Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Forward modeling linear variable filter data: High resolution PAH and Paschen-α maps from SPHEREx

Andrew Saydjari

saydjari@nmsu.edu

New Mexico State University

Connor Hainje

New York University

Benjamin Dodge

Stanford University

Philipp Frank

Max Planck Institute for Astrophysics

David Hogg

New York University

Susan Clark

Stanford University

Research Article

Keywords:

Posted Date: November 25th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11682608/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Forward modeling linear variable filter data: High resolution PAH and Paschen-α maps from SPHEREx

Andrew K. Saydjari [1,2,*] Connor Hainje [3] Benjamin Dodge [4,5] Philipp Frank [6,4,5] David W. Hogg [3,7,8] and Susan E. Clark [4,5]

[1] Department of Astronomy, New Mexico State University, Las Cruces, NM 88001 USA
[2] Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544 USA
[3] Center for Cosmology and Particle Physics, Department of Physics, New York University, 726 Broadway, New York, NY 10003, USA
[4] Department of Physics, Stanford University, Stanford, CA 94305, USA
[5] Kavli Institute for Particle Astrophysics & Cosmology, P.O. Box 2450, Stanford University, Stanford, CA 94305, USA
[6] Max Planck Institute for Astrophysics, Karl-Schwarzschild-Str. 1, 85748 Garching, Germany
[7] Max Planck Institute for Astronomy, Königstuhl 17, Heidelberg 69117, Germany
[8] Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Ave, New York, NY 10010, USA

ABSTRACT

Polycyclic aromatic hydrocarbon (PAH) emission is a powerful tool for the study of the interstellar medium (ISM) in both Galactic and extragalactic contexts. It can reveal ionization state, radiation field, metallicity, composition, and ISM structure as sculpted by processes like turbulence and stellar feedback. Here we use NASA's SPHEREx to produce clean, high-resolution PAH maps at spectral resolution R $\sim$ 35 (compared to WISE at R $\sim$ 2). Prior treatments of the SPHEREx linear variable filter data seek to convert them to more conventional band-averaged maps that can leave striping artifacts associated with spatially variable wavelength sampling, especially at emission lines. In contrast, we deliver a scalable forward-modeling method capable of generating high-quality all-sky SPHEREx emission-line maps at 8.″5 FWHM angular resolution (on 6.″4 pixels), while naturally producing maps without visible striping. We reformat the data into a pixel-level database, mask stellar contamination, and infill missing pixels with a conditional, data-driven, anisotropic Gaussian process. We apply our method to produce 3.3 μm and 3.4 μm PAH and Paschen-α emission maps of a set of fields spanning a range of Galactic and extragalactic conditions including the Rosette nebula, the Corona Australis molecular cloud, M33, and M31. We demonstrate the quality and synergies of our SPHEREx 3.3 μm PAH maps by comparisons with other available gas and dust tracer maps and achieve sensitivities up to $\sigma(I_{3.3}) = 2.7 \times 10^{-8}$ erg s^{-1} cm^{-2} sr^{-1}. We estimate that a two-year extension of the SPHEREx survey, with modest modifications to the scanning strategy, would match or exceed the PAH sensitivity of WISE over the whole sky.

Keywords: Astrostatistics (1882), Polycyclic aromatic hydrocarbons (1280), Interstellar line emission (844), Interstellar dust (836)

1. INTRODUCTION

The interstellar medium (ISM) is a chemically and dynamically rich environment. Polycyclic aromatic hydrocarbons (PAHs) are a ubiquitous class of interstellar organics that play an outsized role in the energetics of the ISM (A. G. G. M. Tielens 2008). PAHs are stochastically heated by the interstellar radiation field, and produce prominent ro-vibrational emission lines in the mid-infrared that are widely used as tracers of interstellar physics, both in the Milky Way and extragalactically (D. Baron et al. 2025). In the Milky Way, diffuse PAH emission is an exquisite tracer of high-angular resolution structure, enabling constraints on interstellar turbulence (M.-A. Miville-Deschênes et al. 2016; G. Vigoureux et al. 2026), the relative orientation of density and magnetic field structures on small scales (R. Córdova Rosado et al. 2024), and other physics that is encoded in the small-scale spatial structure of the ISM. These analyses require separating the diffuse PAH emission from other near-infrared emission, including from stellar and extragalactic sources. A. M. Meisner & D. P. Finkbeiner (2014) produced 12 μm diffuse emission maps from Wide-field Infrared Survey Explorer (WISE; E. L. Wright et al. 2010) data that have been widely used as tracers of diffuse PAH emission. In nearby galaxies, spatially resolved maps of PAH emission can be used alone or in conjunction with other tracers to study the ISM properties, star formation, and the radiation field across galactic environments. Table 1 shows a selection of PAH features and the types of molecules predominantly responsible for the observed emission.

Corresponding author: Andrew K. Saydjari
Email: aksaydjari@gmail.com
* Hubble Fellow

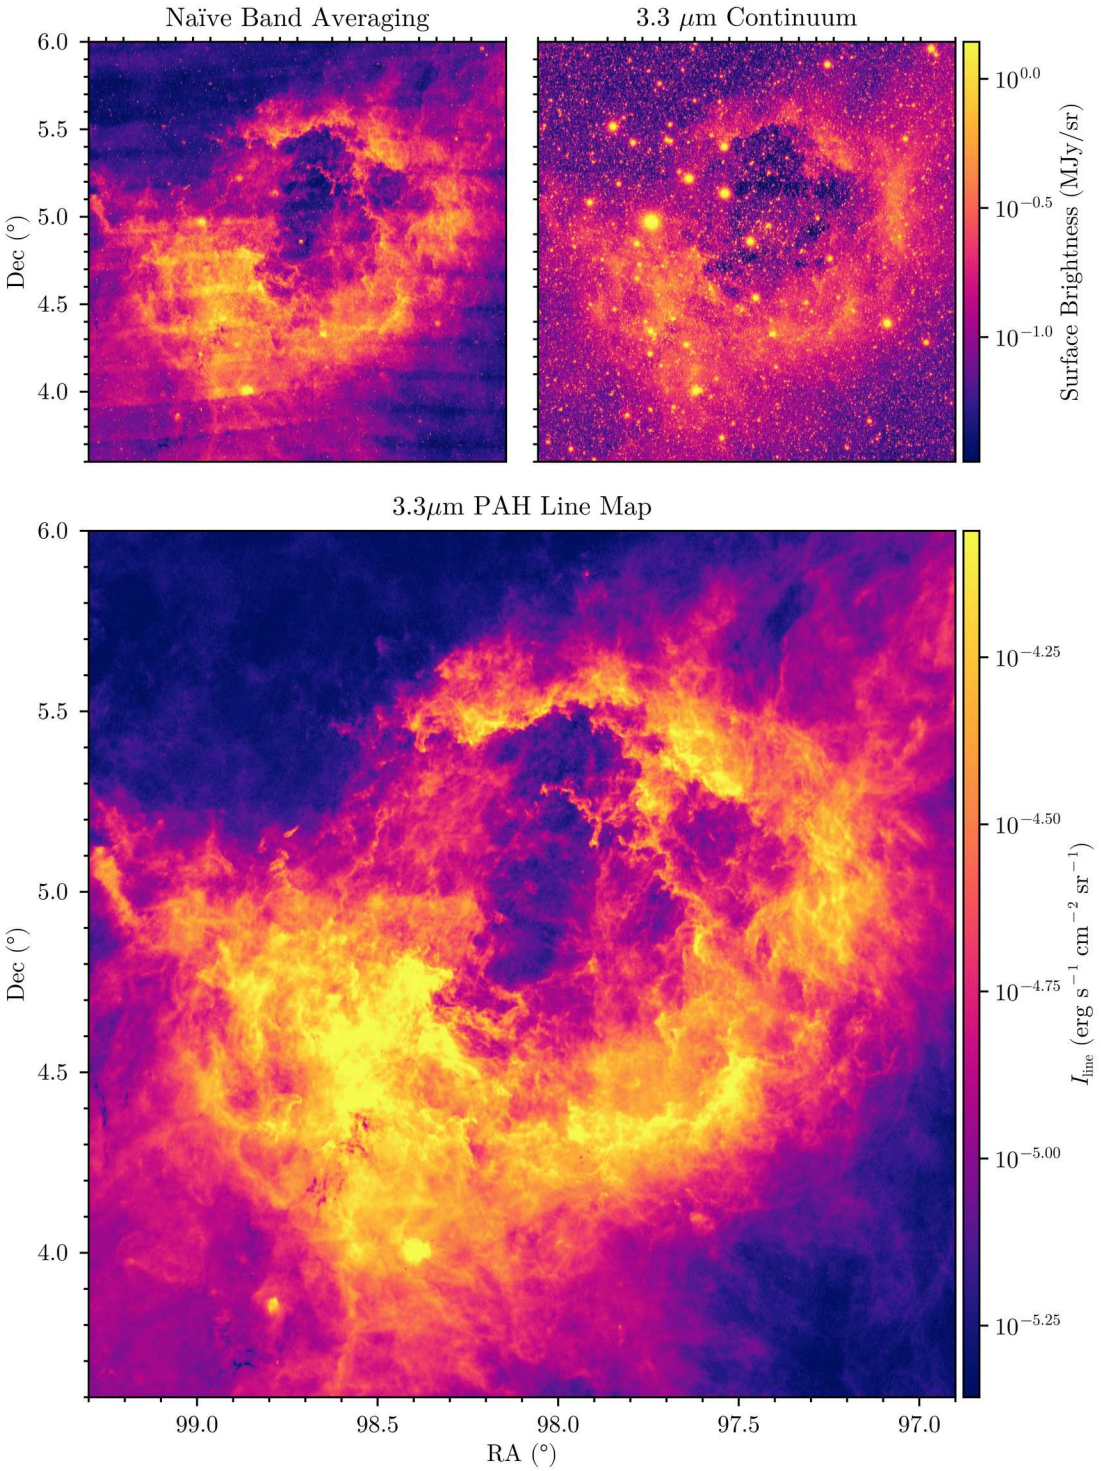

Figure 1. A demonstration of our forward-modeling approach applied to the Rosette nebula: (top right) surface brightness of our continuum model at 3.29 μm and (bottom) the 3.3 μm polycyclic aromatic hydrocarbon (PAH) line surface brightness. For comparison, a straightforward band average of observations near 3.3 μm minus continuum bands straddling the line is shown in the top left.

Table 1. Relevant PAH emission features compiled from B. T. Draine & A. Li (2007a) and A. G. G. M. Tielens (2008). Charge and size refer to the molecules predominantly responsible for the observed emission and are not a property of the features themselves. Our work maps the 3.3 and 3.4 μm features, but the other features are relevant for comparisons with maps from WISE and the Spitzer Infrared Array Camera (IRAC; G. G. Fazio et al. 2004) (see Section 4).

λ (μm)	Mode	Bond type	Charge	Size
3.3	C–H stretch	Aromatic	Neutral	Small
3.4	C–H stretch	Aliphatic	—	Small
6.2	C–C stretch	Aromatic	Ionized	Medium
7.7	C–C stretch + C–H bend	Aromatic	Ionized	Medium
8.6	C–H bend (in-plane)	Aromatic	Ionized	Medium
11.3	C–H bend (out-of-plane)	Aromatic	Neutral	Large
12.7	C–H bend (out-of-plane)	Aromatic	Neutral	Large

The Spectro-Photometer for the History of the Universe, Epoch of Reionization and Ices Explorer (SPHEREx) is a NASA Explorer satellite that is currently conducting an all-sky, near-infrared (0.75 to 5.0 μm) spectral survey (J. J. Bock et al. 2026). These data present rich opportunities for ISM science: the SPHEREx spectral range encompasses the bright 3.3 μm aromatic PAH feature, as well as the 3.4 μm aliphatic hydrocarbon feature, hydrogen recombination lines, several ro-vibrational transitions of molecular hydrogen, and ice absorption features. In this work, we develop a method for forward-modeling the SPHEREx data that enables high-fidelity mapping of diffuse-emission spectral features, and apply our method to construct high-resolution maps of the 3.3 μm and 3.4 μm PAH features and Paschen-α. For simplicity, we refer to this method as LVFit throughout the paper.

SPHEREx uses linear variable filter (LVF) spectrometers, which gradually vary the spectral transmission across the surface of the detectors. This means that each pixel in the resulting spectral image corresponds to a slightly different wavelength and bandwidth. The SPHEREx team defines 17 bands for each of the 6 LVFs, which divides the data into 102 spectral channels (J. J. Bock et al. 2026; SPHEREx Science Data Center 2026). Much of the work on tracing PAH emission with SPHEREx to date has operated in this paradigm, by binning the measurements into spectral channel maps (see Section 6.3). Naïve binning of the data can give rise to banding artifacts as shown in Figure 1 (zoomed in Figure 2), which are also prominently visible in the low-resolution SPHEREx PAH maps produced by R. Cepeda-Arroita et al. (2026) (see Section 6.3).

Here we produce large-area maps from many individual exposures using a forward-modeling method. Standard methods in optical astronomy for this kind of "mosaicing" or map-making traditionally involve interpolation and median (or similar) averaging of overlapping imaging; the current state of the art is drizzle (A. S. Fruchter & R. N. Hook 2002). These approaches are inappropriate here because the LVFs make it such that every pixel of the detector has (in principle) a slightly different spectral response. If images are made by binning pixels into spectral channels and combining by traditional methods, banding appears. Since the readout from any pixel of any detector in any good astronomical telescope is linearly related to the surface brightness on the sky, it is principled to make maps of the surface brightness by a classical "inverse problem" (M. Bertero & P. Boccacci 1998) or forward model, from which the map is found by optimizing an objective function based on a justified likelihood function.

These forward-model map-making methods have excellent properties: The output maps are linearly (or nearly linearly, depending on details) related to the input data, such that measurements made in the map are simply related to measurements made in the original raw data. The output maps can saturate the Cramér–Rao bound (C. R. Rao 1945; H. Cramér 1946) on the delivered image information. The input data can be arbitrarily heterogeneous in spatial structure, noise properties, and spectral response, provided that the likelihood function or forward model accounts for these properties correctly. Forward-model methods for map-making have had an extremely high impact in cosmic microwave background experiments (notably pioneered by E. L. Wright et al. 1996, and used in different forms by the Wilkinson Microwave Anisotropy Probe (WMAP; C. L. Bennett et al. 2013), Planck (Planck Collaboration et al. 2020), and the Atacama Cosmology Telescope (ACT; S. Naess et al. 2020)). Recently these ideas have been introduced into optical imaging and spectroscopy (D. Lang 2014; A. K. Saydjari et al. 2023; D. W. Hogg & A. R. Casey 2024).

One challenge or question with forward-modeling approaches is: Do we choose to deconvolve (in any sense) when we combine the raw data, or do we create a map at the native resolution of the raw data? Since a forward model is in play, the parameters of that model can be of the surface brightness field at any resolution (within reason), in either the angular or spectral sense. When the data and model are both good, the inverse problem can deliver a map at higher resolution than the raw data, at the cost of dependency on detailed assumptions. Here we use beliefs about the PAH and recombination lines to effectively perform only a light deconvolution of the SPHEREx data in the spectral direction. That is, we measure the PAH and recombination lines as precisely as possible given what we know about those line positions and shapes, and about the spectral response of each pixel of the SPHEREx hardware.

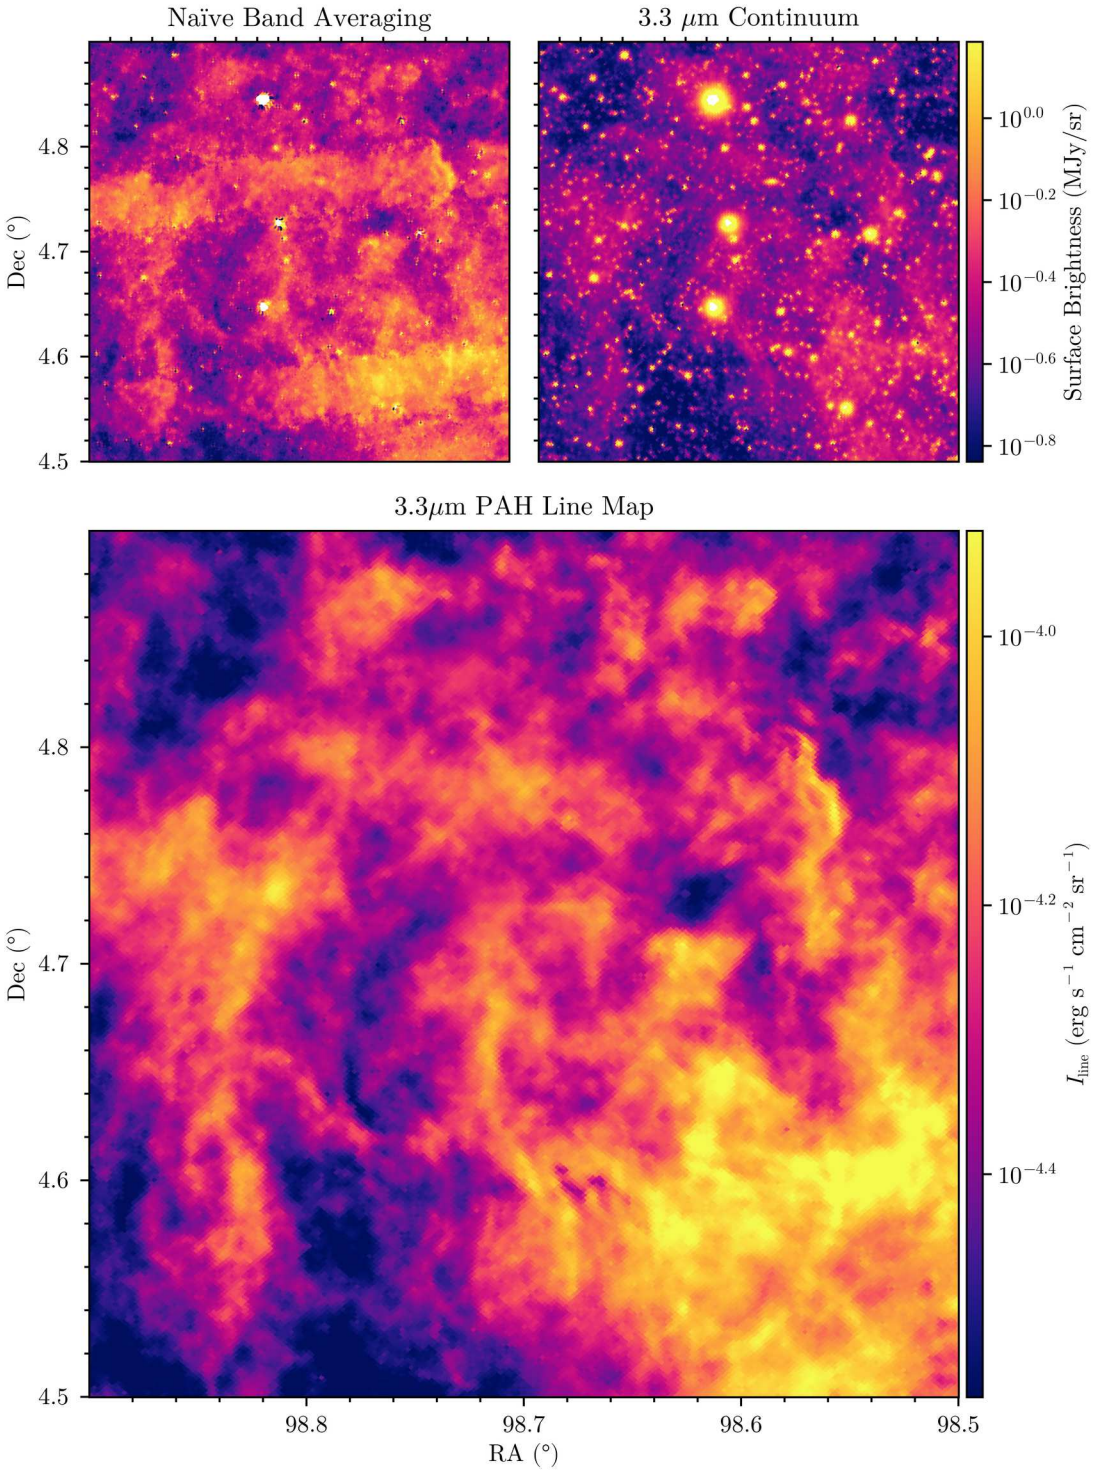

Figure 2. Same as Figure 1, but zoomed in to the middle left.

The rest of this paper is structured as follows. In Section 2 we describe the SPHEREx data products used in this work, and our processing of these data to prepare them for use in our forward model. We describe our methodology in detail in Section 3. We provide a detailed look at maps of select Galactic regions (Section 4: the Rosette nebula, Corona Australis, Cygnus X, and the Heart and Soul nebulae) and external galaxies (Section 5: M31 and M33). We measure the sensitivity of our map and compare to other literature SPHEREx maps in Section 6. We discuss the limitations of our methodology and future improvements in Section 7 and make these maps publicly available in Section 8.

2. DATA MANAGEMENT

This work makes use of the SPHEREx Quick Release 2 (QR2) and Quick Release 3 (QR3) data products. These data are the Level 2 products of the SPHEREx data processing pipeline (R. Akeson et al. 2026), which performs astrometry, flux calibration, and bad pixel flagging. The resulting data are "spectral images", consisting of calibrated surface brightness measurements, variance and zodiacal light (zodi) estimates, and status flags.

The dataset used in this work includes all exposures through 2026 week 36 (2025 April 24 to 2026 September 7). QR2 consists of all exposures taken up to 2026 week 29; QR3 is all data taken starting 2026 week 30 (SPHEREx Science Data Center 2026). The two release versions differ by a major upgrade to the reduction pipeline, which updated the calibrations, applied a new data-driven rescaling to zodiacal light estimations, and enhanced the bad-pixel flagging. The updated calibrations and zodiacal light rescalings are accompanied by calibration products (absolute gain matrix, dark current map, flux correction, zodi scaling) that allow retroactively updating the QR2 data. For a measurement from detector pixel j with surface brightness F, variance V, and zodi Z, the QR3-recalibrated measurements are given by

$$\begin{aligned} F_{\mathrm{QR3}} &= c_j\, F_{\mathrm{QR2}} + A_{\mathrm{QR3},j}\,(D_{\mathrm{QR2},j} - D_{\mathrm{QR3},j}), \\ V_{\mathrm{QR3}} &= c_j^2\, V_{\mathrm{QR2}}, \\ Z_{\mathrm{QR3}} &= \alpha(\lambda_j)\, Z_{\mathrm{QR2}}, \end{aligned} \tag{1}$$

where c is the flux correction (which rescales the gains from QR2 to those of QR3), A is the absolute gain matrix, D is the dark current map, and α is the zodi scaling, linearly interpolated to the central wavelength λ_j of pixel j.

