal{K}^*}{|\mathcal{K}|^2} \right], \quad (2)$$

where $\mathcal{I}_K \equiv \mathcal{F}(I_K)$ and $\mathcal{K} \equiv \mathcal{F}(K)$ so that, if K is a PSF, then \mathcal{K} is the corresponding modulation transfer function (MTF) and the $*$ operator represents complex conjugation. In practice, most astronomical images have at least some zero-magnitude Fourier coefficients, so instead of the direct reciprocal one may use a regularized reciprocal \mathfrak{R} :¹

$$\mathfrak{R}_{\alpha,\epsilon}(\mathcal{A}) \equiv \frac{\mathcal{A}^* |\mathcal{A}|^{\alpha-1}}{|\mathcal{A}|^{\alpha+1} + \epsilon^{\alpha+1}}, \quad (3)$$

where ϵ is a smallness parameter and α controls how sharply the expression drops to zero for small values of \mathcal{A} . Then Equation (2) becomes

$$I' = \mathcal{F}^{-1}[\mathcal{I}_K \mathfrak{R}_{\alpha,\epsilon}(\mathcal{K})], \quad (4)$$

where the prime on I' indicates that it is a regularized approximation rather than an exact reconstruction of the ideal image I . Regularized reciprocals are a common tool used to approximate the solution of ill-posed inverse problems; a good treatment may be found in Claerbout (2014). Equation (4) is suitable to deconvolve images that have been convolved with a fixed PSF K . The ancillary parameters to $\mathfrak{R}_{\alpha,\epsilon}$ control the balance between direct analytic inversion of the original PSF and the need to avoid strongly amplifying noise.

Breaking a source image into small neighborhoods allows the inversion in Equation (4) to work for slowly varying PSFs, such as are caused by optical aberration. The neighborhoods can be deconvolved by individually tailored PSFs K_i for the i th neighborhood. However, the optimally achievable degree of deconvolution will in general vary with the PSF across the image plane.

To achieve PSF uniformity, one may convolve the deconvolved images by a band-limited “target PSF” P , with

¹ Note the distinction between a “regularized reciprocal,” which approximates a true reciprocal, and “PSF regularization,” which improves and/or homogenizes a slowly varying image PSF.