Because of the LVF design of SPHEREx, a spectral image mixes spatial and spectral information. In order to more directly query the data by sky position, we preprocess all QR2 and QR3 data into a pixel-level database called the `talltable`, the details of which are in preparation. To summarize, the pixels of every image are flattened into a set of tuples containing the RA, Dec, surface brightness, variance, flags, zodi, and detector position (detector number, pixel row, pixel column) of each pixel. We call these tuples "pixel measurements." We then compute the NESTED, $N_{\mathrm{side}} = 2^{22}$ HEALPix index (K. M. Górski et al. 2005) of the sky position and sort the pixel measurements from all images according to this index.

There are in total 6.3×10^{12} pixel measurements in our dataset, so the data are split into chunks called "partitions". Partitions are defined by coarser HEALPix indices: $N_{\mathrm{side}} = 2^6$ over most of the sky, with adaptive refinement up to $N_{\mathrm{side}} = 2^9$ in the two deep fields at the ecliptic poles, designed to keep the number of measurements per partition under 200 million. In each partition we store the pixel measurements in Parquet files (Apache Software Foundation 2026), which enable very fast queries and filters. Data that apply only to individual images (such as image metadata) or to individual detector positions (such as the wavelength solution) are stored in separate Parquet files. Each pixel measurement carries an ID for its image and detector position, so these data can be joined with the pixel measurements at query time.

3. METHODS

3.1. *Linear Fit*

At its core, our map-making process LVFit is a least-squares linear fit of a set of basis vectors to all of the SPHEREx observations at a given location on the sky. We make a map by measuring the amplitude of each basis vector on a grid of HEALPix $N_{\mathrm{side}} = 2^{15}$ ($\sim 6\farcs4$) pixels, which have an area of 41.5 arcsec^2. For each pixel, we gather all pixel measurements from the `talltable` (described in Section 2) that lie within $5\farcs5$ great-circle distance of the pixel center (Figure 3). We also select only observations with wavelengths within a model-dependent window around the features of interest, measured in units of each measurement's own bandwidth (see Section 3.4), and we require measurements to have a finite surface brightness, positive variance, and no bit of the SPHEREx quality flag (SPHEREx Science Data Center 2026) set other than `SOURCE` (bit 21), which only marks pixels mapped to a known catalog source. We keep those observations because stars are handled by our own mask (Section 3.5).

We additionally give zero weight to the measurements of the detector pixels in the SPHEREx non-functional pixel list (SPHEREx Science Data Center 2026), in our own list of pixels with anomalous residuals (released in Section 8), and to the measurements rejected by the contamination tests of Section 3.2. We subtract the Level 2 zodiacal light estimate from every measurement (Section 2). We weight the observations as a function of their great-circle distance from the center of the pixel being fit, using a Gaussian with a FWHM of one pixel, $6\farcs44$ ($\sigma = 2\farcs73$), truncated at $2\sigma = 5\farcs5$. Adding the $\sim 6''$ pixel-integrated point spread function (PSF) of SPHEREx (J. J. Bock et al. 2026) in quadrature to this spatial weighting predicts an angular resolution of $\sim 9''$

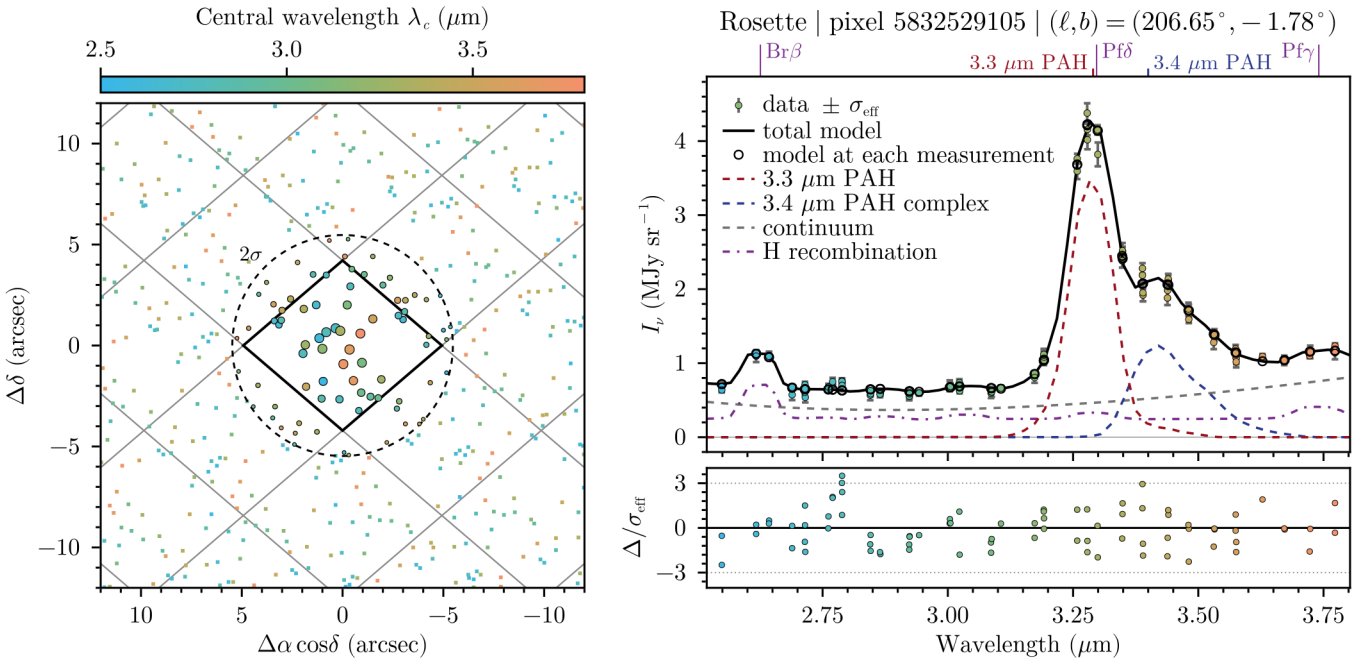

Figure 3. The LVFit model of one map pixel in the Rosette nebula (Index 5832529105, HEALPix $N_{\rm side} = 2^{15}$, NESTED). Left: The pixel of interest (black) and its neighbors (gray), with the 2σ limit circle used to decide on including measurements. Every SPHEREx measurement within the field of view that passes our wavelength cut for the PAH model is shown, colored by its central wavelength. Measurements inside the circle that enter the fit are drawn as circles, with areas that scale with their spatial weight. Measurements not included in the fit are drawn as squares of one size. Right: The same 80 measurements in spectral space with their effective uncertainties. We show the total model (black) and its components as the LVF would see them. Open circles mark the model evaluated with the response of each individual measurement. The positions of the PAH features and of the three strongest hydrogen recombination lines in the window are marked above the panel. Pfδ falls on the 3.3 μm feature. The lower panel shows the normalized residuals. For only two of the 80 measurements do the effective uncertainties include a down-weighting due to the iteratively reweighted least squares procedure of Section 3.1.1.

FWHM. This agrees with the 8.″5 FWHM we measure empirically for isolated stars in the LVFit continuum maps, which we quote as the angular resolution of the maps.

Figure 3 illustrates this selection and the resulting fit for one pixel in the Rosette nebula. This weight is multiplied by the reported inverse variance to form $w_i^{\rm data}$ for each measurement i. The amplitude of each basis vector is then obtained in a first pass simply by weighted least squares, which is refined as described in Section 3.1.1. Our basis vectors are defined as functions of wavelength (on a 2.5 nm grid), but are compared to the data after a transformation given by the spectral response model of the SPHEREx LVF (see Section 3.3). In adopting this procedure, we are making several simplifying assumptions that are not true in reality. We discuss these assumptions, their limitations, and possible solutions in Section 7.

It is important to note that the spatially varying banding we describe in Section 1 that is inherent to naïve binning of the data can also appear in a linear forward model, and will if the model is misspecified. If either the basis vectors or spectral response cause the model to poorly describe the data, then the phase of the wavelength sampling relative to the spectral features of interest will bias the recovered amplitudes in a spatially varying manner. As a concrete example, if the intrinsic width of a basis vector peak is underestimated, then the amplitude you would infer for that peak will depend on whether the wavelength samples you observe straddle either side of the maximum or if a measurement falls at the wavelength center of the peak. This is relevant for the SPHEREx survey strategy because within one survey pass, the dithers and consecutive pointings step a given sky position through the LVF in a regular ladder of wavelengths whose rungs are a fixed 2.7% apart in wavelength, close to one resolution element ($R \approx 35$), because the LVF wavelength is nearly exponential in detector position (J. J. Bock et al. 2026; SPHEREx Science Data Center 2026). Except in the deep ecliptic fields, each later pass lays down a new ladder at a different phase rather than repeating the earlier wavelengths. The phases of those ladders relative to a feature therefore vary smoothly across a field, and any misspecification of the feature shape turns into smooth, spatially varying stripes.

3.1.1. *Iteratively Reweighted Least Squares*

In order to better handle outliers, for example from unmasked bad pixels, transient unmasked cosmic rays, or stray PSF wings, we solve the linear problem above using iteratively reweighted least squares (IRLS; P. J. Green 1984; P. W. Holland & R. E. Welsch 1977). This involves solving the linear problem, updating the weights to down-weight measurements that have large residuals, and repeating until some convergence criterion is met or a maximum number of iterations is reached.

$$w_i^{\rm robust} = \min\left(1,\ \frac{k}{\sqrt{w_i^{\rm data}}\,|r_i|}\right), \qquad (2)$$
$$w_i^{\rm new} = w_i^{\rm robust}\, w_i^{\rm data}$$

Here $w_i^{\rm data}$ is the original data weight for measurement i, in our case the combined spatial and inverse variance weight, r_i are the residuals, and $w_i^{\rm new}$ is the weight to be used in the next round of least squares fitting. When the normalized residual $\sqrt{w_i^{\rm data}}\,|r_i|$ is less than k, the weight is unchanged. When it is greater than k, the measurement is down-weighted in proportion to $k/(\sqrt{w_i^{\rm data}}\,|r_i|)$, so that its contribution to the normal equations grows only linearly with the residual. We choose $k = 3$, which is often called Huber weighting in the literature (P. J. Huber 1964). The weighting optimizes the Huber loss, which transitions from quadratic to linear at 3σ. The Huber loss is convex, avoids modifying the uncertainties within 3σ, and is well-suited for situations such as this one where we expect few outliers as opposed to a heavy-tailed outlier distribution. We repeat this weight update until the largest relative change in the fit linear coefficients is $< 10^{-7}$ or 50 iterations are reached. A pixel is only fit if it has more measurements than basis vectors. Otherwise it is left empty and masked (Section 3.5). For every pixel we store the fit amplitudes, their formal covariance from the inverse of the final (Huber-weighted) normal matrix, the reduced χ^2, and a robust estimate of the scatter of the normalized residuals, which the masking stage uses. The line surface brightnesses we report are the integrals over frequency of the fitted feature basis vectors, in erg s^{-1} cm^{-2} sr^{-1}. The formal uncertainties are calibrated against the data in Section 6.1.

3.2. *Other Contamination Removal*

Some forms of contamination are not easily detected in single-pixel fits. To handle them we defined two custom rejection methods described below for the most problematic cases.

Scattered moonlight—We find that one edge of each of the long-wavelength detectors (4, 5, and 6; 2.4–5.0 μm) needs to be masked for some exposures. After investigating, we find that this edge-outlier occurrence depends strongly on the angular separation between the Moon and the boresight, and on the position angle of the Moon on the sky relative to the detector edge in question. The SPHEREx team has also noted this feature and attributed it to the Moon shining on the exterior of the instrument (J. J. Bock et al. 2026; SPHEREx Science Data Center 2026). We compute the Moon direction for each exposure from the JPL DE440 ephemeris (R. S. Park et al. 2021) through SPICE (C. H. Acton 1996), at the exposure mid-time, from the spacecraft position recorded in the image headers, and the orientation of the detector edge on the sky from the image world coordinate system (WCS). In detector coordinates, this edge runs parallel to the direction of wavelength variation, at the "right" edge of the detector (pixels near the highest x-axis values), so the affected strip spans the full wavelength range of the detector.

In Figure 4, we show the strength and angular dependence of this effect for every exposure of detector 6 in our dataset, measured as the excess of the outermost 8 pixels along this edge over a linear trend fit to the interior of the detector at the same wavelength. The excess appears only when the Moon is 30–60°from the boresight and lies on the side of this edge, in two narrow lobes of position angle, and it recurs every lunar month. We show the ramp profile of this excess as a function of distance from the edge for the exposures in the core of the lobe (Figure 4, bottom left). This shows that the effect is larger, and penetrates further in from the edge, for the longer wavelength detectors. This effect is about 20 times weaker on detector 4 than on detector 6 (0.004 versus 0.087 MJy sr^{-1} in the outermost 8 pixels). For the exposures in the core of the lobe, the excess on detector 6 is still 0.01 MJy sr^{-1} at 128 pixels and beyond about 300 pixels becomes indistinguishable from the interior linear trend fit.

Until we better understand how to model this contamination, we mask the outermost 128 pixels from this edge on detectors 4, 5, and 6 (setting those measurements to have zero weight) for all exposures for which the angular separation to the Moon is less than 60°and the position angle of the Moon on the plane of the sky is within $-20°$ to $+70°$ of the outward normal of this edge (Figure 4, thick dashed box, right). This purely geometric flag encloses the lobe of the excess in Figure 4 at the cost of 5.2% of detector 6 exposures (4.4% of the detector 4 exposures in our dataset). In the fields of this work it removes between 0.0% and 1.3% of the measurements, M31 (29 flagged exposures) being the most affected. Given the orientation, our hypothesis is that this is scattered moonlight that manages to leak in through the $\sim 100\,\mu$m gap between the filter assembly and the detector (P. M. Korngut et al. 2018) for these specific Moon orientations.

Fast moving objects—Single-pixel outlier flagging often struggles to identify faint moving objects because their per-pixel signal falls below detection thresholds even when the integrated trail is significant (P. J. Whidden et al. 2019; G. Nir et al. 2018). They then create coherent biases on the plane of the sky that are especially nefarious for studies of filamentary structure or

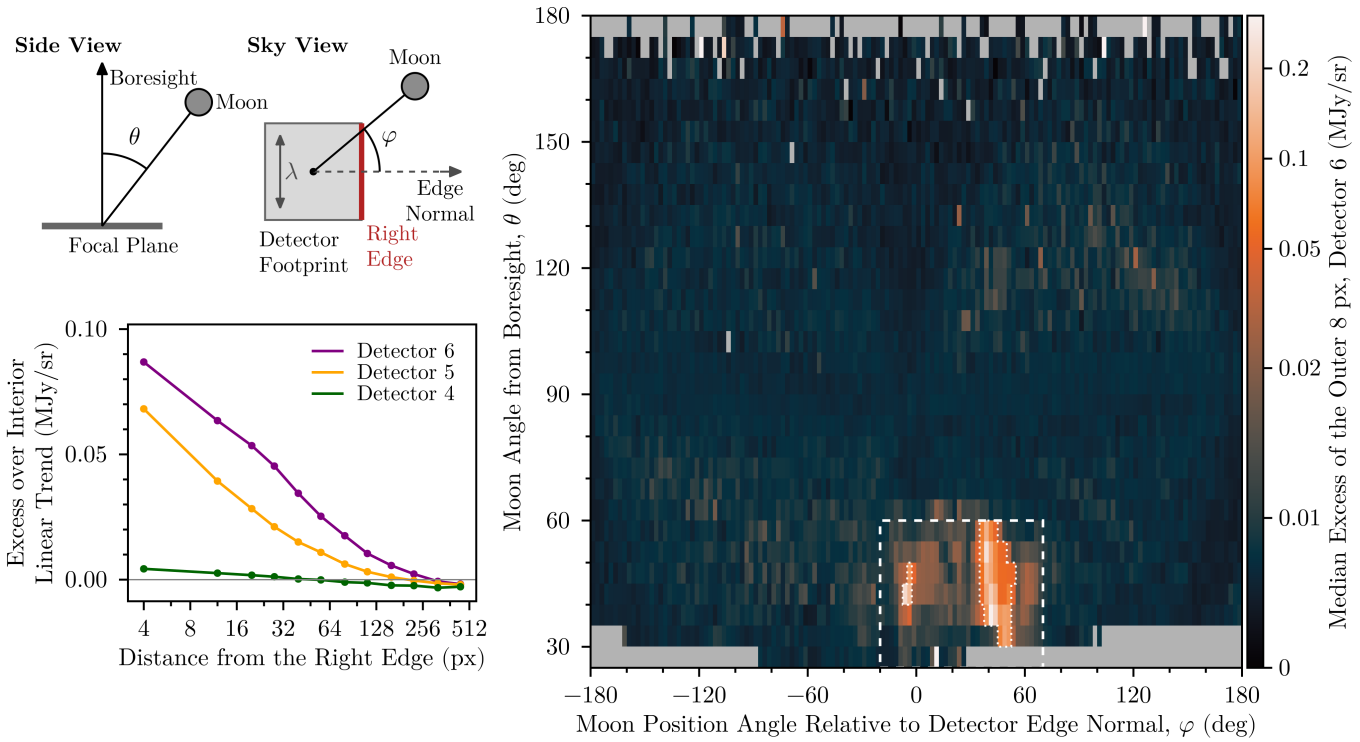

Figure 4. Identifying the effect of scattered moonlight on the right edge of the long-wavelength detectors. Left, top: Definition of the angle θ of the Moon from the boresight and the Moon's position angle φ on the sky relative to the outward normal of the right detector edge. Left, bottom: Excess of a block of pixels at a given distance from the right edge over a linear trend fit to the interior pixels of the same wavelength (more than 128 pixels from either edge), for detectors 4, 5, and 6, as the median over the exposures in the core of the Moon lobe (dotted outlines: bins with a median excess above 0.05 MJy sr^{-1} and at least 20 exposures) minus the median over 4,000 random exposures outside the geometric flag. Each point sits at the center of its block, and the blocks widen with distance from the edge, from 8 to 128 pixels wide. Right: The excess for the outermost 8 pixels of detector 6, for all 248,000 detector 6 exposures, binned by φ (2.°5 bins) and θ (5° bins) (gray indicates bins with < 5 exposures). We use a logarithmic color stretch above 0.02 MJy sr^{-1}. The dashed rectangle is the geometric flag used in this work ($\theta < 60°$, φ from $-20°$ to $+70°$).

image anisotropy. To mask pixels impacted by these faint moving objects, we do a first pass of the IRLS, and then run a Radon transform (S. R. Deans 1983) on the per-exposure residual maps to identify linear streaks of outlying pixels. For each exposure and each HEALPix $N_{\rm side} = 2^7$ ($\sim 0.°5$) pixel of the fit, the normalized residuals $r_i\sqrt{w_i^{\rm data}}$ of that exposure are placed at their detector coordinates, median-subtracted, clipped at ± 5, and summed along lines of width 1–3 pixels and length 100, 200, 400 pixels, or the full footprint, at angles spaced by 0.°25. Each sum, divided by the square root of the number of pixels it contains, is normalized by the median absolute deviation over all lines of the same width and length in that exposure, which gives a Radon peak score S in robust standard deviations.

Because the sky tiles cut long tracks into fragments, candidate fragments of the same detector line in the same exposure (within 1° in angle and 5 pixels in offset) are merged, with their scores summed and divided by the square root of their number. We impose a threshold of $S > 30$ on the merged Radon peak, and additionally require a total extent of at least 200 pixels along the line, a distance of at least 128 pixels from the detector edge (which separates tracks from the edge ramps above), and that the same sky pixels in the other exposures of the tile have a mean residual below 30% of the residual on the track (which separates tracks from static features such as stellar wings or nebular edges). These thresholds were set using the two tracks found by eye. We then mask ± 3 pixels around the line over the full extent of the detector, since the detected extent is only a lower bound on the length of the track (which may continue through tiles with too few measurements to be tested) and the strip costs only the equivalent of six columns of a single exposure. We then refit any HEALPix pixels that used any of the measurements with weights now set to zero by the linear feature mask. Across the Rosette, Corona Australis, Cygnus X, Heart and Soul, M31, and M33 maps, we found 53 linear tracks in 49 of the 1,784 exposures and masked an additional 0.02% of the measurements for this reason (0.001–0.06% per field). In the $1° \times 1°$

depth-test region at the north ecliptic pole (NEP), we found 237 tracks in 230 of the 6,183 exposures.

3.3. *Spectral Response of the LVF*

To forward model theoretical spectra to the SPHEREx measurements, we need to know the spectral response. Because SPHEREx uses a linear variable filter, this is not just a matter of looking up a static filter curve.[9] We instead need a spectral response model that varies across the detector. The SPHEREx team divides the instrument's wavelength coverage into 102 channels (17 per detector), and the ground spectral calibration of H. Hui et al. (2026) further divides each detector into 68 "sub-channels" of pixels that are co-spatial with respect to the variation in center wavelength of the linear variable filter, each with a measured spectral response.

Instead of making this (sub-)channel approximation, we model the linear variable filter as continuous, reflecting its physical fabrication. We use the highest spatial resolution measurement of these sub-channels, the 68 measured sub-channel response profiles per detector of the ground spectral calibration (H. Hui et al. 2026), for the shape of the response only, and continuously interpolate the mapping of pixels to sub-channels, introducing decimal sub-channels. This allows us to define the spectral response of a given pixel on the detector by interpolating the measured (non-Gaussian) sub-channel spectral responses. We blend the two bracketing sub-channel profiles linearly, and the blend is shifted in wavelength so its half-transmission wavelength equals the center wavelength (band center) of that pixel in the public SPHEREx spectral WCS (the per-pixel `spectral_wcs` calibration product of QR3 (SPHEREx Science Data Center 2026), in the pixel convention of the Level 2 images).

We test the efficacy of the interpolation by leave-one-out tests, predicting each measured sub-channel profile from its two neighbors, which is twice the separation of a real interpolation, so it is an upper bound on the interpolation error. We find average errors of 1.8% of the peak transmission in the interpolated spectral response functions of detector 4, though only 0.5% in the response integrated over the bandpass for a flat spectrum. The maximal errors are 12% of the peak and 4% in the bandpass-integrated response, which occur for the spectral response at the short wavelength edge of detector 4 (2.45 μm), where the dichroic spectral response causes rapid variations in the effective spectral response. We distribute a package in Section 8 that implements this continuous spectral response model and can return the spectral response function of an arbitrary pixel in the focal plane in ~ 1 μs having only loaded into RAM 49 MB of precomputed coefficients per detector (287 MB for all six).[10]

When forward modeling narrow, well-characterized lines, we use this spectral response model to predict how these components will contribute to SPHEREx's measurements, for example for hydrogen recombination lines or molecular hydrogen ro-vibrational lines. In this case, the lines are well-known and the biggest source of modeling uncertainty is in the spectral response model. For the PAH emission, the shapes of the underlying spectral lines are not fundamentally well-determined and can vary. As a result, we learned data-driven corrections to basis vectors for our 3.3 and 3.4 μm PAH components. In this case, we are not sensitive to the LVF in terms of the quality of our modeling of the data. The LVF only impacts our ability to do the under-determined task of deconvolving the data-driven basis vectors to infer a higher-resolution pre-filter-curve-convolution PAH spectrum, which we defer to future work.

Because the spectral response function model and wavelength solution are intimately coupled, our modifications to the spectral response model above could introduce wavelength offsets compared to the SPHEREx wavelength solution when forward modeling theoretical spectra. We do find some small wavelength offsets are necessary for detectors 3 and 4. The hydrogen recombination lines in the Rosette nebula are fit best when the response of each pixel sits 4 nm redward of its center wavelength on detector 4 and 2 nm redward in detector 3. This is about 4% of the bandwidth in both bands. We apply these offsets to the narrow-line basis vectors by shifting their line wavelengths by the same amount in the opposite direction. We do not shift the PAH basis vectors, since we see that their data-driven corrections absorb the same offset. We have not attempted to determine whether these offsets reflect issues with the SPHEREx wavelength solution, the response profiles, or our interpolation of them.

3.4. *Models*

Our main spectral model focuses on making PAH maps. For this, we consider all observations within ± 10 bandwidths around both the 3.3 and 3.4 μm PAH peak centers, all of which fall within detector 4. To model the continuum, we use three basis vectors: a constant, λ^2, and λ^{-2}, each normalized to unity at 3.3 μm. These are motivated by our expectation of being in the Rayleigh-Jeans tail of starlight, which falls like λ^{-2}, and the presence of continuum emission from very small grains with a color temperature near 1000 K (K. Sellgren 1984), which λ^2 approximates over detector 4. The three terms together describe the smooth continuum across our 2.5 to

[9] Though, in classical photometry it has been recognized for a long time that the apparent spectral response can vary across the field of view due to angle-dependence in the transmission function of the filter (P. H. Lissberger & W. L. Wilcock 1959; M. Doi et al. 2010; J. L. Tonry et al. 2012; T. S. Li et al. 2016).

[10] This timing depends on the hardware and load. It is reported for a single thread on an otherwise idle Intel Xeon Platinum 8168 (2.7 GHz) node of the Flatiron Institute Popeye cluster, averaged over 2×10^6 random detector 4 pixels. A pixel whose maps are already in cache returns in 0.7 μs.

3.9 μm window with curvature of either sign. To handle regions where ionized gas is bright, we add one basis vector for hydrogen recombination emission. It contains all Brackett, Pfund, and Humphreys lines between 2.3 and 4.5 μm together with the nebular continuum, with relative strengths fixed to Case B at 10^4 K and $n_e = 10^2$ cm^{-3} from PyNeb (V. Luridiana et al. 2015; P. J. Storey & D. G. Hummer 1995). As a result, the hydrogen recombination basis vector has a single amplitude, which we report in terms of the corresponding Pfδ line surface brightness. We represent each line with a Gaussian of width 1.5 nm, far narrower than the spectral response, and we shift this basis vector by -4 nm to compensate for a detector 4 wavelength offset we find necessary to optimize the goodness of fit for these narrow lines (see Section 3.3).

3.4.1. *Data-Driven Corrected PAH Basis Vectors*

We found that simple analytic forms for the PAH emission were insufficient to describe the data well and left significant banding in our maps. To solve this problem, we initialized our PAH basis vectors using higher resolution spectra of PAH emission from JWST NIRSpec, obtained by the PDRs4All program toward the Orion Bar (O. Berné et al. 2022; E. Peeters et al. 2024). We use the published template spectrum of the atomic PDR region from the PDRs4All website.[11] We then separated the 3.3 and 3.4 μm emission into two basis vectors. After masking the emission lines listed by E. Peeters et al. (2024) and subtracting a linear continuum, we fit their Gaussian decomposition to the spectrum, which includes Gaussians at 3.246, 3.29, 3.33, 3.40, 3.42, 3.46, 3.52, and 3.56 μm plus a broad 3.42 μm plateau. We use these Gaussians only to apportion the emission. At each wavelength, the observed spectrum is divided between the two basis vectors in proportion to the Gaussians of the 3.3 μm group (3.246, 3.29, 3.33 μm) and the 3.4 μm complex (3.40, 3.42, 3.46, 3.52, 3.56 μm), so the basis vectors keep the observed line shapes rather than a combination of Gaussians approximation. We keep no separate plateau component and its emission is shared between the two templates in the same proportion.

Each template is normalized by its integral in frequency, so the fitted amplitude is the integrated surface brightness of the whole feature. These two spectra are our initial set of PAH basis vectors. Because of some combination of issues with the spectral response, wavelength solution, or difference in the average shape of PAH emission compared to that PDR, we still observe some banding in our PAH maps when using these basis vectors. As a result, we learn data-driven corrections to these basis vectors on five training fields: 2° boxes toward ρ Ophiuchi, the Rosette nebula, Chamaeleon I, NGC 7023, and IC 348, chosen to span ecliptic latitude, scan angle, and surface brightness. Five more fields (Lupus I, NGC 2068, TMC-1, the Polaris flare, and the Ursa Major cirrus) were held out, and we used the amplitude of banding in those maps to set hyperparameters and refine the procedure presented here. We make the resulting basis vectors public in Section 8.

We learn the corrections in the space of the measurements rather than on a wavelength grid. For PAH basis vector m (the 3.3 or the 3.4 μm basis vector) and a measurement whose spectral response is centered at wavelength λ, the model value is

$$B_m(\lambda) = \mathcal{F}\left[T_m^{\rm init}\right](\lambda) + \sum_j C_{jm}\,\beta_j(\lambda), \tag{3}$$

where $\mathcal{F}[T_m^{\rm init}]$ is the initial basis vector convolved with the response of that measurement, β_j is the jth cubic B-spline in λ on a grid of knots 10 nm apart, and C_{jm} is the coefficient of spline j for basis vector m. The sum runs over the splines that lie entirely within the range in which we allow corrections, 3.04–3.50 μm for the 3.3 μm basis vector and 3.24–3.72 μm for the 3.4 μm basis vector, 43 and 45 splines respectively, each spanning 40 nm. The model for measurement i of map pixel p is then $\sum_m a_{pm} B_m(\lambda_i)$ plus the continuum, where λ_i is the center wavelength of that measurement's response. The correction is thus a smooth function of wavelength with 88 free coefficients in total. The coefficients C are the same for every map pixel and every field with only the amplitudes a_{pm} varying from pixel to pixel.

We find C by minimizing the Huber loss, with two added regularization terms,

$$\sum_p \sum_{i\in p} w_i^{\rm new}\, r_i^2 + s_1\,\|C\|^2 + s_2\,\|D^2 C\|^2 \tag{4}$$

while optimizing both the per-pixel amplitudes and C. Here p runs over the training map pixels and i runs over the measurements of pixel p. The residual r_i of measurement i is relative to the model of Section 3.1 with its PAH columns defined by Equation 3, and $w_i^{\rm new}$ is its weight after the IRLS of Equation 2 has converged, so the first term is the same Huber-weighted loss that the production map-making code minimizes. To this, we also add a ridge term that encourages smaller corrections and a penalty that encourages smoothness along the knot sequence,

$$\begin{aligned} \|C\|^2 &= \sum_m \sum_j C_{jm}^2, \\ \|D^2 C\|^2 &= \sum_m \sum_j \left(C_{j+1,m} - 2C_{jm} + C_{j-1,m}\right)^2, \end{aligned} \tag{5}$$

where the second is the squared second difference of the spline coefficients. The hyperparameters chosen for the amplitude of these terms are $s_1 = 1 \times 10^{-3}$ and $s_2 = 1 \times 10^{-1}$, respectively. Each correction is constrained to integrate to zero, $\sum_j C_{jm} \int \beta_j\, d\nu = 0$ for each basis vector m, which we impose via a Lagrange multiplier in solving for C, so the normalization of the initial basis vectors is preserved. The penalty strengths

[11] https://pdrs4all.org

and the wavelength ranges of the corrections were chosen by scoring the maps of the five held-out fields with a model-independent banding metric, discussed below, breaking ties using the best χ^2.

We alternate between the per-pixel IRLS fits with C fixed and a weighted least-squares update of C with the amplitudes and the weights $w_i^{\rm new}$ fixed, on every third unmasked pixel of the training fields (2.9 million pixels), until the largest change in C is below 10^{-3}. The robust χ^2 per degree of freedom of the training pixels falls from 1.40 with the initial basis vectors to 1.35 with the corrections.

Because the corrections of Equation 3 are defined only through the LVF, we do a final solve for smooth, nonnegative basis vectors on the 2.5 nm wavelength grid shared by the rest of our basis vectors that reproduce the B_m. This is a quadratic program that matches each B_m through the response of every training measurement, keeps the integral of the initial basis vectors, and penalizes curvature and distance from them. We take the smoothest solution within 5×10^{-4} of the best χ^2 per degree of freedom. While we infer smooth, high-resolution PAH basis vectors, we caution against interpreting their high resolution structure. We only show and use the LVF-convolved PAH basis vectors. Understanding how to deconvolve the data-driven corrections and structure we find goes beyond the scope of this work. Conceptually, deconvolution always requires a regularization choice, though one possible route would be to regularize based on a sparse sampling of higher resolution spectra.

Because Equation 4 is the data likelihood of the training fields, our data-driven corrections improve the χ^2 of our fit by construction. As an independent check, we use a banding metric that we never optimize directly. It asks how much of the small-scale structure in a map can be predicted from the wavelength sampling alone. To measure it, we compute for every map pixel seven descriptors of how its measurements sample the spectrum, which serve as regressors: (1) the fraction of the measurements whose center wavelength lies within about one bandwidth of each of three reference wavelengths (3.29, 3.40, and 3.46 μm), (2) the mean offset of those measurements' center wavelengths from the 3.29 and 3.40 μm peaks, in units of the bandwidth, and (3) the signed offset of the nearest measurement from the same two peaks. We bin our PAH maps and these descriptors on a $14\farcs4$ grid, remove structure larger than $6'$ from both using a high-pass filter, and divide the map by its smooth version, so that the result is a fractional fluctuation. We then try to predict the map of interest from the regressors. We take the root-mean-square of the model fit using these regressors as our measure of the banding strength, the fraction of the local surface brightness that the sampling pattern alone explains. Because real structure can correlate with the descriptors by chance, we repeat the regression with the descriptor maps reflected, rotated, and shifted, and subtract the median of these null results in quadrature, so a map with no banding scores zero within the noise of the metric. We tracked these scores on both the summed $3.3+3.4\,\mu$m surface brightness and the $3.4/3.3\,\mu$m ratio maps.

For the Rosette nebula, adding the data-driven corrections to our JWST PDRs4All initialization halves the banding in the $3.4/3.3\,\mu$m PAH map, from 4.7% to 2.2% of the local surface brightness. The banding of the summed $3.3+3.4\,\mu$m map in Rosette is $\sim 2\%$ and only marginally detected with either set of basis vectors. Relative to the initial basis vectors, those with data-driven corrections reduce the banding in every field and in both maps. Across the five training and three bright held-out fields, the banding of the summed $3.3+3.4\,\mu$m surface brightness drops from 0.9% to 0.6% of the local surface brightness. At that level it is only marginally detectable against the null tests (1.5σ). The banding of the $3.4/3.3\,\mu$m ratio in the same fields halves, from 1.8% to 0.9%. It remains detectable (8σ), driven by the Rosette and TMC-1 fields, and the three bright held-out fields alone show the same halving but remain 5σ detectable. On our two faintest cirrus test fields, the Polaris flare and the Ursa Major cirrus, the data-driven corrections reduce the banding in the ratio maps from 3.0% to 2.0%. This is still well above the null tests (6σ). These fields are the hardest case for a ratio map. Both features are faint, so the per-pixel ratio divides noise by noise, and sampling-dependent biases that are negligible against bright emission become a measurable fraction of the signal. Improving our model sufficiently to bring banding below the detection threshold in all fields will be the subject of future work, and may require higher dimensional representations of the PAH features and learning the basis vectors jointly with the map (at inference time) to capture spectral shape variation that correlates with environmental changes. For now, visual comparisons of our maps with others in the literature (see Section 6.3) demonstrate the substantial improvement we have achieved in reducing banding.

3.4.2. *Fit Quality*

In Figure 3, we show the fit of this model to an example pixel in the Rosette nebula. In the left panel, we demonstrate which measurements are included, the spatial weighting scheme, and how the central wavelength of the measurements varies spatially. In the right panel, we show how the basis vectors combine to give the total model and compare the model to the data, with the uncertainty normalized residuals shown in the bottom subpanel. Here, the uncertainties we show include the final multiplicative updates to the weights from the IRLS procedure. We provide markers showing the locations of the center wavelength of the main PAH features, as well as the three strongest hydrogen recombination lines in the wavelength range (Brβ, Pfγ, and Pfδ). The Pfδ feature directly overlaps the PAH features.

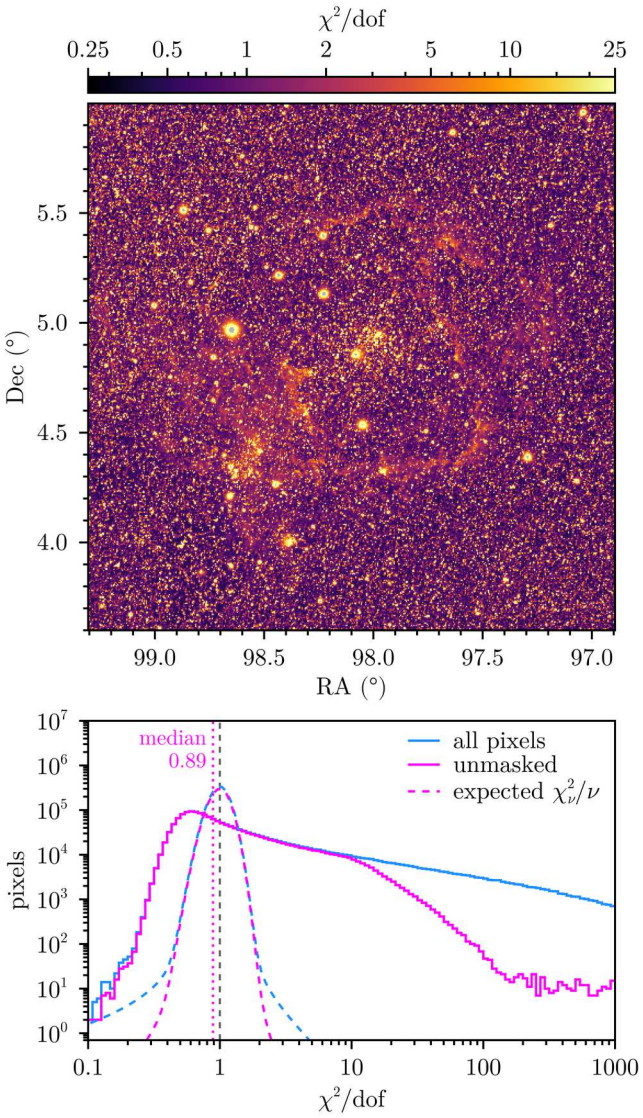

Figure 5. Goodness of fit of our linear forward PAH model for the Rosette nebula. Top: Map of χ^2/dof of each pixel, using the pipeline uncertainties without a model-error floor. Here the dof is the number of measurements contributing to a pixel minus the six model parameters, which has a median of 75, and χ^2 is simply summed over all of the measurements contributing to the pixel. Pixels masked for any reason (Section 3.5) are gray. Bottom: Histogram of χ^2/dof for all pixels (blue) and for unmasked pixels (magenta). The dashed curves show the histograms expected if the χ^2 of every pixel followed a χ^2 distribution given its own number of dof. The dotted line marks the unmasked median of 0.89 and the dashed line marks χ^2/dof = 1.

The right panel of Figure 3 also shows the average wavelength spacing at the current point in the mission. Because the phase of the wavelength sampling is fairly random away from the ecliptic poles, our ability to learn and constrain spectral substructure should improve with increasing duration of the SPHEREx mission, because it will bring an increase in the spectral sampling density. From this perspective, the deep but uniform wavelength sampling of the ecliptic poles is a missed opportunity for data-driven refinements of the spectral response model of SPHEREx and pushing the limits of the spectral resolution of the components we can infer from the data.

To demonstrate how consistent our model is with the data on a larger scale, we show the χ^2 distribution for the Rosette nebula testbed in Figure 5. After excluding masked pixels from our fit (e.g. stars, see Section 3.5), we find a median χ^2/dof = 0.89. This suggests that our model is describing the data well on average. The fact that our typical χ^2/dof is less than one might reflect that we added in too much freedom to our model, for example in terms of the number of continuum components. However, looking at the overall distribution of the χ^2/dof, we see that it is extremely heavy-tailed compared to the statistically expected distributions. This heavy-tailing is especially prominent in the direction of pixels with elevated χ^2/dof. We take this bottom panel as a strong indication that we are not fully describing the complexities of the SPHEREx observations, suggesting there is an opportunity to extract much more information from the data. While it can be disheartening to see that your model is misspecified, having access to the residuals is one of the key benefits to a forward model, and is not accessible to even advanced channel mapping methods (J. L. Hora et al. 2026).

In the top panel of Figure 5, we show the plane-of-sky map of χ^2/dof. The relative paucity of spatial structure in this map acts as further validation of our model. In terms of spatially correlated structure, we do see χ^2/dof ~ 3 along the ridges of brightest PAH emission. This could indicate that the PAH lineshape varies in regions of extremely high UV flux, which we speculate may be because these higher UV fluxes stimulate vibrational hot bands (J. R. Barker et al. 1987) or modify the underlying carrier distributions.

In addition to the model described above, we also consider two other models used in this paper. For some of the test fields where we do not expect significant emission in hydrogen recombination lines, we simply use the model described above, but without the hydrogen recombination basis vector. We also consider a pure hydrogen recombination line model that we apply only to the strongest lines in the SPHEREx wavelength range. We query ± 4 bandwidths around Paα (1.876 μm), which is a subset of detector 3 that also contains the Brδ (1.945 μm) and Brϵ (1.818 μm) lines. We model the three lines as a single basis vector with Case B line ratios (D. G. Hummer & P. J. Storey 1987; V. Luridiana et al. 2015) and report the fit amplitude in terms of the Paα line surface brightness. As above, we represent each line with a Gaussian of width 1.5 nm, shifted here by -2 nm to account for our apparent detector 3 wavelength offset (see Section 3.3). For the continuum, we use the same

three basis vectors as above, a constant, λ^2, and λ^{-2} term, here normalized to unity at 1.876 μm.

3.5. *Masking*

Not every pixel fit is a measurement of diffuse emission. After fitting every HEALPix $N_{\rm side} = 2^{15}$ pixel, we mask those pixels where the fit likely captures something else, such as a star or a poorly modeled spectrum. The pixel mask is applied after the fit and is distinct from the measurement-level rejections of Section 3.2 (Moon-contaminated detector edges and moving-object tracks, which enter the fit with zero weight) or IRLS down-weighting of outliers. The mask involves six different cuts applied in the order listed. The cuts were designed around simple principles and evolved over time in response to by-eye inspection of our test fields. These cuts are admittedly somewhat contrived and this is one of our primary motivations for modeling both the diffuse emission and point sources at the same time, as discussed in Section 7.

Coverage—The fit itself requires more measurements than basis coefficients for the linear system to not be singular. For models including PAH components, we also mask pixels with three or fewer measurements with a central wavelength within one bandwidth of 3.29 μm. This ensures we have sufficient sampling of our spectral feature of interest. For the Paα fit, we apply the same rule about 1.876 μm. This cut is mild and removes less than 0.6% of every field.

Star—We seed the star list with unWISE sources brighter than $W1 = 9.5$ mag (E. F. Schlafly et al. 2019). Where unWISE is saturated ($W1 < 8$ or flagged), we use the 2MASS K_s magnitude of the counterpart within 3″ instead (M. F. Skrutskie et al. 2006), and we add 2MASS stars brighter than $K_s = 8$ mag that unWISE lacks. To predict how far each star contaminates the line map, we measure its continuum halo and the fraction of that halo that leaks into the line map. We model the halo with a radial template $h(r) \propto [1 + (r/r_c)^2]^{-\beta/2}$, which we fit to the stacked halos of the field's bright stars. For fields with fewer than 100 bright stars with measurable halos, we use the median template of the other fields instead. For each star we then measure the amplitude A of this template on the fit 3.29 μm continuum map. The leakage fraction F comes from a regression of the stacked line-map profiles on the stacked continuum profiles, taking the larger of the 3.3 and 3.4 μm values since we use one mask for both maps.

In nebular regions, part of the continuum excess around a star is ISM emission rather than halo. Where the $I_{3.3}$/continuum color of the halo region is closer to the ISM than to a star, we split the excess into a stellar part with color F and an ISM part with the local ISM color, and use only the stellar part when fitting A. The star is masked out to the radius where its predicted contamination of the line map, $A\,F\,h(r)$, falls below $0.2\max(\sigma_I, 0.2\,I_{\rm loc})$, with σ_I and $I_{\rm loc}$ the scatter and median of $I_{3.3}$ in an annulus around the star, respectively. The annulus we use scales with the brightness of the star, spanning 3–4 R_s, where $R_s = 10^{2.33-0.082\,m}$ arcsec and m is the seed magnitude, with a floor of 120″–200″ that applies to all stars fainter than 9 mag.