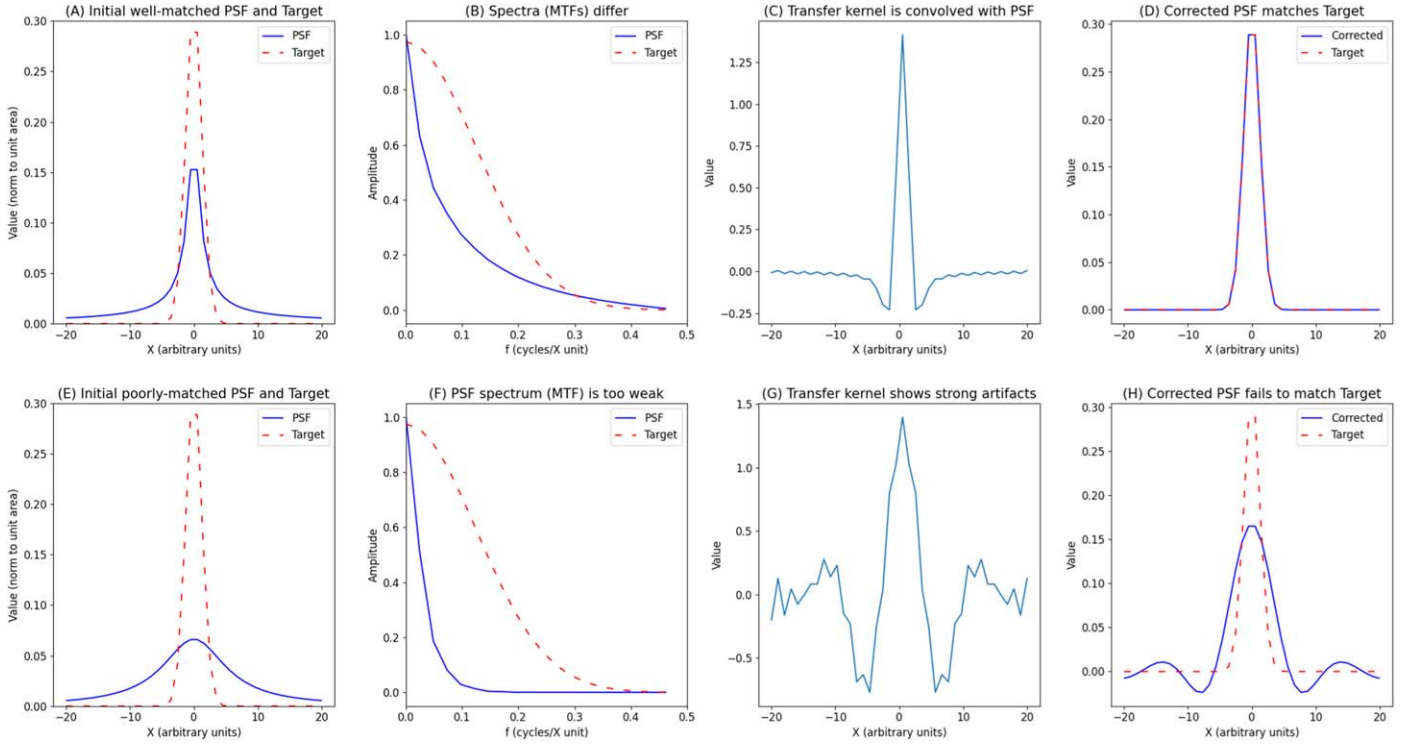


Figure 1. Sample deconvolutions in 1D demonstrate the effectiveness and limits of PSF correction. Top row: a PSF with broad wings (A) moderately attenuates Fourier coefficients at most spatial frequencies, compared to the target PSF (B). Scaling the MTF and Fourier transforming yields a transfer PSF that, when convolved with the original data, matches the target. Bottom row: a PSF with quite broad wings (E) attenuates the MTF too much for effective correction (F), as seen with the heavily artifacted transfer PSF (G) and poor match between the final effective PSF and the target (H). Note that no noise is added to the data in the figure.

the property that P must have zero Fourier coefficients nearly everywhere that each K_i does. Performing this operation in the Fourier domain makes clear the importance of band limiting. Zeroed Fourier coefficients in the band-limited target PSF yield zero coefficients in each deconvolved image:

$$I'_{P,i} = \mathcal{F}^{-1}[\mathcal{I}_{K,i} \mathfrak{R}_{\alpha,\epsilon,\mathcal{P}}(K_i) \mathcal{P}], \quad (5)$$

where the i subscripts run across neighborhoods of the original image, \mathcal{P} is the desired, constant, band-limited MTF (associated with a target PSF P) for all neighborhoods, and $I'_{P,i}$ is the regularized image neighborhood with the desired uniform PSF applied. The new quantity $\mathfrak{R}_{\alpha,\epsilon,\mathcal{P}}$ is just

$$\mathfrak{R}_{\alpha,\epsilon,\mathcal{P}}(\mathcal{A}) \equiv \frac{\mathcal{A}^* |\mathcal{A}|^{\alpha-1}}{|\mathcal{A}|^{\alpha+1} + (\epsilon|\mathcal{P}|)^{\alpha+1}}, \quad (6)$$

in which ϵ is multiplied by $|\mathcal{P}|$ to scale the “smallness threshold” to match the band-limited target PSF, limiting overall amplification to roughly ϵ^{-1} .

The expression $\mathcal{F}^{-1}[\mathfrak{R}_{\alpha,\epsilon,\mathcal{P}}(K_i) \mathcal{P}]$ is interesting, because it is convolved with $I_{K,i}$ in Equation (5) (implicitly, via Equation (1)), to obtain $I'_{P,i}$. We call this expression a “transfer PSF” and it can be explicitly convolved with the original image to shift its intrinsic PSF from K_i to P .

Figure 1 illustrates a correct and incorrect use of transfer PSFs to regularize a 1D PSF. In panels (A)–(D), data in the form of a delta function have been convolved with a complex curve, based on a Lorentzian, with sharp core and broad wings (bold purple solid line), compared to a desired target PSF with slightly broader core and truncated wings (light red dashed line). Panel (A) shows both this original PSF and a target Gaussian PSF. Panel (B) shows the MTF of each curve; while

the target MTF is higher than the original, the ratio of the two quantities does not exceed roughly 3 for any spatial frequency. Panel (C) shows a derived transfer PSF, calculated using Equation (4). Panel (D) shows the result of convolving the original PSF with the transfer PSF in panel (C), matching the target PSF.

In panels (E)–(H) of Figure 1, the same process is attempted as in panels (A)–(D) but with an original PSF whose core is quite broad compared to the target Gaussian curve (and which therefore has less high-spatial-frequency content). The original MTF is essentially zero for a good fraction of the frequency range in panel (F). Full transfer deconvolution would require amplifying high-frequency Fourier components by large factors, and the calculated transfer kernel (G) has many strong artifacts related to regularization of these components’ reciprocals. In panel (H), the “corrected” PSF does not match the target PSF, because the reciprocal regularization prevents strong amplification of the missing or near-missing Fourier components (above 0.1 cycles unit⁻¹, in panel (F)).

In practice, aberration-driven image PSFs retain a sharp core even though the wings might be wide or irregular; these PSFs preserve high spatial frequencies, and therefore afford PSF regularization to an appropriately chosen fixed target PSF that approximately preserves the resolution of the instrument as a whole, without greatly amplifying a large set of the corresponding Fourier components.

3. A Numerical Method for Point-spread Function Regularization

The method described in Section 2 is implemented and available as a Python package (with Cython core implementation for speed-up) called `regularizePSF` (Hughes 2022). Here

we define this reference algorithm, then discuss the parameters needed for the algorithm and how to pick values for them.