For stars that leave a coverage hole of radius $r_{\rm hole}$, for example due to saturation, or are brighter than 6.5 mag, we then grow the mask outward in 15″ rings. Here $r_{\rm hole}$ is the radius of the star's coverage hole or, where larger, of a connected patch of at least 30 cut pixels around it. At each step we compare the median $I_{3.3}$ of the 15″ ring just outside the current disk with that of the 60″-wide reference annulus beyond it, using only unmasked pixels attributed to this star. The growth continues while the ring deviates from the reference annulus by more than twice the standard error of the difference, estimated from the robust scatter in the reference annulus, and by more than half that scatter. The second condition prevents small but statistically significant gradients from growing the disks of stars in bright regions. We grow on deviations of either sign, since the leaked continuum can be oversubtracted or undersubtracted. If a ring is brighter than the reference annulus and the local ISM color can be measured, we also confirm that the ring has a continuum excess and that the ratio of the $I_{3.3}$ excess to the continuum excess is below half the local ISM ratio. This check stops the growth when the emission is dominated by real diffuse emission, which has the local ISM color. We cap the radius of each star mask at the larger of 60″ and $3r_{\rm hole}$. For stars brighter than 6$^{\rm th}$ mag, we make the cap magnitude dependent and it instead grows with brightness as $60'' \, 10^{0.2(6-m)}$, which doubles every 1.5 mag (380″ at 2$^{\rm nd}$ mag). For the few stars brighter than 2$^{\rm nd}$ mag, we relax the cap to 0.°5.

Residuals—For each pixel the fit stores the robust scatter (interquartile range divided by 1.349) of the normalized residuals $r_i/\sqrt{\sigma_i^2 + (f_{\rm m}\,m_i)^2}$ over the measurements shortward of 3.36 μm, where m_i is the model prediction and $f_{\rm m}$ a per-measurement model-error floor. We mask pixels where this scatter exceeds 4. This restriction to only use the shortward side is an attempt to reduce sensitivity to detailed modeling of the 3.4 μm feature shape. The floor is calibrated per field on star-free pixels as the value at which the residual scatter becomes independent of the model signal-to-noise. It ranges from 0.05 (Cygnus X) to 0.18 (Heart and Soul). The uncertainty on each value is 0.003–0.018, based on a spatial jackknife that compares different regions of the image. In two fields, Corona Australis and the NEP, the calibration does not bound $f_{\rm m}$ below 0.2. For these we use the median of the other fields, 0.12.

Fainter stars not in the seed list are caught by a second residual cut. The floor $f_{\rm m}$ can hide a poorly fit point source. Near a star the model m_i includes the star's continuum, so the floor term $f_{\rm m} m_i$ grows with the star's brightness. A misfit of order the star's signal can then

leave the robust scatter below 4 with the floor while far exceeding 4 without it. As a result, we also compute the residual scatter without the floor and mask pixels where it exceeds 4, but only for pixels within $9''$ of an unWISE $W1 < 16$ source that is bright enough to matter, meaning its expected continuum peak from the field's $W1$-to-continuum scaling exceeds four times the floor term $f_{\rm m}|m|$.

Sources whose core has a color consistent with diffuse emission are exempt from this second residual cut. To evaluate the color, we compare the median $I_{3.3}$ and continuum of the pixels within $6\farcs5$ of the source with those in a $12''$–$24''$ annulus. The source is exempt when its core continuum excess is significant, at three times the annulus continuum scatter, and the ratio of the $I_{3.3}$ excess to the continuum excess exceeds half the field's median $I_{3.3}$/continuum ratio. This check works because a star adds continuum but no $3.3\,\mu$m PAH emission, so its PAH/continuum ratio is near zero, in contrast to a compact H II region or PAH knots. This exception protects, for example, the H II regions of M33.

Finally, because the $3.4\,\mu$m map is more sensitive to stellar continuum leakage than the $3.3\,\mu$m map, the disks of stars brighter than $W1 = 6$ mag are grown further using the $3.4\,\mu$m map. At each step we compare the $15''$ ring just outside the current disk with the $60''$-wide reference annulus beyond it, using only unmasked pixels. In both we measure $q = I_{3.4} - \rho I_{3.3}$, where ρ is the ratio of the median $I_{3.4}$ to the median $I_{3.3}$ in the reference annulus. The idea is that q vanishes for emission with the local ISM color, so it isolates stellar excesses. We define Δq as the median q of the ring minus the median q of the reference annulus, and estimate its standard error from the robust scatter of q in the reference annulus. The disk is grown by $15''$ while $|\Delta q|$ exceeds both twice its standard error and 5% of the median $I_{3.4}$ of the reference annulus. The second condition is an attempt to stop the growth on bright backgrounds. There, even small color gradients are statistically significant and would otherwise keep the disk growing. We cap this growth at $\min(4R_0, R_0 + 300'')$, where R_0 is the disk radius before this step. As a final step, we mask any unmasked pixel with at least three of its four neighbors masked, and repeat until no pixel changes. This closes the isolated holes that the residual cuts leave at the cores of compact sources.

The masked fraction of each displayed field is 2.3% (M33), 5.4% (M31), 6.9% (NEP), 11.0% (Rosette), 15.0% (Heart and Soul), 15.3% (Corona Australis) and 34.5% (Cygnus X). In the Galactic plane fields, the two residual cuts dominate. For the Paα map of the Rosette nebula, we apply the mask of the PAH fit in addition to its own coverage cut. In the released maps, the mask is a bit field, so a user can tell which cut masked each pixel.

3.6. *Infilling*

To create clean maps with statistically consistent values in the masked pixels described above, we infill them with a data-driven, anisotropic, conditional Gaussian process method that is a slightly modified version of the method presented in A. K. Saydjari & D. P. Finkbeiner (2022). Briefly, we use unmasked pixels as "training data" of what the structure in an $N_p \times N_p$ pixel patch looks like ($N_p = 31$ pixels, $3\farcm3$) over a surrounding region of widx × widx pixels (widx = 81 pixels, $8\farcm7$). The 5,101 shifted copies of the patch within that region serve as samples from which we build the $N_p^2 \times N_p^2$ pixel-pixel covariance matrix, which does not impose any isotropy or functional form, but instead simply measures the correlations of pixels for each set of 2-dimensional pixel separations.

Then, for the $N_p \times N_p$ patch centered on a masked pixel, we predict the values of all masked pixels in the patch at once, conditioned on the unmasked pixels of the patch, using this covariance matrix as the prior. We store the conditional mean, which is what the maps in this work show, two independent draws from the conditional distribution, and a per-pixel standard deviation, which we report as the uncertainty on infilled pixels in place of the measurement error. Because pixels infilled by earlier patches are treated as known when conditioning later ones (the mean on their means and each draw on its own earlier draw, so that the draws remain a consistent sequential simulation), the conditional variance of a single patch underestimates the uncertainty, by up to 10% in our most heavily masked field. We therefore add the variance propagated from the earlier-infilled pixels in the conditioning set, approximating their posterior covariance by their stored variances and the sample correlation of the patch. Checked against 32 conditional draws, the resulting standard deviation is calibrated to a few percent. Regardless, the uncertainty we report for infilled pixels is a prediction uncertainty rather than a noise estimate so the mask column of the released maps identifies which pixels are infilled and carry this kind of uncertainty. The choice of N_p is driven by computational cost and the choice of widx is driven by the spatial scale at which the structure in the image is changing, the scale at which approximating the samples as stationary breaks down, and the need to have enough samples to stably estimate the $N_p^2 \times N_p^2$ covariance matrix (here 5.3 samples per matrix dimension).

Because we are infilling one patch at a time, and pixels infilled earlier are treated as known when conditioning later patches, the order of choosing these patches introduces some determinism to the map, especially when the masking fraction is large or there are large contiguously masked regions. We infill the masked pixels within the displayed region of each field in independent tiles of $0\fdg75$, each with $0\fdg2$ padding and a fixed random seed, so that the result is reproducible and tiles can be run in

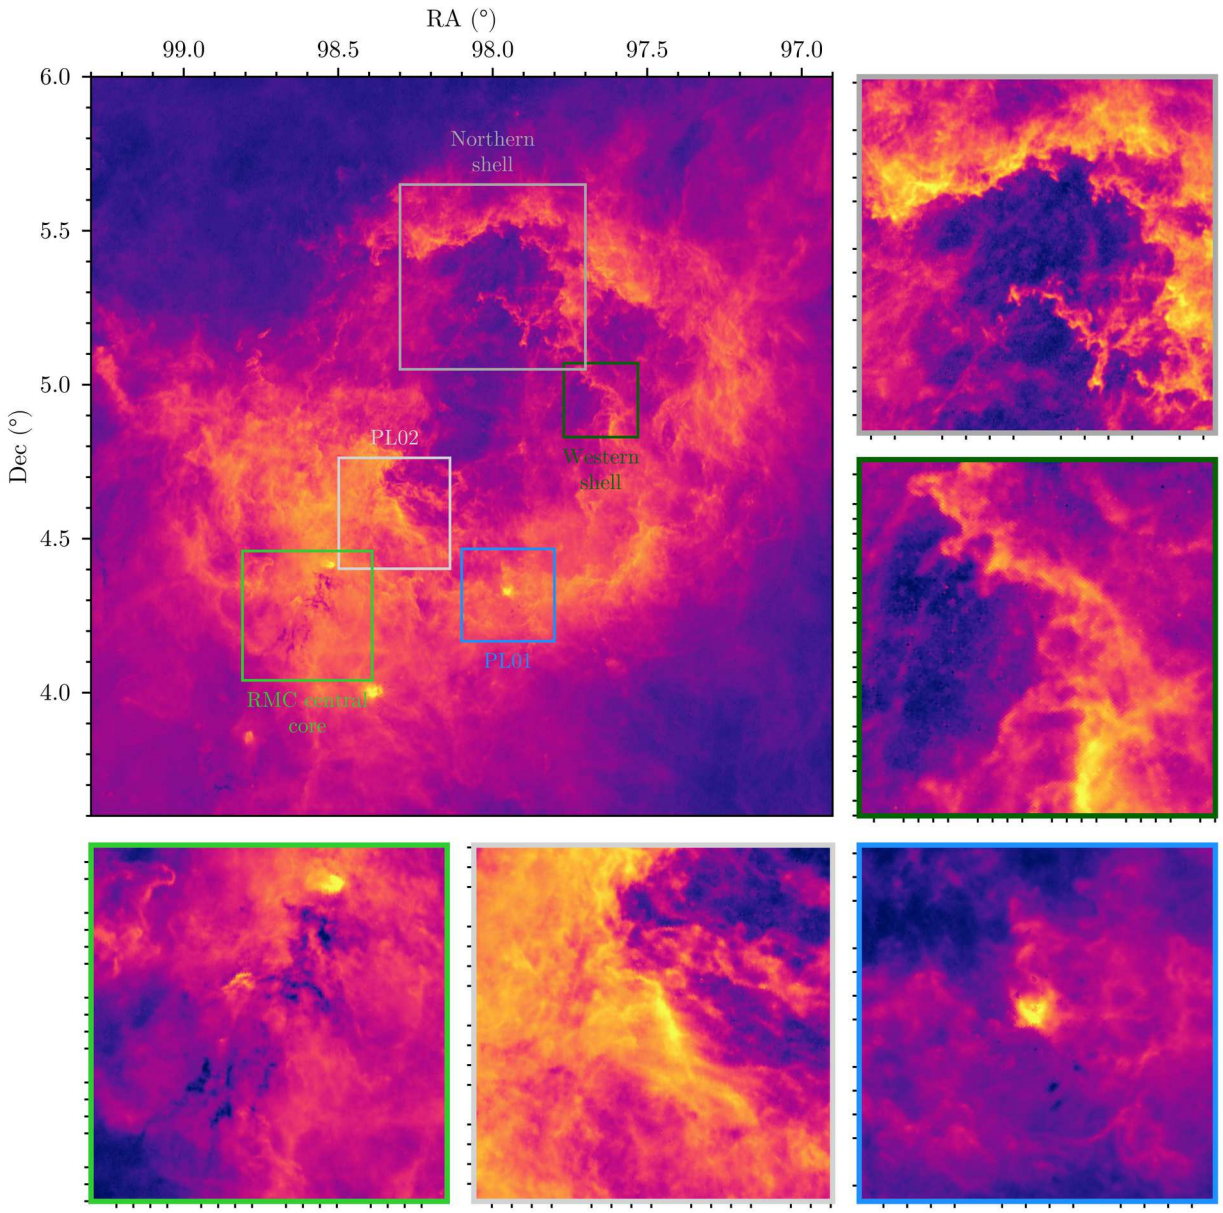

Figure 6. Overview of a few regions of interest in the Rosette nebula indicating their location on our full Rosette image and providing zoom-in subpanels for each region: the northern shell (top right), the western shell (middle right), and, along the bottom row from left to right, the central core of the Rosette molecular cloud (RMC; L. Blitz & P. Thaddeus 1980), and the embedded clusters PL02 and PL01 (R. L. Phelps & E. A. Lada 1997; C. G. Román-Zúñiga et al. 2008). Each subpanel has an independent log color scale from the 0.1–99.9$^{\rm th}$ percentile, with brighter colors representing more emission.

parallel. Within a tile the masked pixels are visited in the order of the NESTED HEALPix index, which is a hierarchical space-filling curve, so that the fill proceeds through compact sub-regions in turn rather than sweeping across the tile in a single direction.

The major change compared to A. K. Saydjari & D. P. Finkbeiner (2022) is adapting this algorithm to work on a HEALPix grid as opposed to a rectilinear grid. For RING-ordered HEALPix maps in the equatorial zone, below absolute latitudes of 41.°8, every ring holds the same number of pixels ($4N_{\rm side}$) and consecutive rings are offset by half a pixel, so an $N_p \times N_p$ patch and its shifted copies are a fixed pattern of index offsets. We previously used this trick to infill masked pixels in the DECaPS2 dust map (C. Zucker et al. 2025).

To handle the higher latitudes where the number of pixels per ring changes, we rotate each field so that its center lies on the equator of a new HEALPix frame, resample the map onto that frame (each rotated pixel takes the value of the original pixel containing its center), infill there, and carry the infilled values back to the original HEALPix pixels. Since our fields currently

extend at most a few degrees from their centers, the whole field lies deep within the equatorial zone of the rotated frame regardless of its declination. Because the two grids are different tessellations wherever the field is off the equator, a rotated pixel is treated as masked both when its center lies in a masked pixel and when it is the pixel that will supply the value of a masked pixel on the way back, so that every masked pixel receives a prediction (with its standard deviation) rather than the value of a neighboring unmasked pixel.

4. GALACTIC TESTBEDS

4.1. *Rosette Nebula*

The Rosette nebula is a well-known H II region that is ionized by OB stars in the star cluster NGC 2244, at a distance of $\sim$ 1.4 kpc (B. Lim et al. 2021). The intricate structure of the nebular emission provides an ideal testbed for our methodology, and allows us to investigate the detailed morphology of the 3.3 μm and 3.4 μm emission across a region that is illuminated by well-characterized photoionization sources.

Figure 1 shows the LVFit 3.3 μm PAH line surface brightness map of the Rosette nebula, along with the surface brightness of the LVFit continuum model at 3.29 μm. For comparison, we show a simple band-averaged map of the same region at 3.3 μm. This band-averaged, continuum-subtracted map is made by taking the inverse-variance weighted average of all SPHEREx observations in a given HEALPix pixel within two times the LVF bandwidth of the PAH central wavelength of 3.29 μm minus the average of observations between two and five bandwidths away. The naïvely band-averaged map showcases the severe banding artifacts discussed above. The LVFit 3.3 μm PAH line map shows a wealth of detailed morphological structure, and we highlight our spatial resolution with a zoomed-in version of Figure 1 in Figure 2. Adopting a distance to the Rosette nebula of 1.4 kpc, the physical resolution of our maps (8.″5 FWHM) is $\sim$ 0.06 pc. We also highlight several regions of interest in Rosette discussed in the literature in Figure 6.

4.1.1. *Paschen-α*

Understanding the multiphase spatial structure of the H II region is essential for interpreting the observed morphology of PAH emission. SPHEREx observations can also trace the warm ionized medium at the center of the H II region via hydrogen recombination lines, the strongest of which in the SPHEREx wavelength range is Paα. Notably, Paα is not accessible from the ground due to water vapor absorption as it falls between the H and K infrared atmospheric windows. In Figure 7, we show the LVFit SPHEREx Paα line surface brightness map of the Rosette nebula made by jointly modeling Paα, Brδ, and Brϵ. While Paα is not accessible from the ground, we can compare our map to recent high angular resolution maps ($\sim$ 6.4–8.″2 FWHM) of Hα from the MDW Hα Sky Survey (N. Aftab et al. 2025) (DR1), which is well-matched to SPHEREx in terms of spatial resolution (Figure 7, middle). In the right panel of Figure 7, we show the logarithmic ratio of Paα to Hα relative to the median ratio over the Rosette nebula. The median ratio of 0.253 exceeds the Case B value ($\approx$ 0.12 at 10^4 K; D. G. Hummer & P. J. Storey 1987; P. J. Storey & D. G. Hummer 1995) by an amount corresponding to $A_V \approx 1.3$ mag for the J. A. Cardelli et al. (1989) extinction curve assuming $R_V = 3.1$, consistent with the $A_V \approx 1.2$–1.5 mag of integrated foreground extinction to the distance of NGC 2244 in a GaiaXP-based 3D dust map (G. Edenhofer et al. 2024).

These SPHEREx PAH and Paα maps thus provide complementary probes of the Rosette nebula. We expect future work will combine high resolution PAH and hydrogen recombination line maps from SPHEREx, high resolution ground-based Hα maps from MDW, and lower-spatial resolution but much higher spectral resolution integral field unit (IFU) maps from the SDSS-V Local Volume Mapper (LVM) (M. A. Villa-Durango et al. 2025) in order to significantly progress efforts in reconstructing the 3D density, temperature, and ionization structure of the Rosette nebula.

4.1.2. *Aliphatic-to-Aromatic Ratio*

Because the 3.3 μm PAH is primarily associated with aromatic carbon-hydrogen bond stretches and the 3.4 μm PAH with aliphatic carbon-hydrogen stretches, the 3.4 μm-to-3.3 μm PAH ratio acts as a tracer of chemical variation. We display maps of the aliphatic-to-aromatic ratio for the Rosette nebula in Figure 8. The morphology of the ratio map is strikingly distinct from the morphology of the 3.3 μm line map: the ratio map shows relatively enhanced 3.4 μm emission in the smaller, clumpier structures that line the boundaries of the H II region, and a relative dearth of 3.4 μm emission both interior to the Rosette shell and in the more diffuse regions in the outskirts of the Rosette nebula. This is generally consistent with trends observed with SPHEREx data toward the Iris nebula (C. Boersma et al. 2026), but LVFit enables a spectacularly spatially resolved view of the relative strength of these two spectral features. Interpreting the 3.4 μm feature as arising primarily from aliphatic groups attached to PAHs (L. J. Allamandola et al. 1985; A. G. G. M. Tielens 2008), we infer that these moieties or molecules are more easily photolyzed when exposed to far-UV photons, and thus only survive in relative abundance when shielded from the far-UV radiation field in denser clumps of dust (C. Joblin et al. 1996). This interpretation is consistent with the particularly strong aliphatic-to-aromatic ratio coincident with the Rosette molecular cloud and the dense edges of the northern shell. We plan to leverage the power of our enhanced spatial resolution in future work to carefully model the destruction thresholds for these PAHs in terms of incident UV flux and dust density.

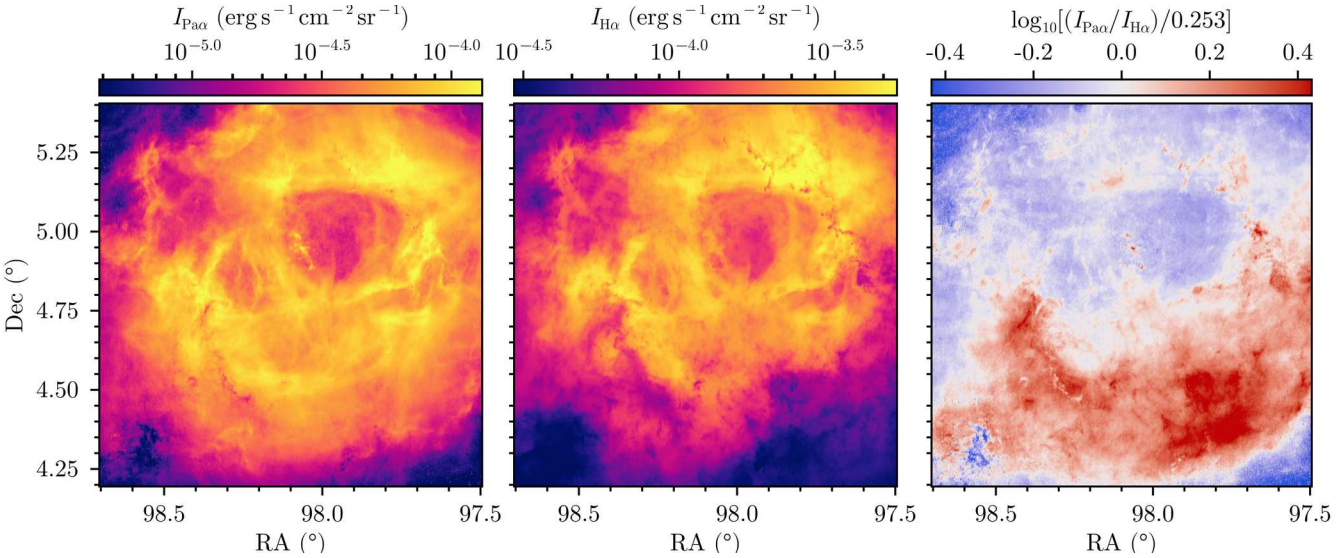

Figure 7. The LVFit Paα SPHEREx map (left) compared with Hα from the star-removed MDW Hα Sky Survey DR1 mosaic (N. Aftab et al. 2025) (center), and their ratio $\log_{10}[(I_{\mathrm{Pa}\alpha}/I_{\mathrm{H}\alpha})/0.253]$ (right), where 0.253 is the median ratio over the displayed region. The Hα surface brightness uses the nominal MDW conversion of 86.7 R per MJy sr^{-1} and includes [N II] and continuum contributions, so its absolute level is uncertain at the $\sim$ 10% level. Red regions toward the Rosette molecular cloud and along dark globules indicate enhanced extinction of Hα. The MDW stacks have a PSF FWHM of $\sim$ 6.4–8.″2 (median over stars in the four fields used), comparable to the 8.″5 FWHM angular resolution of our map.