3.1. Description of Algorithm

We perform PSF regularization on a collection of overlapping neighborhoods, which are recombined at the end of the process to produce a homogenized image. The neighborhoods are selected to be small enough that the PSF is essentially invariant across each one, but large enough that the PSF itself is small compared to the neighborhood. The neighborhoods overlap, to ensure smooth splicing in the final recombination. The algorithm is divided into steps as in the subsections below.

3.1.1. Determine a Point-spread Function

For many astronomical applications, including the PUNCH mission, the scene to be imaged is the dark sky and hence contains many stars. Because stars are point sources, we use them to measure the PSF in each neighborhood of the instrument FOV.

We begin with one or more starfield images with the full FOV of the instrument. We break each image into neighborhoods of size $N \times N$ pixels, with interneighborhood spacing of $N/2$ pixels. This results in sampling each pixel four times into four separate neighborhoods. For each neighborhood, we identify bright stars in the FOV, using a conventional unsharp-mask-and-threshold algorithm, and resample each identified star, together with a small patch of pixels around it, into a common coaligned image patch of size $M \times M$ pixels (with $M < N$). These identified star image patches are normalized to the same overall flux and then aggregated, using a pixel-wise median, into a composite stellar image which serves as a PSF model for the neighborhood. We use the pixel-wise median, rather than the pixel-wise mean, to eliminate nearby stars in the small patches around any of the bright stars. This technique makes use of an implicit assumption that less than half of each star’s nearby pixels are affected by asterisms or nebulosity. Distributed “objects” such as nebulosity, zodiacal light, the galaxy or instrumental stray light may be a concern; using a smaller-percentile sample than the median can help remove these sources also, or they can be removed via smooth-background estimators such as minsmooth (DeForest et al. 2016). Stars near the edge of each neighborhood are sampled in their entirety, even if some peripheral pixels of the corresponding patch are outside the neighborhood boundary; this eliminates boundary effects that otherwise might be associated with each neighborhood.

Optionally, the accumulated PSF images can be further refined using a parametric fit. Fitting greatly reduces the residual noise in each PSF image, at the cost of forcing the model PSF to conform to a priori expectations of the form of the fitting function. Alternatively, one could fit a functional model to either the centered star patches or the averaged stars. Using a functional model may decrease noise in the modeled PSFs and allow for interpolating to fill in regions with insufficient star patches to extract a viable data-driven PSF model. However, it requires determining the correct functional form of the PSF and carefully fitting. Parametric fitting is therefore deprecated when noise levels and star counts do not preclude direct aggregation. Should stellar counts be too low, multiple images from the same instrument may be sampled in each neighborhood, especially with slightly varying pointing,

to increase the population of stellar images in each neighborhood. In applications where stars are not available, other PSF-estimation methods, of which there are many (e.g., Chalmond 1991; Starck et al. 2002), may be used instead of direct stellar aggregation or fitting.

3.1.2. Define a Target Point-spread Function

We define a target PSF analytically and then evaluate it to create a same-sized patch to the measured (or modeled) local PSF patches. In the examples in Section 4 we utilize a symmetric 2D Gaussian as the target PSF model.

The target PSF must be comparable in width to the core of the measured PSFs. To design a target PSF, one must inspect the ratio of the two corresponding MTFs, $\mathcal{P}/\mathcal{K}_i$. If the ratio is large at any location in the Fourier plane, then P is not a suitable target PSF. The goal of target PSF design is to balance the specific application’s need for PSF narrowness against the need to minimize high amplification factors in any one Fourier component.

3.1.3. Apodize Each Neighborhood

Before deconvolving each neighborhood, we apodize it by multiplying it with a root-Hann window:

$$w(x, y) = \sin\left((x + 0.5)\frac{\pi}{n_x}\right)\sin\left((y + 0.5)\frac{\pi}{n_y}\right), \quad (7)$$

where x and y are the current pixel coordinates in the patch and n_x and n_y are the size in pixels of the patches in each dimension.

The root-Hann window is applied twice: once before and once after deconvolution. This minimizes edge effects for the Fourier transform used in the deconvolution, and also produces Hann window profiles in the final homogenized neighborhoods, simplifying recombination of the neighborhoods (Section 3.1.6).

Apodization is important, in particular, because most Fourier transformation algorithms impose periodic boundary conditions on the data being transformed. Combining Hann windows² with 50% overlap is a straightforward and common way to segment continuous data, while minimizing spectral spreading from the apodization window itself, avoiding boundary effects from the deconvolution step (Section 3.1.4) and simplifying the recombination of the segments into a continuous whole after processing (e.g., International Organization for Standardization 1993; Shlien 1994; DeForest 2017). While many Fourier methods ignore apodization entirely (e.g., International Organization for Standardization 1994; Handy et al. 1999; Wuelser et al. 2020), such methods are known to introduce neighborhood-boundary artifacts which are detectable in the resulting images whether or not they are directly visible to a user (Schrijver et al. 1999; Bianchi & Piva 2011).

3.1.4. Deconvolve the Neighborhoods

Using the local PSF models and the target PSF for each neighborhood, we deconvolve as in Section 2. In particular, we Fourier transform the apodized neighborhood, multiply by the MTF ratio as in Equation (5), then inverse Fourier transform to produce regularized neighborhoods.

² Note that many authors have mistakenly referred to the Hann window (Kahlig 1993) as a “Hanning” window—a common eggcorn of the related, but different, “Hamming” window (e.g., Podder et al. 2014).

3.1.5. Apply a Second Apodization

Because multiplication in Fourier space is equivalent to convolution, the original apodizing window is deconvolved along with the actual image. This produces an incorrect windowing function in the modified neighborhood, as described by DeForest (2017). To minimize this effect, we apodize a second time with the root-Hann window as in Section 3.1.3. This yields an overall Hann window envelope on each neighborhood, subject to small perturbations from the deconvolution operation.

3.1.6. Merge the Neighborhoods

Neighborhoods are remerged by direct addition, with the appropriate pixel phase, back into a new image array that begins with zero-valued pixels. Using the Hann window automatically weights each neighborhood to reconstitute the original image without further weighting or scaling to remove the windowing function, because of the identity that $\sin(x)^2 + \cos(x)^2 = 1$.

3.2. Algorithm Parameters

The basic PSF-regularization algorithm has several key parameters, described in the paragraphs below.

The particular target PSF is adjustable. Ideally, the target PSF should be well localized to preserve resolution, with a core that is no narrower than the core of the widest neighborhood-PSF core in the source image plane. A narrow, isotropic, normalized Gaussian is a good choice. In the real-world examples we considered (Section 4), a 3 pixel FWHM proved to be a reasonable minimum width.

The size of each neighborhood must be selected based on the source image properties. The size should be small enough that the PSF does not change appreciably within each neighborhood, but large enough that edge effects from the PSF itself do not dominate the deconvolution operation. We found that one-eighth to one-sixteenth the original image size was a suitable range in each of our example images.

The ϵ and α parameters of the $\mathfrak{R}_{\epsilon,\alpha}$ operator set, respectively, the “smallness” criterion and the hardness of the transition from reciprocal to zero. The maximum amplification of each Fourier component is roughly (though in general somewhat less than) $1/\epsilon$, so a value of $\epsilon = 0.1$ allows Fourier components to be amplified by roughly $10\times$. ϵ should be set based on the original S/N of the source image, and the acceptable S/N in the final image. In our examples, we maintained $\epsilon = 0.1$. The α parameter sets the hardness of the transition from amplification to attenuation of weak components. We used $\alpha = 10$; with $\epsilon = 0.1$, this allowed up to approximately $9\times$ amplification of each Fourier component if needed.

4. Sample Applications of Point-spread Function Regularization

We applied PSF regularization to a variety of images to demonstrate its effect.

4.1. Model Starfield Data

We first homogenized a simple model starfield to demonstrate the method in a noiseless setting under controlled conditions.

First, we generated point source “stars” in a regular square grid across a 256×256 pixel image. Next, we separately convolved each star with an individual synthetic instrument PSF calculated from a parametric analytic expression, to create a synthetic observation (Figures 2(a) and (b)). We applied the algorithm of Section 3.1 to the image, with a 3 pixel FWHM Gaussian target PSF, resulting in a homogenized corrected image that matches the convolution of the initial image in panel (A) with the target PSF (Figures 2(c) and (d)). This demonstrates reconstruction of a uniform-PSF image of moderate resolution (panel (d)) from a modeled variable-PSF image (panel (b)).

The details of the synthetic PSF are described in the Appendix.

4.2. The Demonstration Airglow-Subtracted Heliospheric Imager

The Demonstration Airglow-Subtracted Heliospheric imager (DASH) was a prototype instrument to demonstrate wide-field heliospheric imaging from near Earth. DASH was deployed approximately 10 meters from Earth, on a rooftop in the Colorado Rocky Mountains for 90 days in 2014. It was used to test deep subtraction of multiple foregrounds/backgrounds including atmospheric effects and the starfield (DeForest & Howard 2015; DeForest 2015). The instrument comprised an automated observatory housing, a corral-type baffle, a commercial wide-angle lens, a commercial cooled-CCD camera, and control electronics. DASH collected 30 s cadence images nightly in the western sky. Data reduction involved very precisely tracking all visible stars as they crossed through the DASH FOV over the course of each night’s 3 hr observing “run.”

The Nikon wide-angle lens exhibited significant aberration at wide angles, complicating analysis, as shown in Figure 3. While the central portion of the image is close to ideally sampled (with a PSF core FWHM of approximately 2 pixels), the edges of the FOV exhibit some spherical aberration, astigmatism, and broad sagittal rays; this meant that the shape of a star changed considerably as it crossed the FOV. In the 2015 analysis, the DASH team blurred the starfield to a FWHM of 7.5 pixels to render these artifacts irrelevant to the analysis. We now demonstrate repair of the DASH starfield to a uniform target Gaussian PSF with a FWHM of just 3.5 pixels, increasing effective resolution by a factor of 2 while maintaining PSF uniformity across the full FOV.