While the absolute ratio of 3.4/3.3 μm PAH emission depends strongly on convention for partitioning overlapping spectral features, including the continuum, we can more easily compare the relative values we measure here. Compared to the PDR shell (0.°3–0.°6), the median $I_{3.4}/I_{3.3}$ inside the H II cavity ($r < 0.°25$) is 15% lower and within 0.°3 of the Rosette molecular cloud central core is 14% higher. We find 3.4/3.3 μm ratios are $\sim$ 15% higher than the Rosette average in our molecular cloud testbeds such as Corona Australis (Section 4.2). Mapping and understanding the variations of this ratio in detail across all Galactic environments would be a valuable topic for future work.

Under our 3.4/3.3 μm convention, the median absolute ratio in Rosette is 0.44. To compare to T. I. Mori et al. (2014), who studied the 3.4/3.3 μm ratio in 36 Galactic H II regions, we estimate that assigning the PAH plateau beneath the 3.3 and 3.4 μm PAH features to the continuum similar to their method would reduce our ratios by a factor of 0.42, to 0.16–0.21 for Rosette, which agrees well with their quoted ratios of 0.25 ± 0.1. These can then be converted into an upper bound on the fraction of C atoms in aliphatic chains versus aromatic rings, which gives $\sim$ 3% (X. J. Yang et al. 2016). This conversion leverages theoretical oscillator strengths to predict the 3.4/3.3 μm ratio if there were equal numbers of aliphatic and aromatic C-H bonds, though the conversion factor varies by a factor of two depending on assumptions about what particular aliphatic species are present. We use the PAH continuum-subtracted ratio for this conversion because the PAH continuum in this region is often attributed to anharmonic hot bands (J. R. Barker et al. 1987) and other transitions not included in the theoretical calculations of X. J. Yang et al. (2016).

4.2. *Corona Australis*

Corona Australis is a nearby ($\sim$ 150 pc, C. Zucker et al. 2020) molecular cloud that is notable for its position well off the Galactic plane. Its proximity combined with its gradient in density makes it a great testbed for validating morphological recovery of an isolated cloud with variable signal-to-noise. We show the LVFit 3.3 μm PAH map in the left panel of Figure 9 and the 12 μm WISE PAH map from A. M. Meisner & D. P. Finkbeiner (2014) in the right panel, for comparison. While the WISE W3 PSF is 6.″5 FWHM (E. L. Wright et al. 2010), A. M. Meisner & D. P. Finkbeiner (2014) smoothed their map to 15″ FWHM and so we present the LVFit map smoothed to 15″ FWHM as well to facilitate the comparison.

In general, we find that the LVFit 3.3 μm PAH map matches the 12 μm WISE PAH map morphologically, but with fewer artifacts. The off-axis optical design of SPHEREx (J. J. Bock et al. 2026) significantly reduces diffraction spikes compared to the on-axis design of WISE (E. L. Wright et al. 2010), allowing us to more easily remove stars. This is especially apparent for the bright sources in the Coronet cluster at the top left of Figure 9, which also show significant pupil ghosts below the bright stars in the WISE map that are not present in our SPHEREx map. The WISE map also has a large defect at the middle left of the image associated with a satellite transit. On our first processing of the SPHEREx data, a faint linear streak was present in the bottom

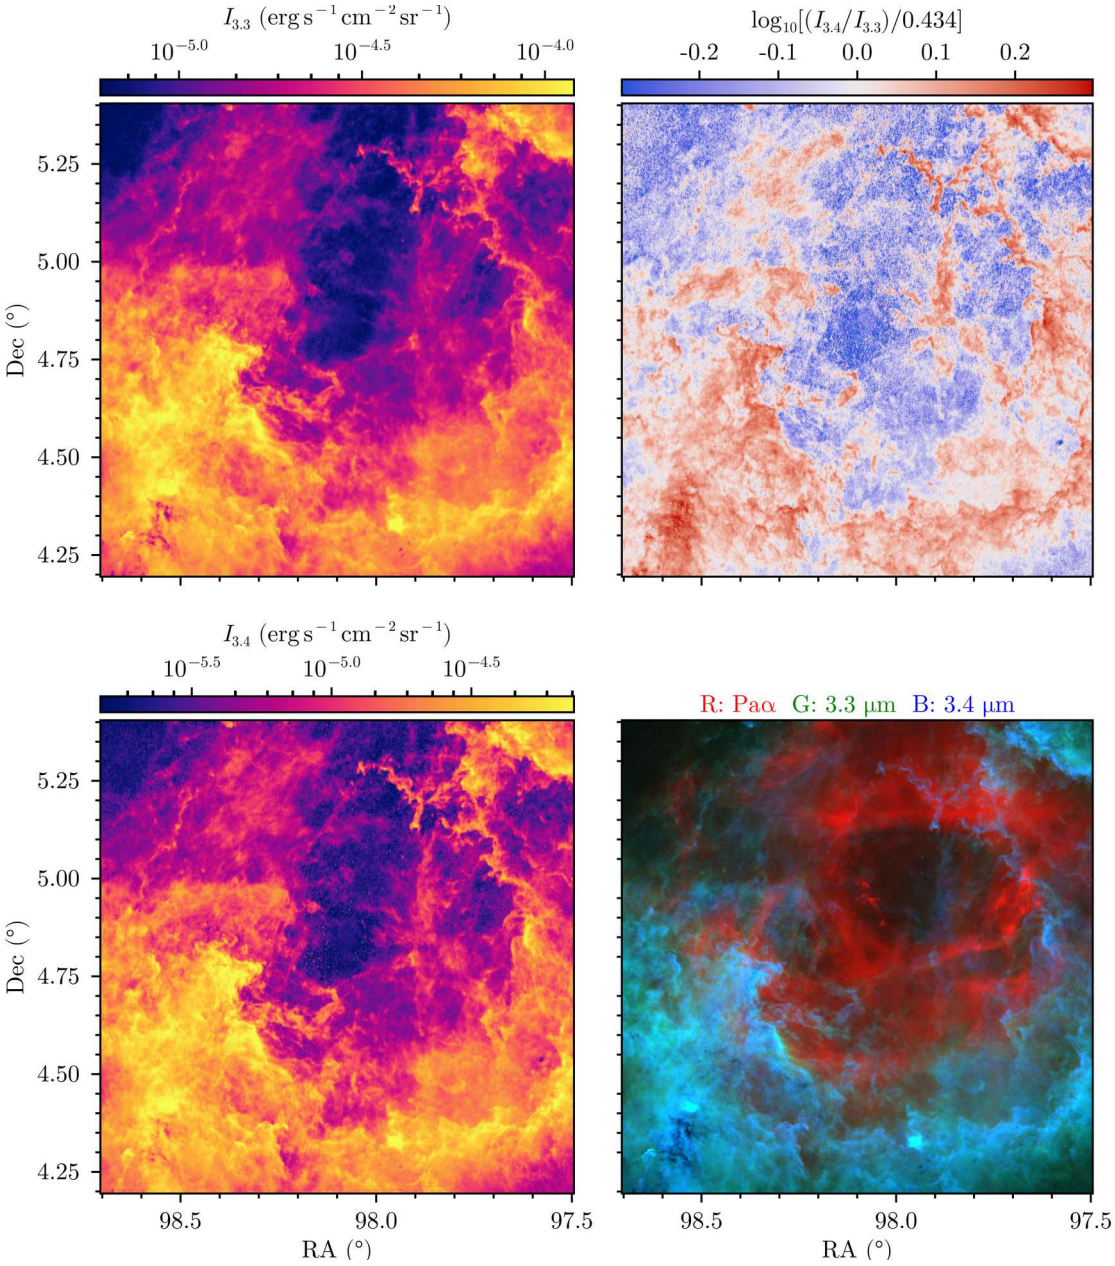

Figure 8. PAH 3.3 μm (top left) and 3.4 μm (bottom left) line maps for the Rosette nebula. The aliphatic-to-aromatic ratio $\log_{10}[(I_{3.4}/I_{3.3})/0.434]$ is shown in the top right, where 0.434 is the median ratio over the displayed region. Red regions indicate locations with enhanced 3.4 μm compared to the average. Bottom right is an RGB composite (R: Paα, G: 3.3 μm, B: 3.4 μm, each normalized by its median and asinh-stretched) which shows how the ionized shell traced by Paα lies inside the PAH-emitting photodissociation region.

right of Figure 9, likely from a moving solar system object such as an asteroid. The affected measurements were subsequently excluded from their respective per-pixel fits by our streak detector (see Section 3.2), which flags six such tracks in five of the 110 exposures of this field.

Another artifact in the WISE map results from A. M. Meisner & D. P. Finkbeiner (2014) adjusting the zero-points to agree with Planck on scales above 2°to handle drifts in the WISE zero-point. This leads to "ringing" over the sharp edge of the cloud, where the 12 μm emission appears to decrease and then return to a value typical of the rest of the low density regions in the image. In contrast, no such background adjustments were required for our SPHEREx processing and the background 3.3 μm emission appears uniform in the top right of the image. We will discuss the sensitivity limits and comparisons in the low-density regime further in Section 6.1.

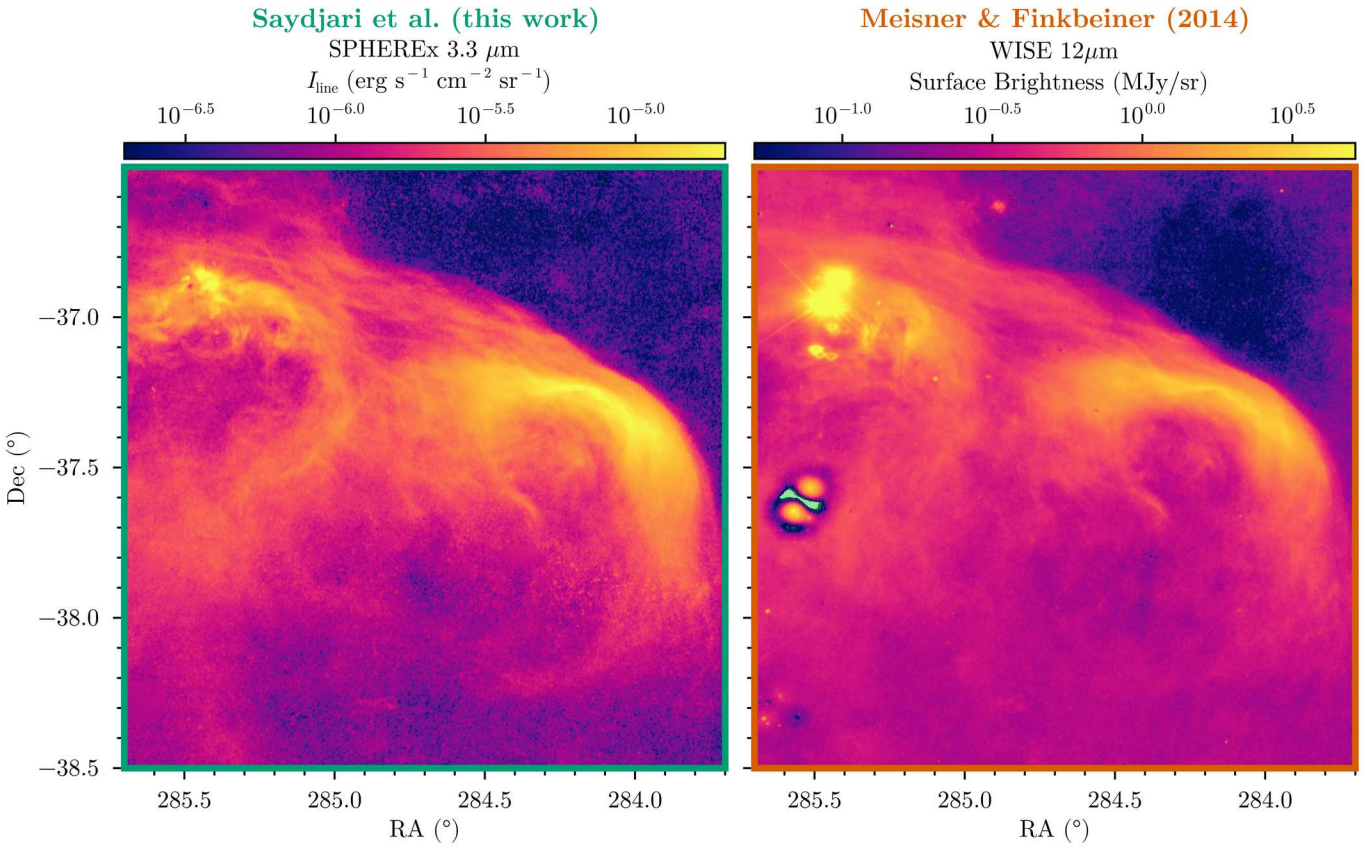

Figure 9. A portion of the Corona Australis dust cloud used to compare (left) the 3.3 μm PAH map from LVFit and (right) the WISE 12 μm PAH map (A. M. Meisner & D. P. Finkbeiner 2014). To facilitate comparison, we smooth the LVFit map to match the 15″ FWHM angular resolution of A. M. Meisner & D. P. Finkbeiner (2014).

After accounting for these artifacts, the ratio of the SPHEREx 3.3 μm and the WISE 12 μm maps could provide valuable insight into chemical variation between molecules traced by 3.3 μm emission versus those traced by WISE W3. However, this comparison is far more fraught than interpreting the 3.4/3.3 μm maps because there are several different PAH emission complexes (7.7 μm, 8.6 μm, 11.3 μm, 12.7 μm) that all contribute to the WISE W3 band (7.5–16.5 μm; E. L. Wright et al. 2010). The relative strength of the 7.7 μm and 11.3 μm features within this band depends strongly on the PAH ionization state (B. T. Draine et al. 2021a).

These initial maps of Corona Australis are especially useful because the cloud will serve as an important testbed for at least two aspects of future work. First, literature H_2 column density maps from Herschel (P. André et al. 2010; HGBS Team 2020) will serve as an important reference for modeling H_2 ro-vibrational transitions in SPHEREx. In addition, the reflection nebulae NGC 6726/6727/6729 near the Coronet cluster provide an important testbed for how to handle strongly spatially localized coupling between point sources and diffuse emission. When trying to make diffuse maps, the conundrum is always how to ensure extended light from point sources in the PSF wings is sufficiently removed while preserving any nearby diffuse structure. For the 3.3 μm PAH map, this is even more pressing since the UV flux from stars may excite PAH emission in the nearby ISM, so there is a causal connection rather than simply a chance overlap between point sources and extended emission. Simple masking and infilling approaches like those described in this work, even after subtracting off our best model of the stellar light, often over-mask and miss some of the nearby substructure in these regimes. More advanced simultaneous modeling of point sources and the diffuse components may be required to better preserve the substructure of the reflection nebulae (see Section 7 for more).

4.3. *Cygnus X and the Heart and Soul Nebulae*

We show two further Galactic testbed regions that have been featured in recent SPHEREx works. In Figures 10 and 11 we show Cygnus X, which was studied by J. L. Hora et al. (2026). Cygnus X hosts many embedded young stellar objects and compact H II regions whose PAH emission our star mask partly removes, the same conundrum we discuss with regards to the Corona Australis reflection nebulae (Section 4.2). We plan to use this field to test the joint modeling of point sources

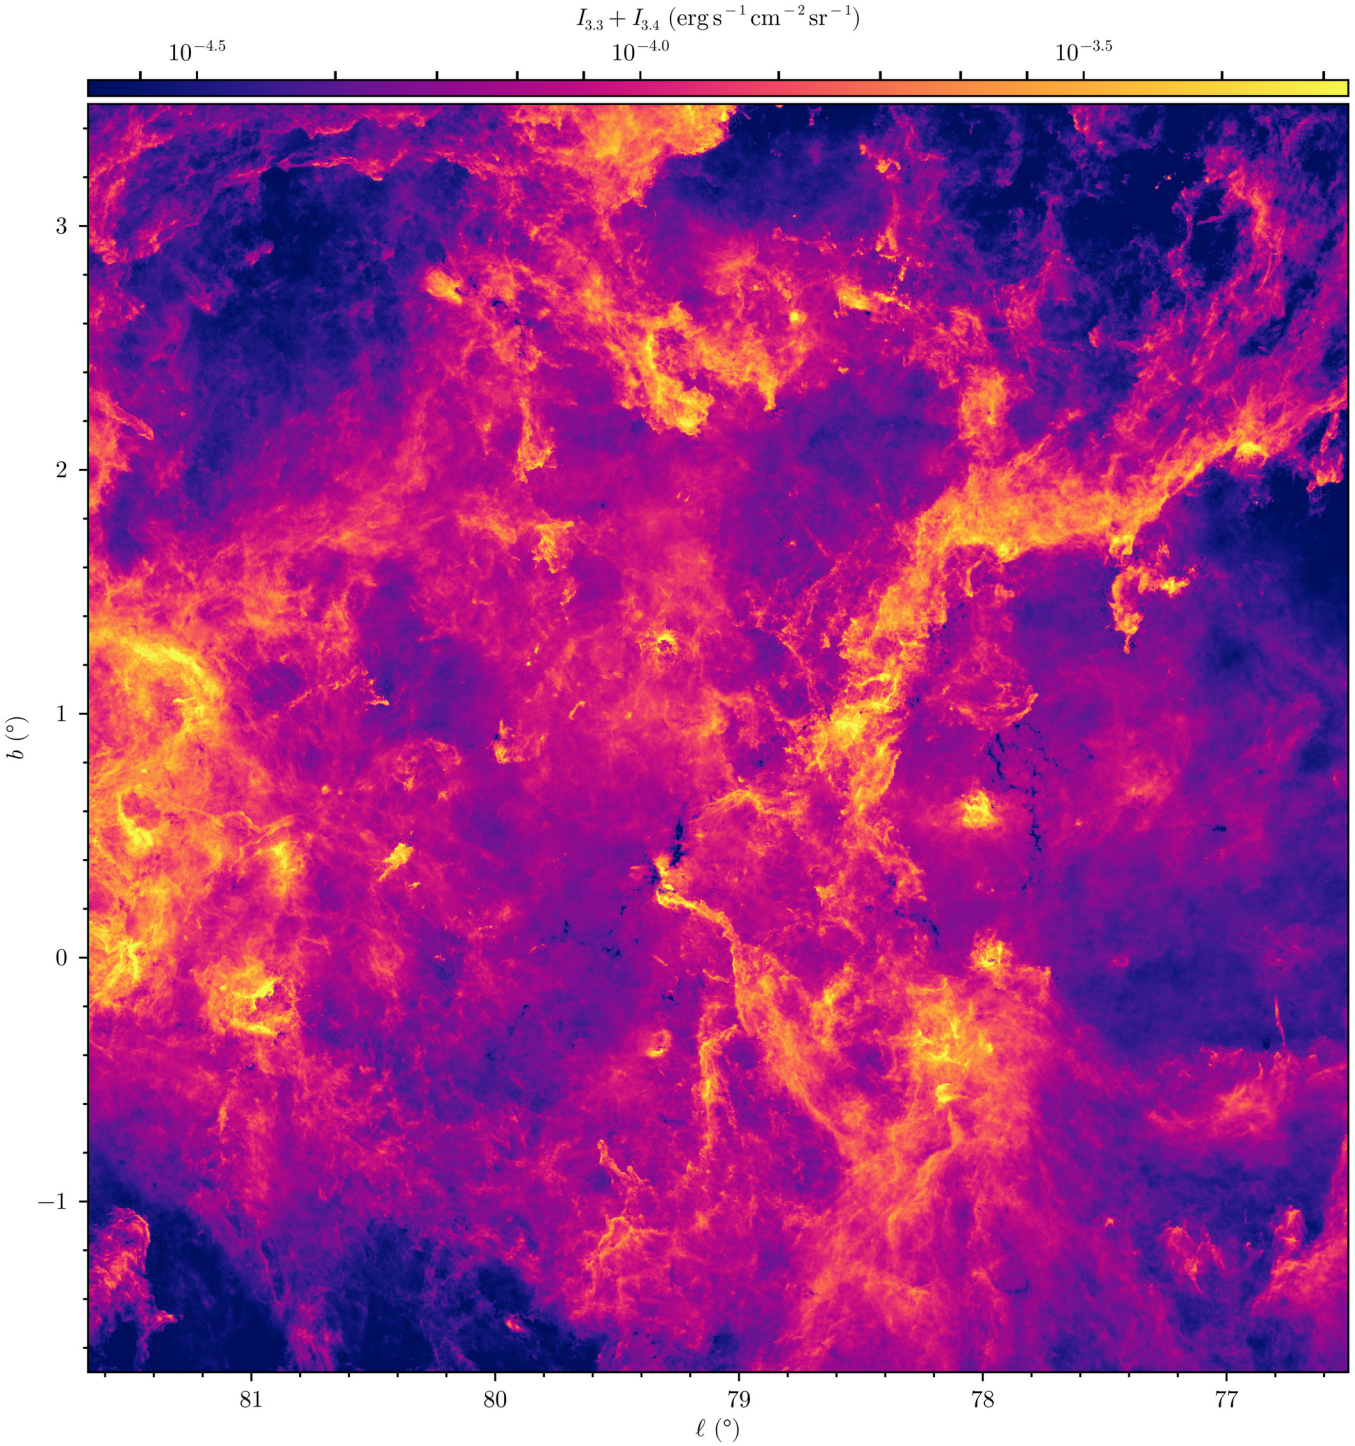

Figure 10. An LVFit map of Cygnus X in Galactic coordinates showing the sum of the 3.3 and 3.4 μm PAH emission, $I_{3.3} + I_{3.4}$, on a log color scale spanning the 2$^{\rm nd}$–99.9$^{\rm th}$ percentiles. This field was studied by J. L. Hora et al. (2026). About a third of the pixels in this crowded field are masked (mostly faint stars) and infilled.