We used a single DASH image to build the PSF model. SEP is a Source Extraction and Photometry library written in Python by Barbary (2016). We used SEP to identify stars in the image. We broke the image into 256×256 neighborhoods, nine of which are shown in Figure 3(a). We stacked patches around all the stars centered in each neighborhood, to produce a median patch PSF image. We used this composite patch image as a PSF model to correct the image, with $\alpha = 2$, $\epsilon = 0.3$, and a symmetric Gaussian target PSF with a FWHM of 3.5 pixels. Figure 3(b) shows the result of the PSF regularization, which greatly reduced the aberrations from the commercial lens.

4.3. The Polarimeter to Unify the Corona and Heliosphere Mission

In 2022 June, an engineering model (EM) of the PUNCH Wide Field Imager (WFI) camera was taken to McDonald

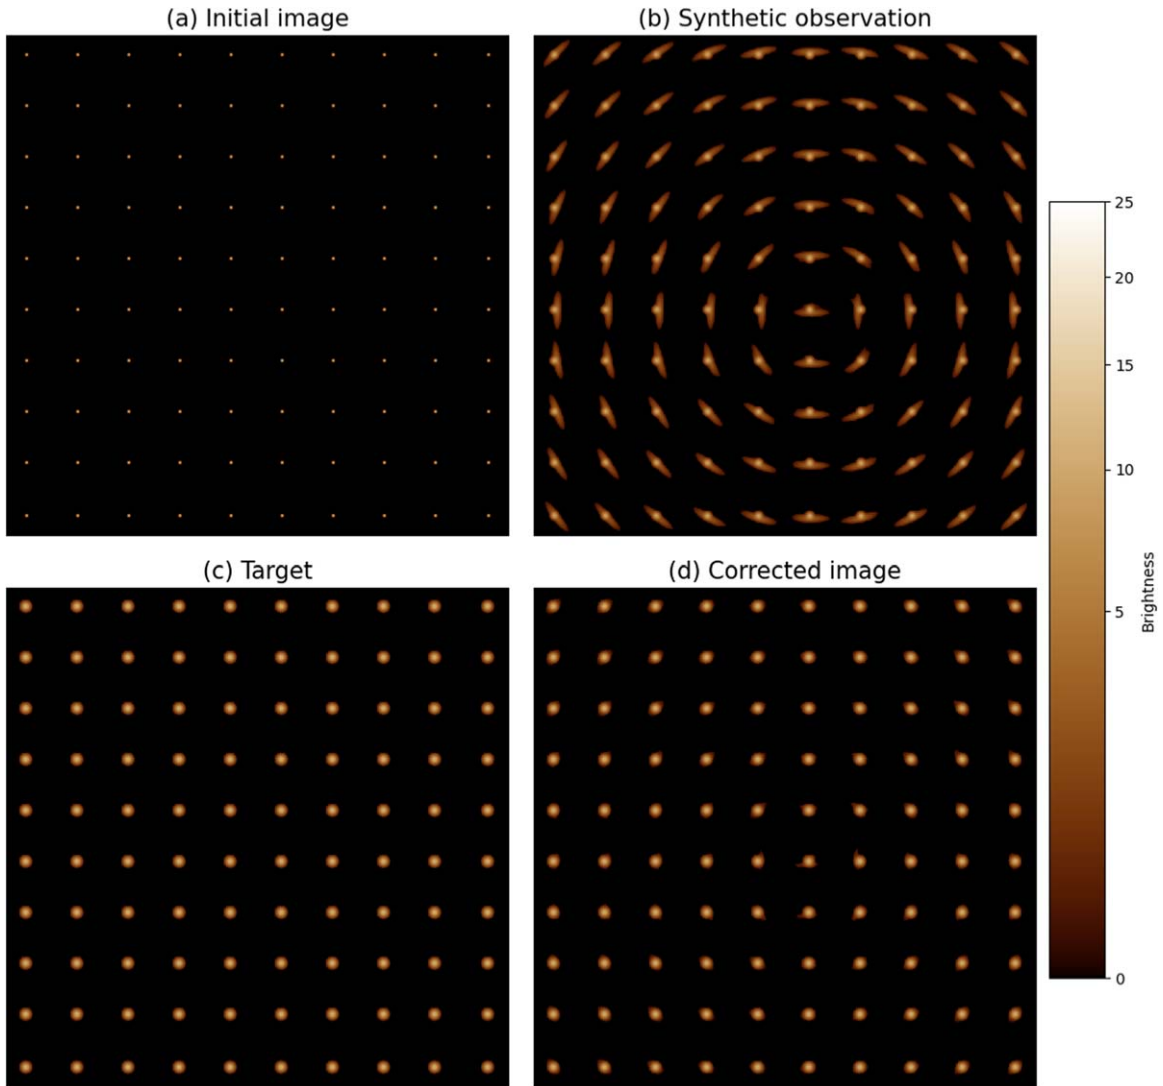


Figure 2. Model data demonstrating PSF homogenization. An initial image of a simplified starfield (a) is modified with a slowly varying PSF (b), then homogenized to a target PSF (c) with the method of Section 3.1. The final image (d) visually matches the target in (c). The panels are gamma-corrected to highlight the periphery of the model PSFs.

Observatory to capture dark-sky images for calibration and test. Five-hundred and forty of these images are used in this study. While these data suffer from high noise as the EM WFI camera was operating at roughly 30 °C, far from the in-flight design temperature of -60 °C, the ensemble data set was sufficient to reveal the variation of the PSF across the FOV.

We broke the WFI images into 256×256 pixel neighborhoods, and accumulated PSF images by coaligning stars within each neighborhood of the FOV, as described in Section 3.1.1, across 50 images in the data set. Similar to Section 4.2, we utilized the SEP library to identify and extract stars. Once we had extracted patches for all the selected stars imaged in each neighborhood, we used the pixel-wise median of all the patches to create an instrument PSF model for that neighborhood. Figure 4(a) shows a sampling of these PSF models for 16 neighborhoods in the WFI FOV.

We then corrected the original images as described in Section 3. We used a Gaussian target PSF with $\text{FWHM} = 2.25$ pixels, $\alpha = 3$, and $\epsilon = 0.3$. Finally, we reassembled new PSF images from the corrected data. Figure 4(b) shows the PSFs after regularization. The wings of the processed PSFs are

reduced in amplitude by a factor of order 30 compared to the original data.

Some residual PSF effects are visible in the peripheral neighborhoods of Figure 4(b); we attribute these to the high noise level in the source images and smaller number of samples of each star near the edges of the original FOV.

5. Discussion

We have demonstrated spliced linear PSF regularization to homogenize variable telescope PSFs in astronomical applications, using direct starfield imaging to recover and deconvolve for variable PSF across the FOV. The method is useful for applications where uniform imaging over a wide field is more important than the highest possible spatial resolution in a small portion of each image. Further, it may improve future telescope design, by allowing designers to trade resources between telescopes with inexpensive optics (e.g., spherical mirrors in a Cassegrain telescope) and a large aperture, versus telescopes that minimize point-spread variability with more expensive optics (e.g., aspheres in a Ritchey–Chretien telescope) with a smaller aperture at the same price point.




Figure 3. Image patches in a DASH image before and after PSF homogenization. Left: stars in each region exhibit a varying PSF before correction. Right: after application of a transfer PSF, the stars are more uniform across the image. The panels are gamma-corrected to highlight the periphery of the original and homogenized PSFs.




Figure 4. Average stars in PUNCH image regions before and after PSF homogenization. Left: accumulated average star in each region (see text) exhibits a varying PSF before correction. Right: after application of a transfer PSF, the stars are uniform across the image and the PSF wings are reduced by a factor of order 30–100. The image is gamma-corrected to highlight the wings of each PSF.

We used a simple PSF-correction method with batch processing of neighborhoods in a source image, each of which we treated as having a single PSF susceptible to direct linear inversion. Direct linear treatment is equivalent to construction of and convolution by a neighborhood-specific “transfer PSF” that converts the image in each neighborhood from its source PSF to a target PSF. The major advantage of this method over more rigorous techniques is its efficiency: two to three orders of magnitude less computation is required than for existing methods that consider a smoothly varying PSF.

Direct linear known-PSF deconvolution has known limitations. In particular, it is equivalent to amplifying particular Fourier components in the source image. While deconvolution to sharpen an image beyond the core of its original PSF is fraught with peril (because it involves strongly amplifying many high-spatial-frequency components, along with any image noise contained in them), deconvolution to remove PSF wings has a far more modest effect on the noise level in the underlying data. This is because PSFs with broad wings and a sharp core do not strongly attenuate most image Fourier

Table 1
Parameters for Variable-instrument PSF Model

Parameter	Value
h_{core}	25
$x_{c,\text{core}}$	16
$y_{c,\text{core}}$	16
$\sigma_{x,\text{core}}$	1
$\sigma_{y,\text{core}}$	1
θ_{core}	0
h_{tail}	8
$x_{c,\text{tail}}$	$1.5 \cos(\arctan((y - 16)/(x - 16)) + 45^\circ) + 16$
$y_{c,\text{tail}}$	$1.5 \sin(\arctan((y - 16)/(x - 16)) + 45^\circ) + 16$
$\sigma_{x,\text{tail}}$	5
$\sigma_{y,\text{tail}}$	1
θ_{tail}	$\arctan(y/x)$

components (DeForest et al. 2009; Poduval et al. 2013), and because we are deconvolving not to an ideal 1 pixel wide PSF but to a target PSF with significant width, so that the final images are slightly oversampled by the pixel grid. Oversampling by 50% or more is known to improve photometry by suppressing artifacts related to feature placement relative to the pixel grid. It also further band-limits the final image to reduce amplification of high-spatial-frequency Fourier components. In our examples, well under half of the Fourier components of each image are multiplied by as much as a factor of 2, so that overall image noise is nowhere increased by a factor greater than 1.5 and generally much less.