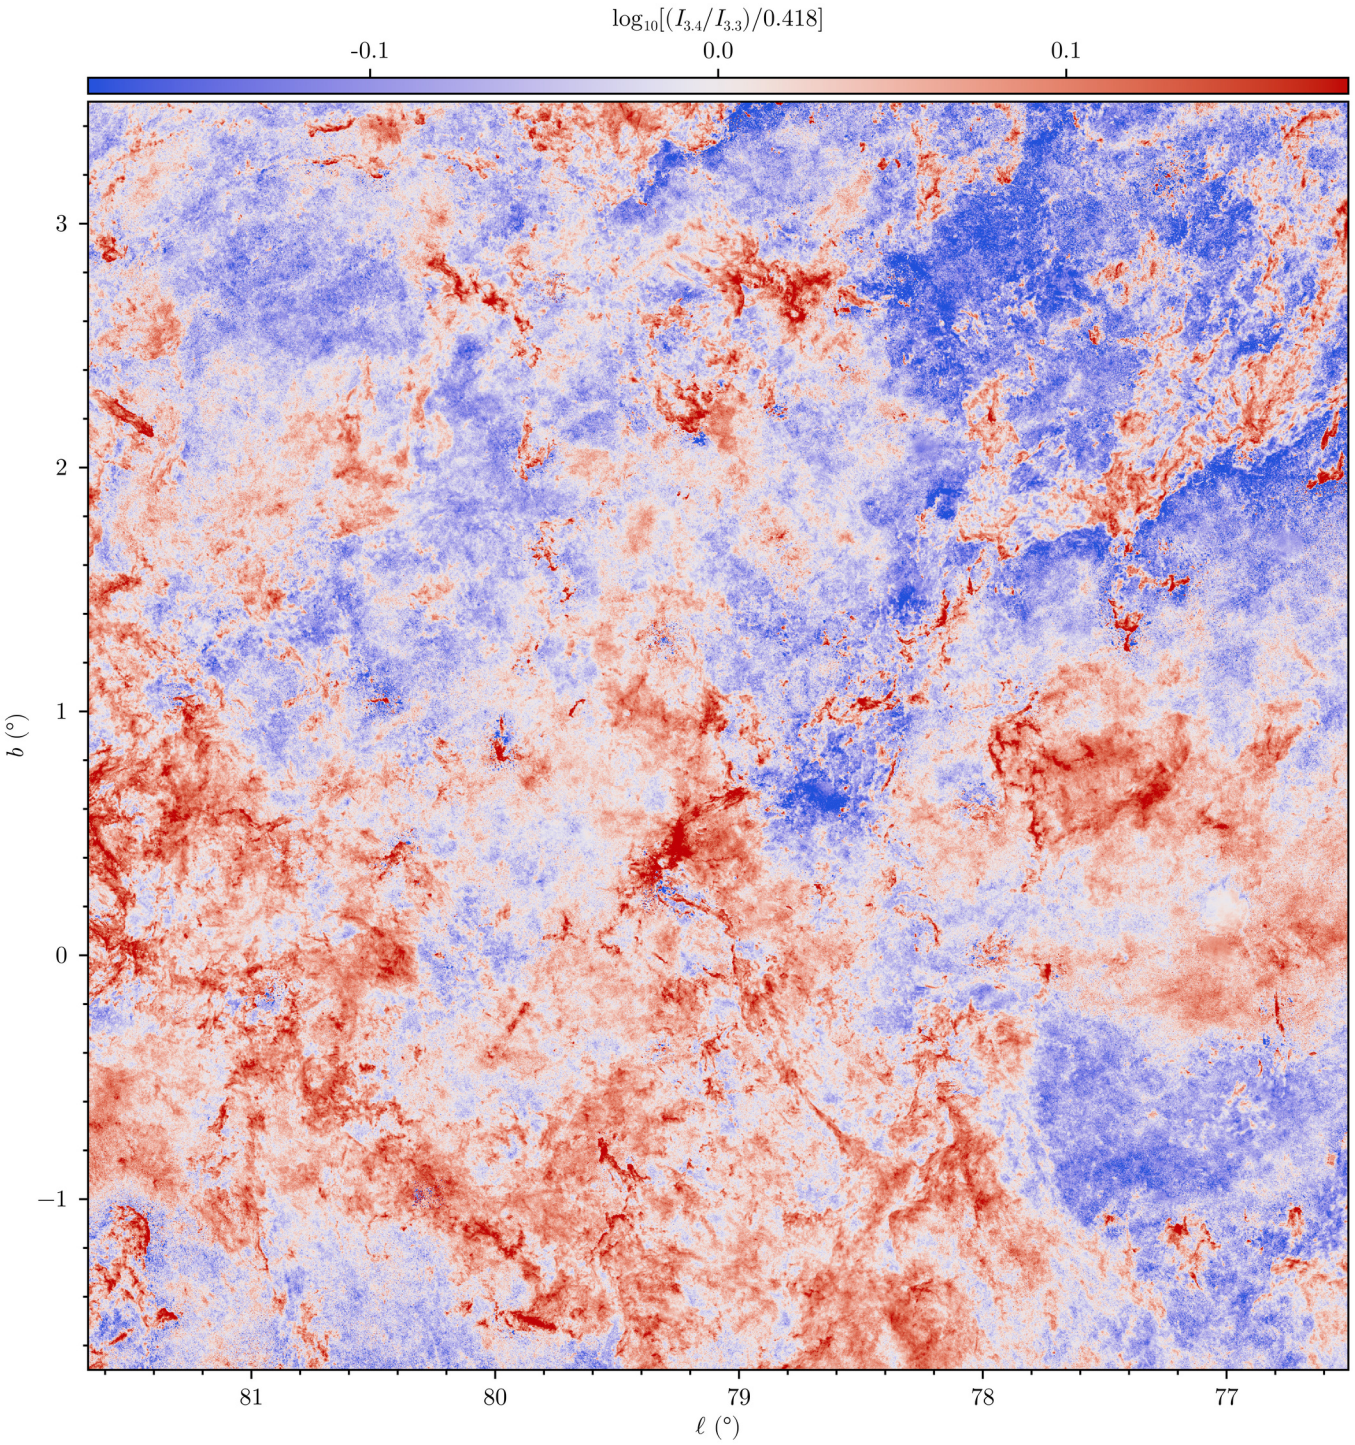

Figure 11. The aliphatic-to-aromatic ratio $\log_{10}[(I_{3.4}/I_{3.3})/0.418]$ for the Cygnus X field of Figure 10, where 0.418 is the median ratio over the displayed region. Red indicates enhanced 3.4 μm emission relative to the median.

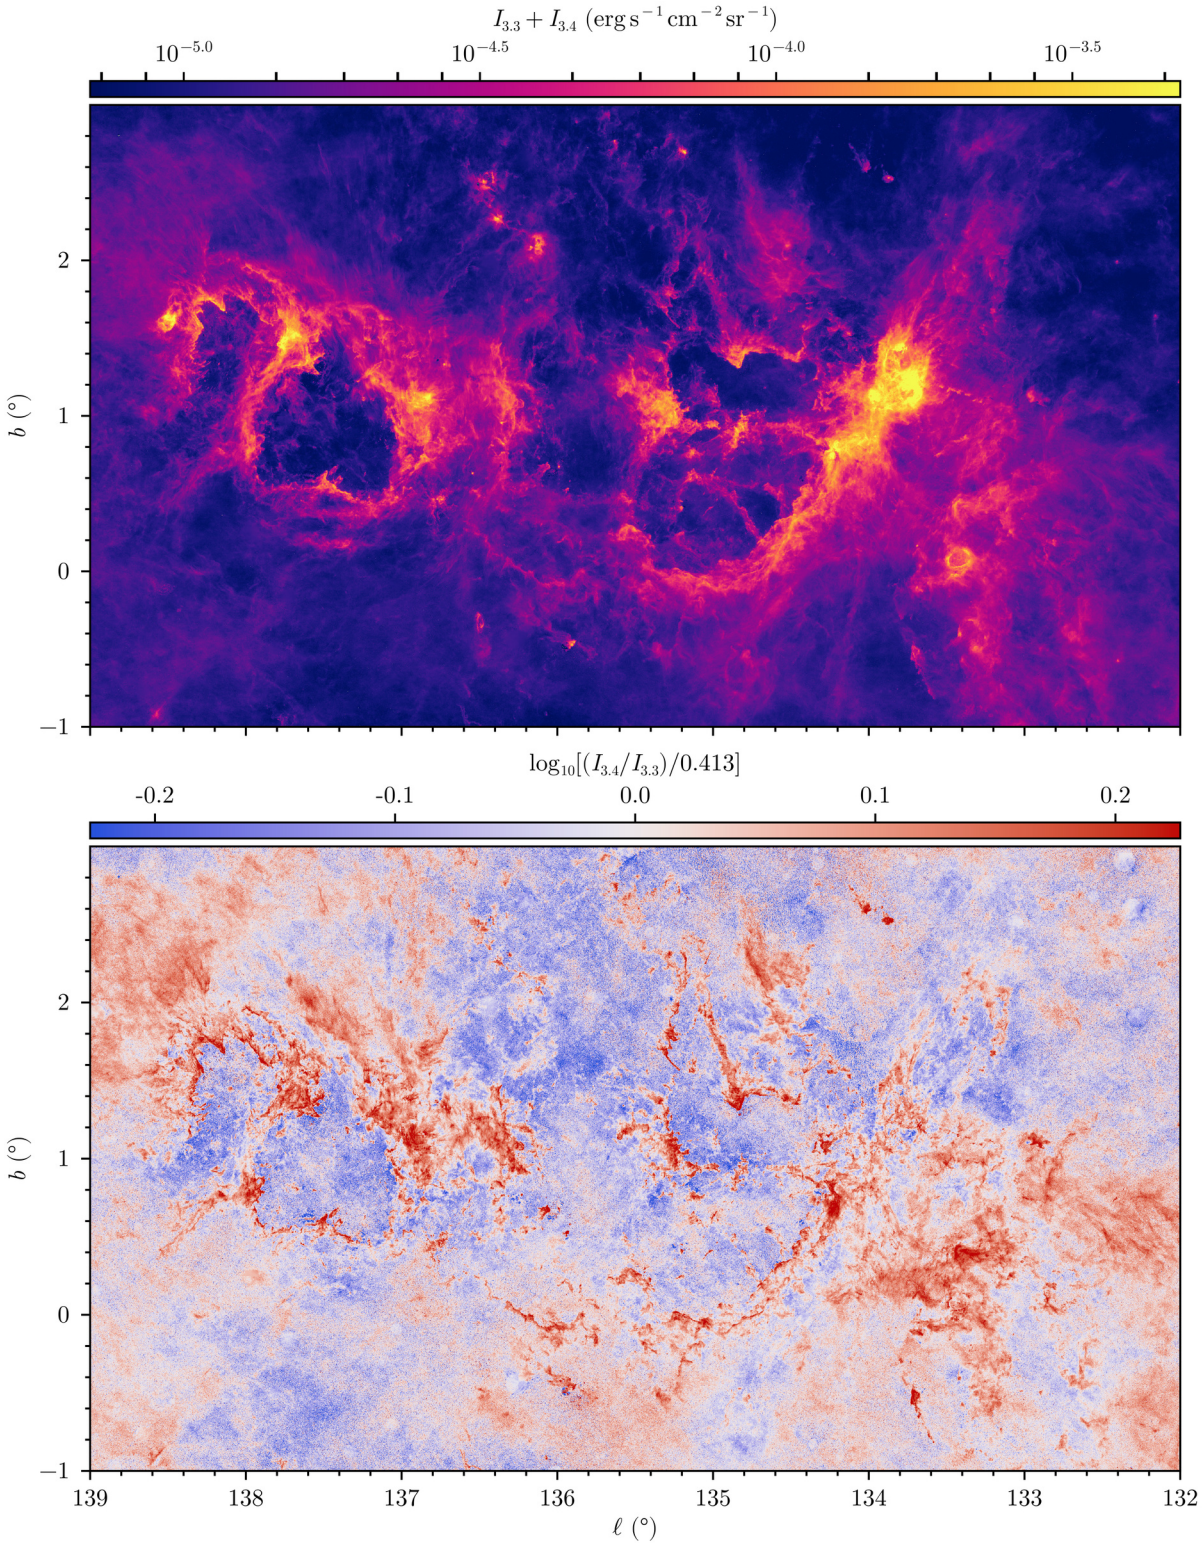

Figure 12. The Heart (IC 1805, W4) and Soul (IC 1848, W5) nebulae in Galactic coordinates, a 7° × 4° field centered on $(\ell, b) = (135.^{\circ}5, +1.^{\circ}0)$, which was studied by A. J. Cukierman et al. (2026). Top: The sum of the LVFit 3.3 and 3.4 μm PAH line maps, $I_{3.3} + I_{3.4}$, on a log color scale spanning the 2$^{\rm nd}$–99.9$^{\rm th}$ percentiles. Bottom: The aliphatic-to-aromatic ratio $\log_{10}[(I_{3.4}/I_{3.3})/0.413]$, where 0.413 is the median ratio over the displayed region. Red indicates enhanced 3.4 μm emission relative to the median.

and diffuse emission with spectral priors in order to better preserve these somewhat compact regions of diffuse emission. In Figure 12, we show the Heart and Soul nebulae, which were studied by A. J. Cukierman et al. (2026). Neither makes the data behind their figures publicly available for comparison.

5. EXTRAGALACTIC TESTBEDS

5.1. *M31: Andromeda*

In the top row of Figure 13, we show the 3.3 μm PAH map of M31 from LVFit, with a zoom-in on the center of M31 in the right column. Twelve moving-object tracks in 11 exposures have their pixels rejected by our streak detector, and 29 of the 253 exposures have the detector edge masked for scattered moonlight (see Section 3.2). For comparison, we show the H I column density at 24″ from the Local Group L-Band Survey (LGLBS) (E. W. Koch et al. 2025). One important difference between these maps is that we integrate over the full line of sight, including foreground PAH emission from Galactic cirrus, while the LGLBS mosaic shown applies their broad "signal mask" to reject H I not associated with M31.

By comparing these maps, we see that our PAH map successfully resolves substructure within the rings of M31 and low density rings/spiral structure just outside the nucleus that are also traced by the H I gas. However, in the nucleus, our map appears to show PAH emission. Due to the lack of H I in that region in the LGLBS image and the low dust mass surface density measured by Spitzer (B. T. Draine et al. 2014), we interpret this as continuum emission from stars that is not fully removed in this extremely crowded limit. This same mechanism of leakage of continuum into the PAH map appears to impact the satellite dwarf elliptical galaxy M110 (top right quadrant of the left column), which is not expected to have significant PAH emission due to its low metallicity (D. R. Gonçalves et al. 2014) near the threshold below which PAH emission is deficient (B. T. Draine & A. Li 2007b), lack of ongoing star formation (L. Monaco et al. 2009), and low dust content (I. De Looze et al. 2012). Jointly forward modeling both the resolved and unresolved point sources with the diffuse emission is one possible solution to this limitation, which we discuss further in Section 7. Revisiting M110 and similar objects after we have better control and quantification of continuum leakage may either provide a detection or valuable upper bound on PAH content in Local Group dwarf ellipticals.

In the bottom row of Figure 13, we compare the ratio of the PAH 3.3 μm emission and H I column density maps. In between the edges of the LGLBS map and the continuum contamination in the nucleus, we appear to clearly trace enhanced PAH emission in M31's ring known for enhanced star formation (K. D. Gordon et al. 2006; B. T. Draine et al. 2014; A. R. Lewis et al. 2015). A deeper investigation of individual sites of star formation and variation in the 3.3 μm and 3.4 μm content of those regions as a function of galactic environment would be an exciting avenue for future work.

5.2. *M33: Triangulum*

Similar to M31, the LVFit 3.3 μm PAH map of M33 (Figure 14, top row) can be used to make maps of PAH-to-neutral-H and trace star formation in a Local Group galaxy. The LVFit map can also be compared to the simple continuum-subtracted SPHEREx 3.3 μm M33 map presented in A. J. Cukierman et al. (2026). As noted in A. J. Cukierman et al. (2026), the SPHEREx data enable analyses of the ratio between the 3.3 μm emission and the point-source-subtracted WISE 12 μm map of M33 presented in T. H. Jarrett et al. (2019), similar to our discussion in Section 4.2.

One other important legacy dataset for comparison is Spitzer. As we show in the middle row of Figure 14, M33 has 8 μm IRAC imaging from Spitzer (M. W. Werner et al. 2004; G. G. Fazio et al. 2004) that looks extremely similar to the LVFit map, but which has a PSF with 2″ FWHM. Here we show the sky level-subtracted Spitzer Enhanced Imaging Products super mosaic (SEIP; SSC and IRSA 2020). To put the mosaic tiles on a common background before forming the ratio, we solve for a constant offset per tile from the tile overlaps, which differ by up to a few tenths of a MJy sr^{-1} at 8 μm. One important difference from the LVFit SPHEREx map is that the 8 μm IRAC image is not point source subtracted. In terms of PAH emission, the 8 μm IRAC band covers 6.4–9.4 μm and primarily traces the 7.7 μm PAH feature, with contributions from 6.2 μm and 8.6 μm PAH features, all of which are typically associated with ionized PAH carriers (Table 1). In contrast, the 3.3 μm PAH is typically associated with aromatic C-H stretches in small neutral PAHs (B. T. Draine & A. Li 2007a; A. Maragkoudakis et al. 2020; B. T. Draine et al. 2021a), and so careful analysis of the 3.3/7.7 μm ratio can trace variation in the charge state of PAHs in the ISM.

In the bottom row of Figure 14 we show the 3.3/7.7 μm PAH ratio after interpolating and smoothing both maps to a common 10.″7 FWHM angular resolution. We apply the stellar continuum correction of G. Helou et al. (2004) to obtain $I_{8,\mathrm{PAH}} = I_8 - 0.232\, I_{3.6}$, which removes the stellar continuum using IRAC 3.6 μm. This should be a good first order correction for variation in the stellar continuum, but the IRAC 3.6 μm band used for correction does also contain 3.3 μm PAH emission that may be important to model to measure the 3.3/7.7 μm PAH ratio in detail. Many of the bright knots in the ratio map are simply due to the lack of point-source subtraction in the 8 μm IRAC image.

In the LVFit ratio map, we see that the 3.3/8 μm PAH ratio increases radially by $+0.048 \pm 0.003$ dex/kpc from the center to the edge of M33. The quoted uncertainty is statistical. The level of the 8 μm background is the dominant systematic, and shifting it by the scatter of the mosaic tile offsets moves the slope by ~ 0.01

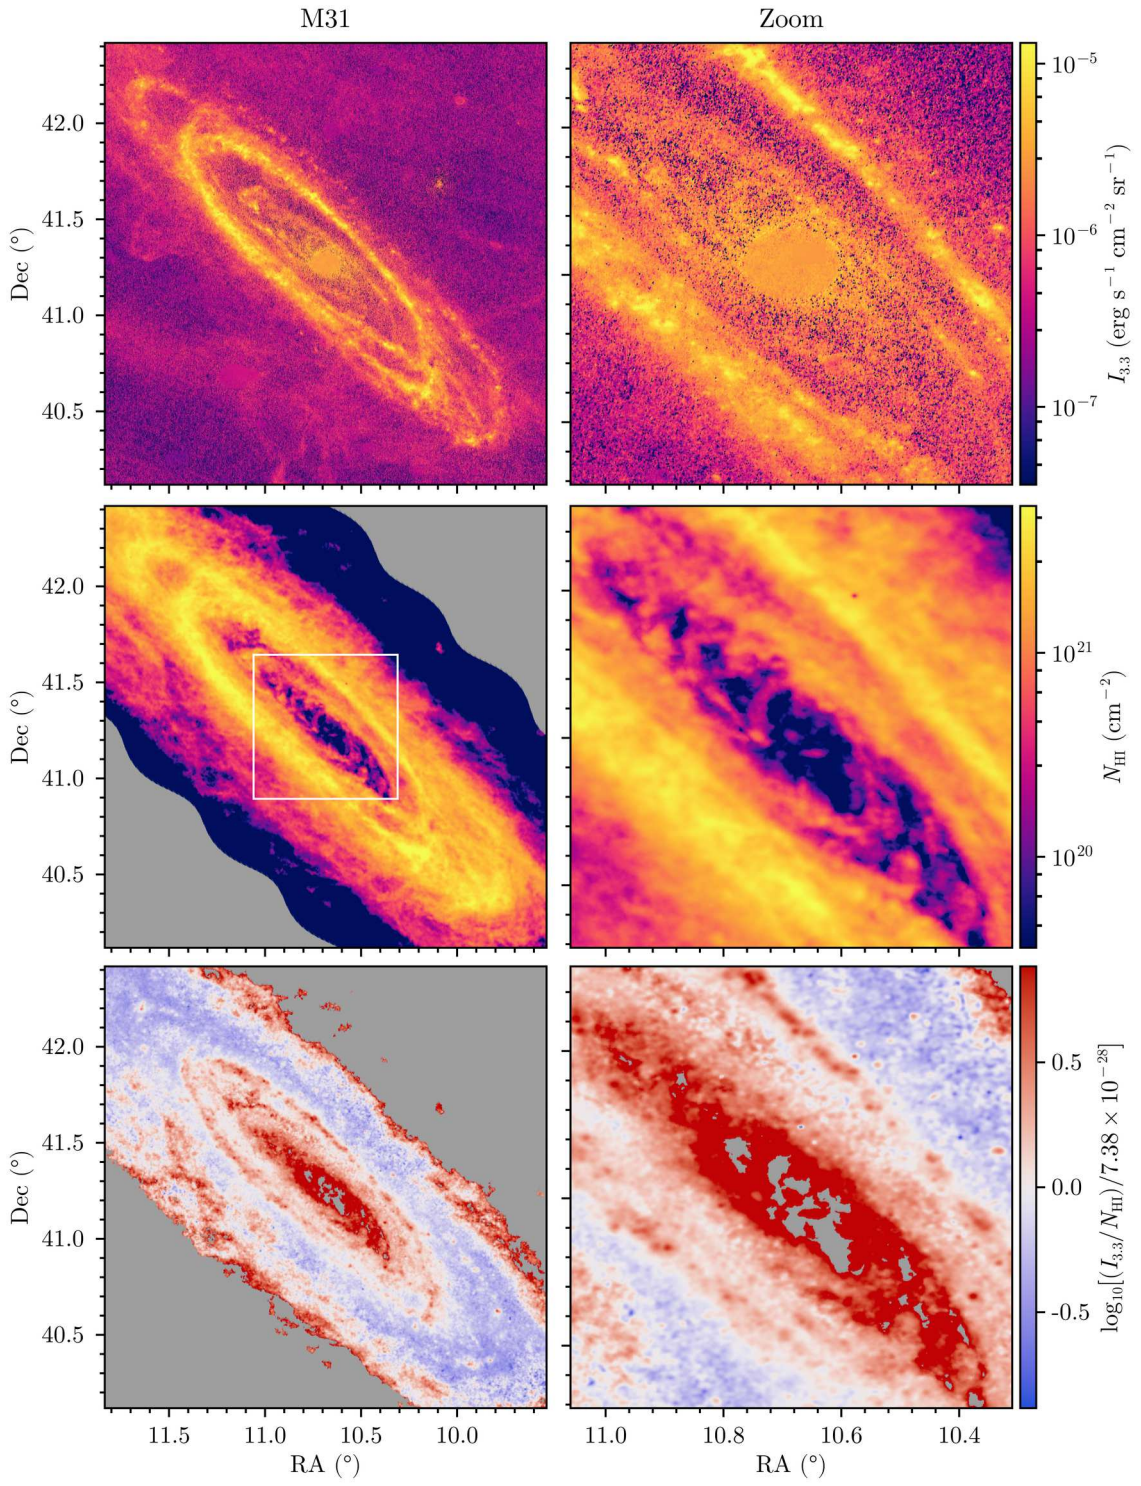

Figure 13. M31 (left) and a zoom-in (right) on the region marked by the box in the middle-left panel. Top: LVFit SPHEREx 3.3 μm PAH line map. Middle: H I column density from the Local Group L-Band Survey (LGLBS) VLA C+D+GBT moment-0 map (E. W. Koch et al. 2025) at its native 24″ beam, assuming optically thin emission. Gray pixels are areas outside the H I mosaic. Bottom: $\log_{10}[(I_{3.3}/N_{\rm HI})/7.38 \times 10^{-28}]$, with our map convolved to 24″ FWHM and restricted to H I S/N > 3, where we normalized by the average erg s^{-1} sr^{-1} per H atom in the median outside the central 10′. The PAH emission per H atom appears enhanced in the star-forming ring and toward the bulge, though starlight may contaminate the bulge measurement.