Deconvolution in general and PSF regularization in particular work by comparative photometry across image pixels. They are therefore limited by the source images' photometric quality. Calibration (flat-field) errors and/or detector nonlinearity induce errors that are amplified by the PSF-regularization process.

PSF regularization cannot significantly narrow the PSF core at any location in the source image, so its best use is to produce more uniform images, compromising optimal resolution at the center of the image in favor of a uniform PSF across the entire FOV. Potential applications include (i) forming mosaics from multiple source images, which can be homogenized for absolute uniformity of the resulting PSF to reduce boundary effects in the final mosaic; (ii) survey studies in which a large ensemble of features is detected either manually or automatically over a wide FOV or even across multiple instruments with different imaging characteristics, to eliminate PSF-based variations in the survey; (iii) background-removal problems such as the heliospheric imaging problem, in which the steady starfield must be removed to reveal the much fainter dynamic solar wind; (d) telescope optimization, in which a larger-aperture telescope with spherical or other inexpensive optics may become preferable to a smaller telescope with more sophisticated (hence expensive) aspheric optics.

The technique we present here is far from the only PSF-correction method that exists. Schuler et al. (2011) demonstrated PSF homogenization using second-order techniques, and Plowman et al. (2022) have recently described a more general, but considerably more computationally intense, method of PSF homogenization. Plowman's method uses sparse-matrix inversion techniques to homogenize the PSF of the SPICE extreme-UV spectral imager on board ESA's Solar Orbiter mission. The present method is more specific to direct

imaging in the particular (and most common) case of a slowly varying PSF that may be treated locally with linear methods, and is 1.5–2.5 orders of magnitude less computationally intense than the more exact Schuler et al. (2011) or Plowman et al. (2022) methods.

6. Conclusions

We have demonstrated a direct method of PSF homogenization that can equalize the PSF performance of a telescope or other optical system whose imaging characteristics vary slowly across the FOV, and that is sufficiently robust and rapid to apply to large data sets. Homogenized images take on the appearance of having been collected by a single ideal imager with uniform PSF, at modest cost in S/N (a few percent increase in noise level for the examples we demonstrate here). The method will be deployed for the upcoming PUNCH mission, which benefits from the extreme uniformity in modeling the starfield as a fixed background to be removed from that mission's images. PSF homogenization is useful for many astronomical and terrestrial applications in which uniformity of imaging performance over a wide field is more important than absolute maximal resolution within a small subfield, and a reference implementation of this rapid method is available for use as an open source package via GitHub³ or Zenodo⁴.

This work was funded through PUNCH, a NASA Small Explorer mission, via NASA grant No. 80GSFC18C0014. The authors also thank the anonymous referee for helping improve the article.

Appendix

Constructing the Analytic Model Point-spread Function

In Section 4.1, we constructed a synthetic instrument PSF from two Gaussians: a symmetrical Gaussian for the core brightness of the star and a second elliptical Gaussian to model aberration effects. Here, for completeness, we reproduce the exact expressions used to produce the PSFs in Figure 2 in the main text.

The general Gaussian function is

$$g(x, y; a, b, c, x_c, y_c, h) = h e^{-(ax'^2 + 2bx'y' + cy'^2)}, \quad (\text{A1})$$

where (x_c, y_c) is the center of the Gaussian with height h , and $x' \equiv x - x_c$, $y' \equiv y - y_c$. The free parameters a , b , and c correspond to a positive definite matrix. More specifically,

$$\theta = \frac{1}{2} \arctan\left(\frac{2b}{a - c}\right), \theta \in [-45^\circ, 45^\circ], \quad (\text{A2})$$

$$\sigma_x^2 = \frac{1}{2(a \cdot \cos^2 \theta + 2b \cdot \cos \theta \sin \theta + c \cdot \sin^2 \theta)}, \quad (\text{A3})$$

$$\sigma_y^2 = \frac{1}{2(a \cdot \sin^2 \theta - 2b \cdot \cos \theta \sin \theta + c \cdot \cos^2 \theta)}, \quad (\text{A4})$$

where θ is the rotation angle of the elliptical Gaussian, and σ_x^2 and σ_y^2 are the variances.

For our instrument PSF, we used

$$f(x, y) = g_{\text{core}}(x, y) + g_{\text{tail}}(x, y). \quad (\text{A5})$$

³ <https://github.com/punch-mission/regularizepsf>

⁴ <https://zenodo.org/badge/latestdoi/555583385>

Both the $g_{\text{core}}(x, y)$ and $g_{\text{tail}}(x, y)$ have a full complement of parameters as described in Table 1. The PSF windows are 32×32 pixels. We used $\alpha = 3$ and $\epsilon = 0.3$. The results are shown in Figure 2 in the main text.

ORCID iDs

J. Marcus Hughes  <https://orcid.org/0000-0003-3410-7650>

Craig E. DeForest  <https://orcid.org/0000-0002-7164-2786>

Daniel B. Seaton  <https://orcid.org/0000-0002-0494-2025>

References

- Alard, C., & Lupton, R. H. 1998, *ApJ*, 503, 325
- Atwood, S., & Kankelborg, C. 2018, *JATIS*, 4, 028002
- Barbary, K. 2016, *JOSS*, 1, 58
- Bianchi, T., & Piva, A. 2011, in 2011 IEEE Int. Workshop on Information Forensics and Security (2011(New York: IEEE), 1
- Boucaud, A., Bocchio, M., Abergel, A., et al. 2016, *A&A*, 596, A63
- Bracewell, R. N. 2000, *The Fourier Transform and its Applications* (Redwood City, CA: Stanford Univ. Press)
- Chalmond, B. 1991, *CVGIP*, 53, 364
- Claerbout, J. 2014, *Geophysical Image Estimation by Example* (Morrisville, NC: Lulu Press)
- DeForest, C., Killough, R., Gibson, S., et al. 2022, in IEEE Aerospace Conf. (AERO) (New York: IEEE), 1
- DeForest, C. E. 2015, arXiv:2212.02594
- DeForest, C. E. 2017, *ApJ*, 838, 155
- DeForest, C. E., & Howard, T. A. 2015, *ApJ*, 804, 126
- DeForest, C. E., Martens, P. C. H., & Wills-Davey, M. J. 2009, *ApJ*, 690, 1264
- DeForest, C. E., Mattheaus, W. H., Viall, N. M., & Cranmer, S. R. 2016, *ApJ*, 828, 66
- Desai, S., Armstrong, R., Mohr, J. J., et al. 2012, *ApJ*, 757, 83
- Handy, B. N., Acton, L. W., Kankelborg, C. C., et al. 1999, *SoPh*, 187, 229
- Hughes, J. M. 2022, <https://github.com/punch-mission/regularizepsf>
- International Organization for Standardization 1994, *Digital Compression and Coding of Continuous-tone Still Images*, ISO 10918-1, Geneva, Switzerland
- International Organization for Standardization 1993, *Coding of Moving Pictures and Associated Audio for Digital Storage Media at up to About 1.5 Mbit/sec—Part 3: Audio*, ISO 11172-3, Geneva, Switzerland
- Jarvis, M., Schechter, P., & Jain, B. 2008, arXiv:0810.0027
- Kahlig, P. 1993, *Interactions Between Global Climate Subsystems: The Legacy of Hann* (Washington, DC: American Geophysical Union), 1
- Kepler, J., Galilei, G., & Pena, J. 1611, *Ioannis Kepleri* (Zurich: e-rara)
- Khorin, P. A., Porfirev, A. P., & Khonina, S. N. 2022, *Photonics*, 9, 204
- King, H. 2003, *The History of the Telescope* (Mineola, NY: Dover Publications), <https://books.google.com/books?id=KAWwzHIDVksC>
- Plowman, J. 2022, arXiv:2209.01753
- Plowman, J. E., Auchère, F., Cuadrado, R. A., et al. 2022, arXiv:2211.16635
- Podder, P., Khan, T. Z., Khan, M. H., & Rahman, M. M. 2014, *IJCA*, 96, 1
- Poduval, B., DeForest, C. E., Schmelz, J. T., & Pathak, S. 2013, *ApJ*, 765, 144
- Scharmer, G. B., Löfdahl, M. G., van Werkhoven, T. I. M., & de la Cruz Rodríguez, J. 2010, *A&A*, 521, A68
- Schrijver, C. J., Title, A. M., Berger, T. E., et al. 1999, *SoPh*, 187, 261
- Schuler, C. J., Hirsch, M., Harmeling, S., & Schölkopf, B. 2011, in 2011 Int. Conf. on Computer Vision, IEEE (Barcelona, Spain) (New York: IEEE), 659
- Shearer, P., Frazin, R. A., Hero, A. O. I., & Gilbert, A. C. 2012, *ApJL*, 749, L8
- Shlien, S. 1994, *ITB*, 40, 206
- Smith, W. J. 2008, *Modern Optical Engineering: The Design of Optical Systems* (4th edn; Maidenhead: McGraw Hill)
- Starck, J. L., Pantin, E., & Murtagh, F. 2002, *PASP*, 114, 1051
- Wöger, F. 2007, PhD thesis, Kiepenheuer Institute for Solar Physics, Freiburg
- Wuelser, J.-P., Löfdahl, M., Tarbell, T., & De Pontieu, B. 2020, IRIS Technical Notes 52, <https://iris.lmsal.com/documents.html>