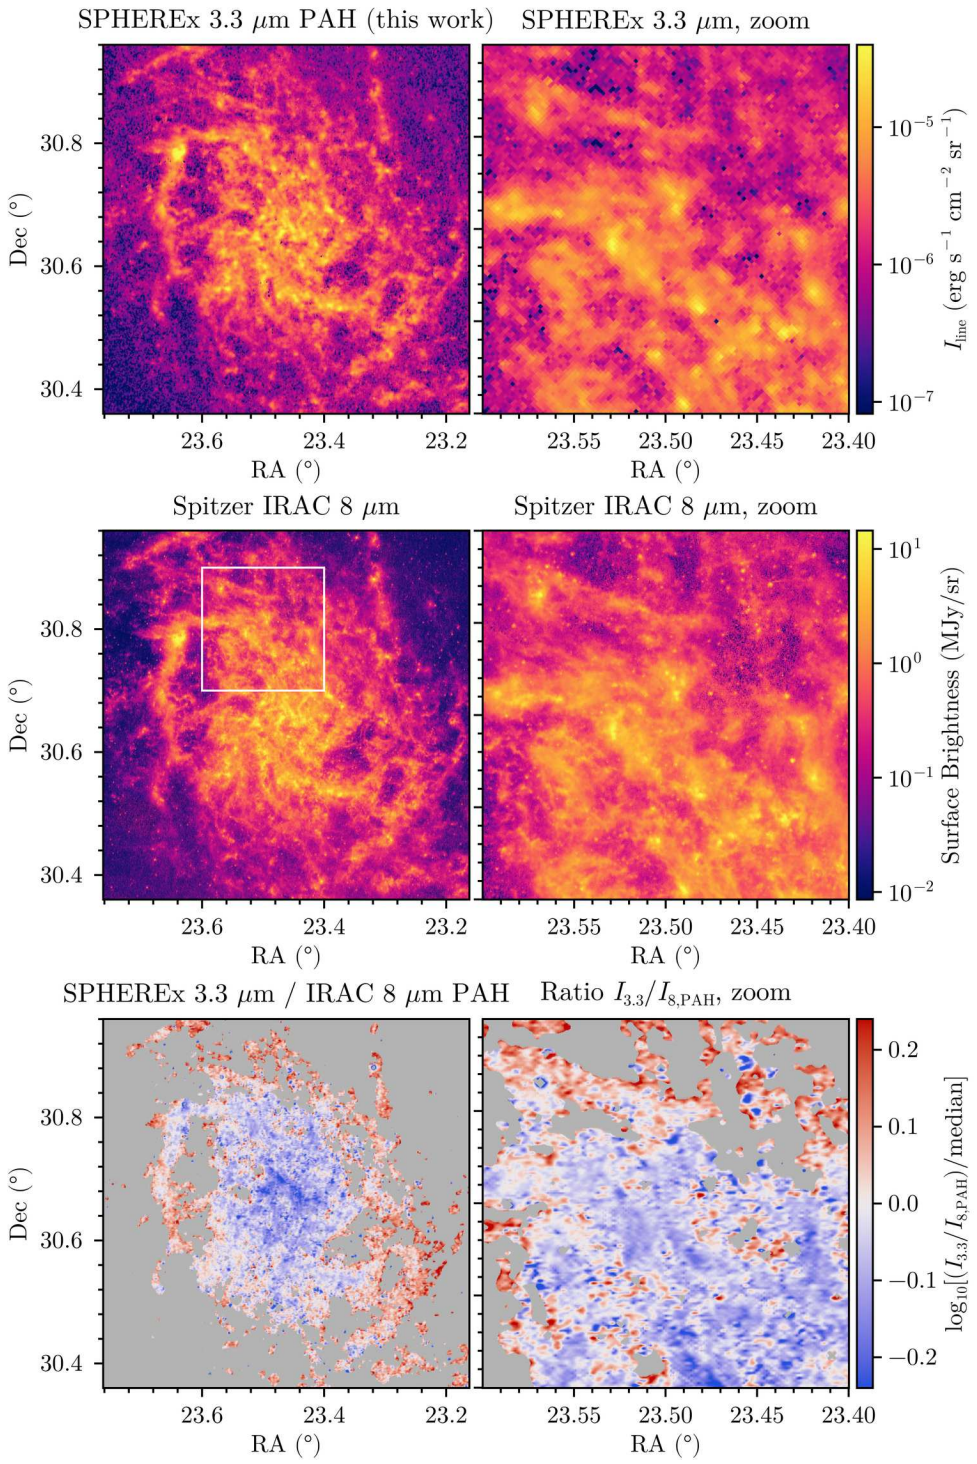

Figure 14. Top: The LVFit SPHEREx 3.3 μm PAH map of M33 (left) and a zoom-in on the region marked by the white box in the middle left panel (right). Middle: The same regions in the Spitzer IRAC 8 μm band, which is dominated by the 7.7 μm PAH complex and is shown at its native $\sim 2''$ FWHM angular resolution without smoothing, as a higher-resolution reference for the structure in the LVFit map ($8\farcs5$ FWHM). Bottom: The ratio of the LVFit 3.3 μm map to the stellar-continuum-corrected 8 μm PAH emission, $\log_{10}[(I_{3.3}/I_{8,\mathrm{PAH}})/\mathrm{median}]$, where both maps are interpolated and smoothed to a common $10\farcs7$ FWHM angular resolution. The median ratio is $4.3 \times 10^{-6}\,\mathrm{erg\,s^{-1}\,cm^{-2}\,sr^{-1}}$ per $\mathrm{MJy\,sr^{-1}}$. Pixels shaded gray are either in our star mask or have $I_{8,\mathrm{PAH}} < 0.25\,\mathrm{MJy\,sr^{-1}}$.

dex/kpc. This gradient could be reflective of either variations in metallicity or the radiation field spectrum. Recently, H. B. Koziol et al. (2026) studied this question with JWST NIRCam imaging of 19 PHANGS galaxies and the PAH emission model grids from B. T. Draine & A. Li (2007a) and B. T. Draine et al. (2021a) to conclude the 3.3 μm PAH radial gradient is driven primarily by metallicity, with a slope of $\mathrm{d}\log_{10}(I_{3.3}/I_{7.7})/\mathrm{d}(12+\log_{10}(\mathrm{O/H})) = -0.85\pm0.08$. We can compare our observations to this using the metallicity gradient measured by the CHemical Abundances Of Spirals (CHAOS) project from nebular and auroral emission lines in 65 H II regions with the Large Binocular Telescope (LBT) (N. S. J. Rogers et al. 2022). They find -0.037 ± 0.007 dex/kpc in $12 + \log_{10}(\mathrm{O/H})$. Combined, this gives us $\mathrm{d}\log_{10}(I_{3.3}/I_8)/\mathrm{d}(12+\log_{10}(\mathrm{O/H})) = -1.3\pm0.3$, which is consistent with H. B. Koziol et al. (2026) to within 2σ. We speculate that some differences could come from comparing a single galaxy to a trend pooled over many, and from our use of the 3.3 μm feature alone, since their 3.3 μm band also contains the 3.4 μm complex. Adding our 3.4 μm map to the numerator does flatten the slope, because the 3.4/3.3 μm ratio itself declines outward in M33, though we leave a more careful exploration of this to future work.

While these specific results are mostly validations of recent JWST results, the spectral information in LVF imaging from SPHEREx can enable better component separation (e.g. continuum subtraction, 3.3 μm versus 3.4 μm emission) compared to imaging with classical photometric filters, which may allow a resolution to open questions about the mechanism(s) driving the radial gradient. In addition, SPHEREx provides access to these 3.3 μm and 3.4 μm PAH emission maps all-sky without expensive targeted imaging/spectroscopy, opening up a wide variety of targets for study.

Beyond these astrophysical motivations for studying M33, the higher resolution and sensitivity of the 8 μm IRAC data make M33 a perfect testbed for (1) attempts at super-resolution, by forward modeling sub-PSF scale shifts between different exposures, (2) validating models of undetected point sources, and (3) validating differentiation of compact diffuse emission from point sources. This will be the subject of future work (see Section 7).

6. DISCUSSION

6.1. *Sensitivity*

One key property of a diffuse emission map is the sensitivity it achieves. We prefer an empirical estimate of this precision measured on the LVFit maps in order to compensate for any limitations in our uncertainty model, which we discuss further in Section 7. We use this sensitivity analysis to calibrate the LVFit per-pixel uncertainty maps delivered with the data products (Section 8).

We have to be careful about which pixels we select for an empirical sensitivity estimate, because each HEALPix $N_{\mathrm{side}} = 2^{15}$ pixel is fit with the 6.″44 FWHM Gaussian spatial weighting truncated at 5.″5, as described in Section 3.1. As a result, only pixels that are at least two pixels apart use a disjoint set of measurements. Assuming that the structure in our image varies slowly relative to the pixel scale, we can then measure the scatter between pairs of pixels separated by two pixels as an empirical estimate of the pixel-to-pixel uncertainty in the map. To account to first order for spatial variation in the map, we actually compute $d_r = I_r - (I_{r-2} + I_{r+2})/2$ for pixels along the HEALPix ring ordering index r and estimate the single pixel noise at index r as $\sigma_r = 1.4826\,\mathrm{MAD}(\mathrm{d_r})/\sqrt{1.5}$, where we use the median absolute deviation as an outlier-robust measure of spread, and convert to σ_r using its relation to σ for a unit normal distribution.

These per-pixel estimates can then be used to estimate uncertainties for averages on large spatial scales by roughly scaling by $\sqrt{N}$ for an average over N pixels. The estimator σ_r is noisy and so we average the σ_r over the parent HEALPix $N_{\mathrm{side}} = 2^{11}$ ($\sim 1.'7$) pixels, four levels up from the map grid. This gives approximately 230 samples per $N_{\mathrm{side}} = 2^{11}$ pixel, which we find yields an estimator of the noise with an uncertainty of $\sim 10\%$. This provides sufficient precision for us to fit analytic models to our noise estimates and assess the quality of those fits.

This estimate of the precision is conservative in the sense that astrophysical structure causing variation between pixels beyond a linear gradient will inflate our estimate. On the other hand, systematics in common between different detector pixels in the same exposure may be shared between neighboring pixels. This is one way in which the pixels we use are not completely independent, even though they are at least two pixels apart. However, given that most pixels have $\gtrsim 35$ exposures, we expect these systematics to predominantly enter as additional scatter between pixels covered by different combinations of exposures. Given the highly structured nature of Galactic emission on all measured scales, we expect true astrophysical structure to dominate the measured variation, thereby making our precision estimate conservative.

We want to establish how the per-pixel precision scales with the number of measurements used in our fits. For this, we will consider the NEP (Figure 15, left), one of the two deep fields where SPHEREx's observational strategy accumulates the largest number of exposures. We can compute σ_r across the field, dropping a varying number of exposures, keeping between 10% and 100% of the exposures in steps of 10% (Figure 15, top right). From this we can estimate that our $\sigma_r = \sigma_0 \times (N_{\mathrm{pts}}/100)^{-1/2}$ with $\sigma_0 = 2.00\times10^{-7}$ erg s^{-1} cm^{-2} sr^{-1}. However, we find that the scaling with N_{pts} is better fit with an uncertainty floor added in quadrature, for which we find $\sigma_0 = 1.82\times10^{-7}$ and $\sigma_f = 1.11\times10^{-8}$ erg s^{-1} cm^{-2} sr^{-1}.

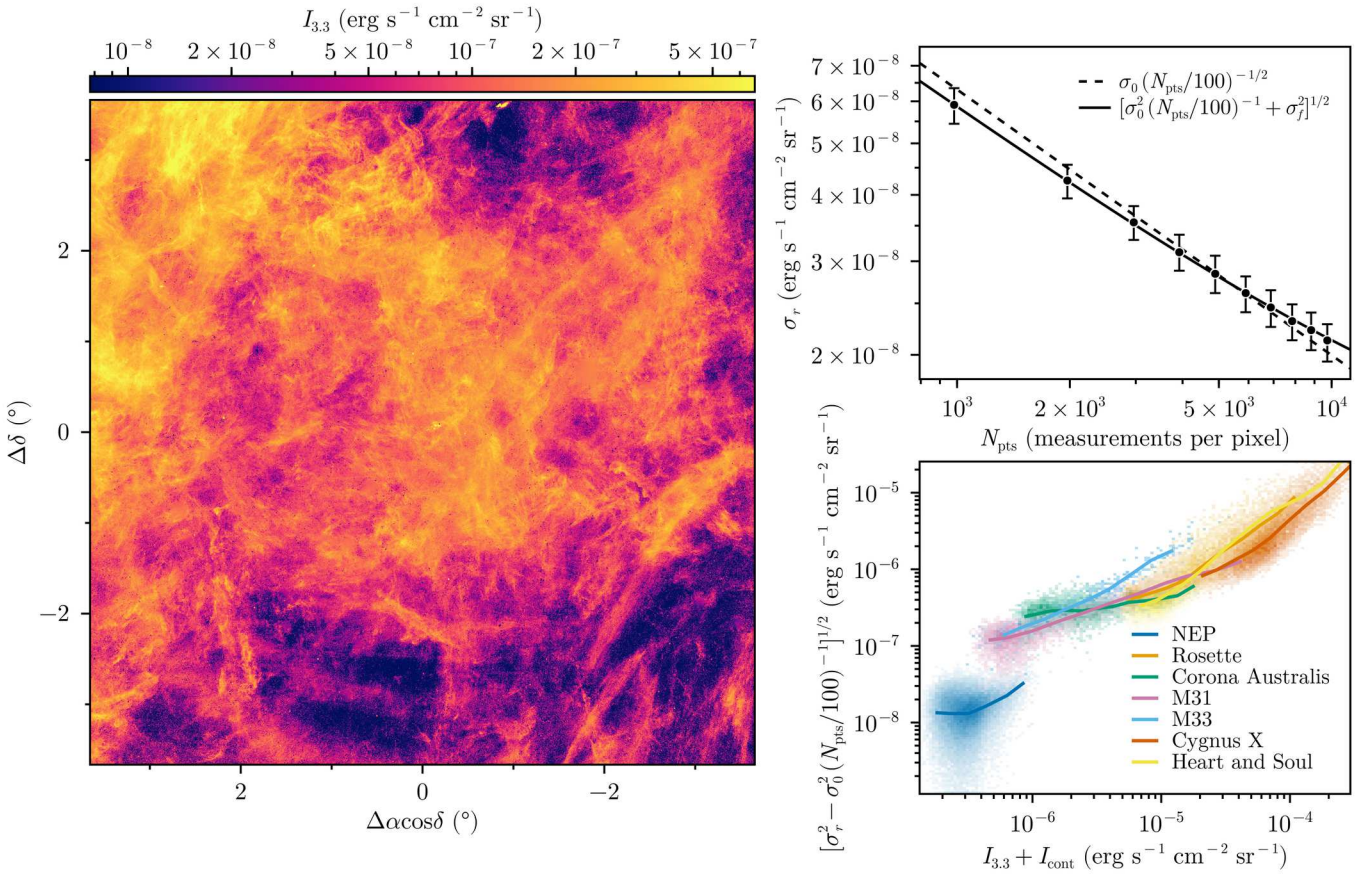

Figure 15. Left: A 7.°3 × 7.°3 field of the LVFit 3.3 μm PAH map centered on the north ecliptic pole in gnomonic offsets from (RA, Dec) = (270.°0, +66.°56) with east to the left. Top right: Average pixel noise σ_r in the central 1° × 1° of this field versus the median number of measurements per pixel $N_{\rm pts}$, while varying the number of exposures kept. The error bars show the interquartile range of σ_r averaged over parent $N_{\rm side} = 2^{11}$ pixels. The dashed line is pure $N_{\rm pts}^{-1/2}$ scaling and the solid line adds a floor in quadrature. Bottom right: The excess noise above the $N_{\rm pts}^{-1/2}$ scaling, using σ_0 from the solid-line fit versus the total 3.3 μm surface brightness $I_{\rm tot}$ in the fields shown in this work. The shading shows density of samples for each field (normalized by peak density of that field) and lines are running medians.

There are several possible origins of this apparent floor. The most likely stems from multiplicative errors in the forward model, such as in the component shapes or the spectral response of the LVF. To evaluate this, we can plot $\sqrt{\sigma_r^2 - \sigma_0^2 \times (N_{\rm pts}/100)^{-1}}$ versus the total surface brightness at 3.3 μm for a variety of our test fields (Figure 15, bottom right). This suggests it is necessary to add a term in quadrature that is a function of surface brightness. However, the fact that not all of the fields fall along the same line suggests that the multiplicative term is different for the different fields. Specifically, the NEP has much higher precision than the other fields. Understanding the origin of these differences, for example whether the NEP benefits from seeing a wider range of detector orientations or if astrophysical differences between regions drives variations in model-data mismatches, will be left to future work. For now, we fit the multiplicative term for each field (Table 2), adopting

$$\sigma_r^2 = \sigma_0^2 \left(\frac{N_{\rm pts}}{100}\right)^{-1} + (f\,I_{\rm tot})^2\,, \qquad (6)$$

with σ_0 and f both fit per field and $I_{\rm tot} = I_{3.3} + I_{\rm cont}$ the total surface brightness through the 3.3 μm feature, where $I_{\rm cont} = c_{3.29}\,\Delta\nu$ is the continuum surface brightness fit at 3.29 μm integrated over the feature width $\Delta\nu = c\sqrt{2\pi}\,\sigma_\lambda/\lambda^2$ ($\sigma_\lambda = 0.017\,\mu$m, so 1 MJy sr^{-1} of continuum corresponds to 1.19×10^{-5} erg s^{-1} cm^{-2} sr^{-1}). The fit σ_0 is 1.9×10^{-7} in the NEP, consistent with the depth test above, and 2.3–3.1 $\times 10^{-7}$ in the other fields, as expected from their brighter zodiacal background at lower ecliptic latitude. The fit values for f range from 0.03 (NEP) to 0.14 (M33). Away from the ecliptic poles, $N_{\rm pts}$ varies by less than a factor of 2 within a field, so σ_0 there is set by the noise of the faintest patches rather than by the $N_{\rm pts}$ dependence, which we only test in the NEP. In Cygnus X every patch is bright

enough that the surface-brightness term dominates and σ_0 is unconstrained.

To convert from $N_{\rm pts}$ used above to the number of exposures, we note that the number of points-per-exposure passing our wavelength and spatial cuts of our PAH model is approximately 2.3. With the current 16.5 months of SPHEREx imaging, pixels in regions at low ecliptic latitude typically have 35 exposures ($N_{\rm pts} \approx 80$), while the ecliptic poles can have two orders of magnitude more.

We can also compare our formal LVFit uncertainties $\sigma_{\rm formal}$ (from the Huber-weighted covariance of the fit) with the empirical σ_r measurements. We recalibrate the per-pixel uncertainties for each field with the same functional form,

$$\sigma_r^2 = (k\,\sigma_{\rm formal})^2 + (f'\,I_{\rm tot})^2\,, \tag{7}$$

fitting k and f' on the same $1.^\circ7$ patches (Table 2). Because $\sigma_{\rm formal}$ tracks the number, variances, and weights of the measurements that enter each pixel, it already carries most of the depth and some of the brightness dependence. A single scale factor k and a small surface-brightness term then describe σ_r with the same or lower scatter than Equation 6 in every field. In the NEP, Corona Australis, M31, and M33 the formal uncertainties are within 5–25% of the empirical noise ($k = 0.76$–0.94) and the recalibrated uncertainties reproduce σ_r with a scatter of 0.05–0.06 dex over patches, which is within the precision of our σ_r estimator. In the bright Galactic plane fields the formal uncertainties underestimate the scatter by a factor of 1.2–2.6 at low surface brightness, and the surface-brightness term dominates. In Cygnus X, the best fit sets $k = 0$ and the noise is simply $0.039\,I_{\rm tot}$ (as in Equation 6), with a scatter of 0.17 dex. A single global recalibration ($k = 0.88$, $f' = 0.038$) reproduces the NEP, Corona Australis, and M31 to within 6% but is biased by 0.1 dex for M33, so we report the per-field values and deliver the recalibrated per-pixel uncertainties with our data products. The same analysis of the tied hydrogen recombination-line amplitude (Section 3) gives formal uncertainties that are already within 20% of the empirical noise ($k = 1.0$–1.2), a surface-brightness term is not required there, and the fit f' values are consistent with zero. For the Paα map of the Rosette nebula (Section 4.1.1, fit on a separate detector) the formal uncertainties are within 10% of the noise at low surface brightness and the empirical noise again requires a surface-brightness term ($k = 1.08$, $f' = 0.024$), with a fraction similar to that of the 3.3 μm maps.

An important point of comparison is the depth of this SPHEREx PAH map compared to the WISE 12 μm PAH map from A. M. Meisner & D. P. Finkbeiner (2014). Trying to achieve comparable sensitivity appears difficult from the outset because the PAH emission in WISE W3 compared to the 3.3 μm PAH feature is $\sim 15:1$ in band-integrated power for standard diffuse-ISM PAH models (B. T. Draine & A. Li 2007b; B. T. Draine et al. 2021a), where $\mathrm{W3}/I_{3.3} \approx 6\times10^4$ MJy sr^{-1} per erg s^{-1} cm^{-2} sr^{-1}, and is larger still for PAH size distributions shifted to larger grains or for softer starlight spectra (1–3×10^5 from the B. T. Draine et al. (2021b) model library).

To convert the units between these two maps, we use the ratio actually measured on the test fields studied here, $\mathrm{W3}/I_{3.3} = 1.8\times10^5$ MJy sr^{-1} per erg s^{-1} cm^{-2} sr^{-1} (measured as a robust slope, with a field-to-field scatter of 0.1 dex), which lies within the model range for larger-than-standard PAHs. These were also the models preferred by JWST observations in nearby galaxies (K. M. Sandstrom et al. 2023). The slope fit is lowest in the Galactic plane fields and highest in the high-latitude cirrus, the direction expected for harder versus softer radiation fields, which could be explored more in combination with recent 3D interstellar radiation field maps (D. Lomeli et al. 2026; H. V. Bish et al. 2026) in future work. We then measured the empirical scatter as above on our PAH map, but smoothed to the $15''$ FWHM angular resolution of A. M. Meisner & D. P. Finkbeiner (2014), with both the WISE artifact mask and our infilling mask applied to both maps.

In the NEP, the WISE W3 precision is 3.6×10^{-3} MJy sr^{-1} corresponding to 1.7×10^{-8} erg s^{-1} cm^{-2} sr^{-1} in $I_{3.3}$. Converting the LVFit precision to the same spatial scale, we have 1.3×10^{-8} erg s^{-1} cm^{-2} sr^{-1} in $I_{3.3}$, which means the LVFit SPHEREx map is already a deeper PAH tracer at the ecliptic poles by a factor of 1.3. For the fainter fields away from the ecliptic poles, the W3 precisions are 8–15×10^{-3} MJy sr^{-1} corresponding to 4.5–8.0×10^{-8} erg s^{-1} cm^{-2} sr^{-1} in $I_{3.3}$, compared to the LVFit range in the same fields at the WISE spatial scale of 7.0–10×10^{-8} erg s^{-1} cm^{-2} sr^{-1} in $I_{3.3}$. Because the polar scan strategy of WISE (E. L. Wright et al. 2010) makes its coverage grow smoothly away from the ecliptic (from 10 exposures at $|\beta| < 20^\circ$ to 25 at $|\beta| = 30$–60° and 220 at the pole based on the A. M. Meisner & D. P. Finkbeiner (2014) coverage maps), whereas the SPHEREx coverage only increases above baseline within $\sim 10^\circ$ of the poles, the relative depth is not a monotonic function of ecliptic latitude.

At 16.5 months, the LVFit PAH map is already deeper than WISE at the ecliptic poles, within 15–35% of the WISE depth in the other low-latitude fields (Corona Australis and Rosette, $|\beta| < 20^\circ$), and 1.4–2 times noisier in the mid-latitude fields (M31, M33, and the Heart and Soul nebulae). Since the noise in these faint regions scales as $N_{\rm pts}^{-1/2}$, we project the depth by the end of the two-year prime survey (J. J. Bock et al. 2026) and after a two-year extension, based on the number of survey visits obtained so far versus the number expected. By the end of the prime survey, we estimate that the low-latitude fields will achieve parity with WISE or better (noise 0.8–1.0$\times$ that of WISE), while the mid-latitude fields will remain 1.2–1.7 times noisier. A two-year extension would bring most fields to parity (e.g. the Heart

Table 2. Per-field noise models. Rows are grouped by field and quantity being mapped, with the basis vectors of the corresponding fit listed in the group header. σ_r is the median empirical pixel noise and $N_{\rm pts}$ the median number of measurements per pixel. We show coefficients and residual scatter for fits to Equation 6 and Equation 7. σ_r and σ_0 are in units of 10^{-7} erg s^{-1} cm^{-2} sr^{-1}; f, f', and k are dimensionless. Values of f, f' below 0.01 are consistent with zero and values of residual scatter below 0.05 dex are consistent with our σ_r estimator uncertainty.

Field	σ_r	$N_{\rm pts}$	Equation 6			Equation 7		
	(10^{-7})	(median)	σ_0 (10^{-7})	f	scatter (dex)	k	f'	scatter (dex)
3.3 μm PAH map: 3.3+3.4 μm PAH + continuum[a]								
NEP	0.27	5515	1.92	0.029	0.053	0.871	0.029	0.049
Rosette	8.37	80	2.92	0.052	0.130	0.716	0.049	0.126
Corona Australis	3.54	82	3.09	0.036	0.066	0.944	0.018	0.051
M31	2.32	120	2.32	0.073	0.066	0.848	0.045	0.047
M33	3.82	150	2.36	0.141	0.061	0.761	0.133	0.061
Cygnus X	23.5	118	< 3[b]	0.039	0.169	< 0.01	0.039	0.169
Heart and Soul	5.04	129	2.34	0.043	0.104	0.689	0.039	0.100
Pfδ map: 3.3+3.4 μm PAH + continuum + Brackett/Pfund/Humphreys lines + nebular continuum								
Rosette	1.01	80	0.83	< 0.01	0.106	1.008	< 0.01	0.057
Heart and Soul	0.78	129	0.85	< 0.01	0.088	1.179	< 0.01	0.061
Paα map: Paα/Brδ/Brϵ + continuum								
Rosette	19.0	44	6.79	0.030	0.115	1.080	0.024	0.103

[a] + hydrogen recombination template for Rosette and Heart and Soul.

[b] The noise is flux-dominated throughout the field, and any $\sigma_0 \lesssim 3 \times 10^{-7}$ erg s^{-1} cm^{-2} sr^{-1} fits equally well.

and Soul nebulae) or within a few percent of it (e.g. M33), but the mid-latitude fields with the fewest visits (e.g. M31) would remain $\sim 25\%$ noisier than WISE, because the current survey strategy places a quarter of all exposures within 15° of the ecliptic poles. Spreading the exposures of an extension more evenly over the sky would raise the exposure count elsewhere by up to 35% (a 16% gain in depth), which would bring even those fields to within a few percent of WISE (e.g. M31 to within $\sim 6\%$) by the end of the extension.

6.2. *Reasons for Banding*

Working with LVF data presents a challenge for classical methods of photometric and spectroscopic analysis, but also an opportunity. Most methods that attempt to treat LVF data like classical photometry result in banding in the scanning direction, as shown in the top left of Figure 1. More generally, any mismodeling of LVF data will result in banding in the resultant images and so banding can be viewed as a clear indicator of model misspecification. Designing models that naturally do not produce unphysical banding in images resulting from LVF data reductions is necessary, but not sufficient, to have a well-specified model that describes the data.

Because more missions with LVF instruments are coming online (e.g. the PRobe far-Infrared Mission for Astrophysics, PRIMA; J. Glenn et al. 2025), it is important we codify some methodological lessons from developing SPHEREx reductions for these future surveys. In our experience, there are three main causes of banding: mismodeling either the instrument, a component you care about, or a component you do not care about. In the first case, the simplest example of this comes from roughly approximating the spatially varying spectral response of the detector, for example by defining discrete spectral channels. This is why we are careful to linearly interpolate the spectral response at the highest resolution possible. This means careful laboratory and on-sky calibration are critical for LVF instruments.

The second case is the one that most astronomers are trained to consider, for example by describing a PAH complex of interest by increasingly complicated substructure. We find that data-driven methods are critical in this case because we aim to discover unknown unknowns about the features of interest from the data. Being data-driven also allows us to defer questions of absolute spectral shape by making inferences only in the

spectral-response-convolved space, which is the one that is actually observed.

The third case is the most nefarious, because it results from a component that we might rather just subtract away and forget about. Unlike classical single-band photometry where you might have to subtract off a constant background or a low-order polynomial background to handle a nuisance component, in LVF data, one has to understand both the spatial and spectral properties of the nuisance component in detail.

In all of these cases, to ignore banding is to ignore a misspecified model that for LVF data results not only in a bias, but a spatially varying bias. These biases can be so large as to obscure even the spatial structure of the features of interest (see Section 6.3), but even when significantly suppressed may also impart more subtle biases that impact studies of the diffuse emission correlation function, such as for quantitative studies of interstellar turbulence (e.g., M.-A. Miville-Deschênes et al. 2016; G. Vigoureux et al. 2026). As a result, it is essential to improve our forward models on all of the above axes. Additionally, it should be an extremely high priority for the community to be producing and consuming estimates of the correlated uncertainties that account for any residual banding.

6.3. *Comparison to Literature*

There have been several previous works that make PAH maps from SPHEREx and we briefly compare LVFit with those works.

- A. J. Cukierman et al. (2026) produce maps up to the native 6.″15 pixel scale of SPHEREx (H. Hui et al. 2026; SPHEREx Science Data Center 2026) for many test fields using a channel-binning approach, which is improved by interpolation between the bins to reduce banding.
- G. Murgia et al. (2026) follow the method of A. J. Cukierman et al. (2026), but smooth to 20′ FWHM to reduce remaining banding effects and study large scale correlations.
- J. L. Hora et al. (2026) use channel maps, but apply a "correction factor" that weights the channels by a fixed spectral shape assumed for the PAH. While not applied in the forward model direction, this should be equivalent to LVFit with a single, non-data-driven component, though it does not give access to the residuals and χ^2 of the fit.
- C. Boersma et al. (2026) fit analytic continuum and Gaussian models of the PAH peaks integrated over top-hat approximations of the spectral response at 7.′2 apertures on a 7×7 grid.
- C. Li (2026) make channel maps of the LMC/SMC at 2′.
- R. Cepeda-Arroita et al. (2026) forward model HEALPix $N_{\rm side} = 2^{12}$ ($\sim 0.'9$) pixels with a 2.′6 FWHM beam, using analytic Drude, Gaussian, and continuum models with a nonlinear fit to channel-binned measurements using per-channel spectral response functions.

The main differences in this work are that we forward model the data by the highest spatial resolution representation of variation in the spectral response of the detector, maintain the full spatial resolution of SPHEREx, and use PAH lineshape models that are initialized from observations at higher spectral resolution with learned data-driven corrections. The combination of these choices naturally produces images without apparent banding.

We compare the LVFit and R. Cepeda-Arroita et al. (2026) maps of the Rosette nebula in Figure 16, because their approach is the most similar to our own. On large scales (2.′6) the maps agree to $\sim 4\%$ for the summed $3.3 + 3.4\,\mu$m surface brightness and $\sim 1\%$ for the $3.4/3.3\,\mu$m ratio. However, residual banding along the SPHEREx scan direction is visible both subtly in the PAH surface brightness map and clearly in the $3.4/3.3\,\mu$m ratio map in the R. Cepeda-Arroita et al. (2026) map (Figure 16, top row). The banding is not visible in either the surface brightness or ratio map using LVFit (Figure 16, bottom row). While we only show the Rosette nebula, this general trend of LVFit suppressing banding in both the surface brightness and ratio maps holds for all of the regions we have examined.

7. CONCLUSIONS

In this work, we have demonstrated LVFit, a linear, data-driven, forward-modeling approach to SPHEREx LVF data that provides access to maps of diffuse emission close to the native pixel scale of SPHEREx (6.″4 pixels, 8.″5 FWHM angular resolution). This approach allowed us to cleanly separate overlapping components of continuum, $3.3\,\mu$m PAH, $3.4\,\mu$m PAH, and hydrogen recombination line emission, without the significant banding artifacts that plague other literature SPHEREx reductions. We applied our method to make maps of a wide variety of testbeds that stress test angular resolution, spectral complexity, sensitivity, and spatial deblending. We validate our results against external datasets tracing PAHs, ionized gas, and neutral H gas from WISE, Spitzer, MDW, and LGLBS, highlighting synergies with those datasets as well as upcoming data releases from the SDSS-V LVM. While we only show maps of $3.3\,\mu$m PAH, $3.4\,\mu$m PAH, and Paschen-α emission for specific regions, the methodology presented here can be straightforwardly extended to other emission features in the SPHEREx wavelength range and is scalable to the entire sky (in $\sim$ 60k core-h).

However, there are several limitations of the methodology in this work that impact the quality of the fi-

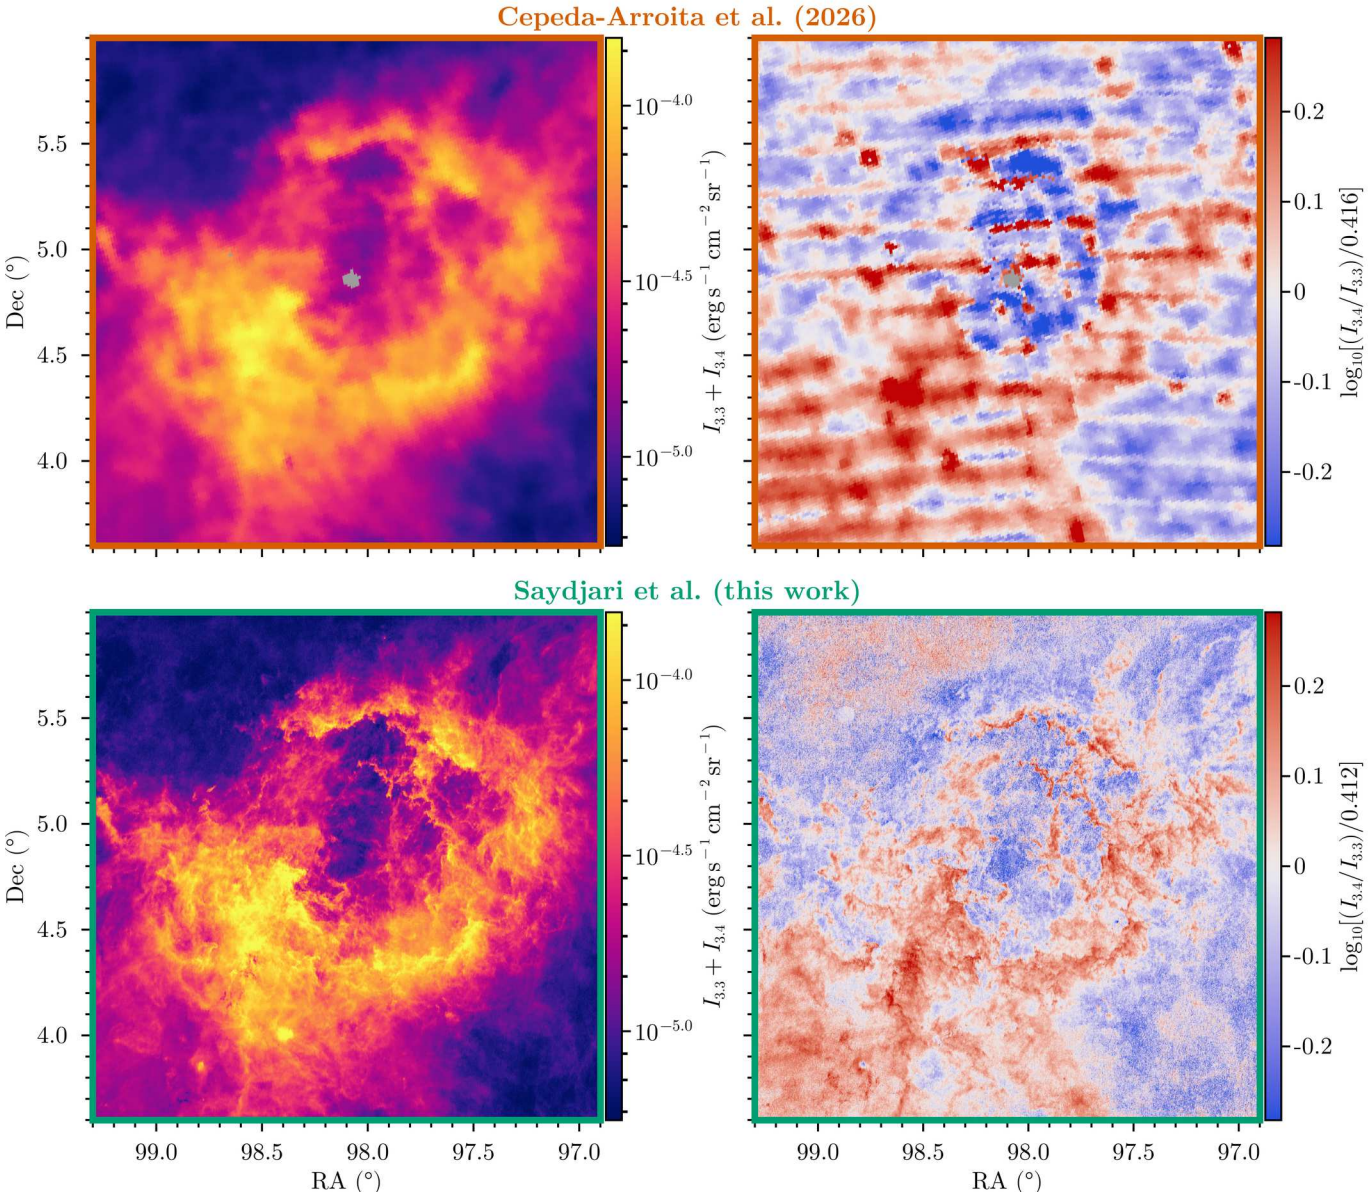

Figure 16. Comparison of the R. Cepeda-Arroita et al. (2026) SPHEREx PAH maps (top) with the LVFit maps produced in this work (bottom) for the Rosette nebula. Left: Summed $3.3 + 3.4\,\mu$m line surface brightness on a common logarithmic color scale. Right: $\log_{10}[(I_{3.4}/I_{3.3})/\text{median}]$, each normalized by its own median over the panel (0.416 and 0.412), on a common ± 0.28 dex scale. The R. Cepeda-Arroita et al. (2026) maps are their unfiltered products at native resolution (HEALPix $N_{\rm side} = 2^{12}$, 2.′6 FWHM beam) and the LVFit maps are shown at native resolution ($N_{\rm side} = 2^{15}$ pixels, 8.″5 FWHM beam).

nal map. Future work would benefit from the following methodological advances.

Model PSF—Here we model the SPHEREx data at the pixel level, ignoring any correlations between measurements introduced by the point-spread function. This significantly impacts our ability to model point sources (see below) and limits our ability to resolve substructures in the diffuse emission smaller than the PSF scale, which is in principle possible via a forward model of different exposures with a variety of subpixel shifts. Furthermore, there are real variations in the point-spread function over the detector plane. Future work could explicitly learn and forward model the spatially varying point spread function to recover more and higher-resolution PAH emission with fewer star artifacts as a result.

Model point sources—We do not explicitly model point sources. We mask and conditionally infill over them. By jointly forward modeling the diffuse emission and the point sources, we expect to better recover structures in the diffuse emission that are coincident with point sources compared to our only statistically accurate infills. Improved PSF models (see above) and wider wave-

length coverage should facilitate better point source "removal."

Fully data-driven basis set—Here we take a practical approach to determining a fixed set of basis vectors. We use two PAH basis vectors learned as data-driven perturbations from high-resolution JWST spectra, three continuum components normalized to unity at 3.3 μm (constant, λ^2, and λ^{-2}), and one theoretically generated basis vector for hydrogen recombination lines in regions with strong ionized gas emission.

Given the large number of SPHEREx measurements available, it should be possible to learn the line shapes directly from the data, including additional types of line variation and components that are not anticipated in advance. These could potentially be estimated in a preprocessing step but would ideally be learned jointly with the maps, reducing the possibility of model misspecification. One difficulty we have run into, when adding rovibrational H_2 emission line basis vectors, is that while their amplitudes may be easy to fit in regions of high signal, it may be necessary to regularize the amplitude of rarer basis vectors so that they do not bias measurements of the other components in regions of low signal-to-noise. This regularization fits naturally into the variational inference approach we describe below.

Data-driven exposure corrections—There are several poorly constrained components in SPHEREx observations that are expected to vary at the exposure level, which are difficult to measure and correct when working at the pixel level and might require alternating passes between pixel-level and exposure-level fit parameters to resolve. The two most obvious examples are "shuttle glow," emission from the spacecraft interacting with atomic oxygen in the atmosphere that is strongest when the pointing has a small angle relative to the direction of motion (SPHEREx Science Data Center 2026), and zodiacal light from scattered sunlight and thermal emission of interplanetary dust.

For zodiacal light, we rely on a model provided by the SPHEREx collaboration based on T. Kelsall et al. (1998) to perform a point-estimate subtraction of zodiacal light from each image before fitting pixel values. However, we also expect deviations from the model which cannot be fully captured by static-sky continuum subtraction as they are determined by orientation of the telescope pointing relative to solar system objects (e.g. Jupiter, Mars) based on the Earth's orbital phase and are not the same in all exposures which point to a given location of the sky. Failure to properly model the zodiacal light is one form of model misspecification that we suspect makes banding at low ecliptic latitudes so difficult to eliminate, because it is where the zodiacal light is strongest. To solve this problem, we are pursuing data-driven models of the zodiacal light or shuttle glow that are constrained to depend only on the differences between exposures that we expect to be physically and causally linked to those signals (e.g., angle relative to spacecraft motion).

Time dependence—We are modeling the astrophysical sky as if it were perfectly static. However, as we discuss above, moving solar system objects violate that assumption and had to be masked. Similarly, parallax and proper motion of point sources cannot be modeled in LVFit. Brightness variations from transient events (e.g., supernovae) or periodic variable stars are also outside of our framework, though they are masked by our current approach that disregards point sources. However, any approach that attempts to model the point sources will have to contend with variable stars. For example, our preliminary efforts in this direction have struggled to use a static model to describe a bright Mira variable in our Rosette nebula footprint. In addition, diffuse structures may have observable proper motions if they are very nearby (e.g., M. Piecka et al. 2025) or brightness variations (e.g., supernova light echoes; E. Dwek & R. G. Arendt 2008), which, if left unmodeled, could decrease the quality of our model fit to the data.

Spatial correlations—Our work treats data-constrained model pixels independently, averaging over a small Gaussian beam. This procedure does not acknowledge the actual power-law spectrum of the diffuse emission, which exhibits more correlation at larger scales and less correlation at scales below the beam size. Proper treatment of these correlations could produce maps with simultaneously less noise, fewer artifacts, higher resolution, and improved separation of diffuse emission into components which differ in their spatial structure. One promising option is to use a Gaussian process prior with an appropriate covariance kernel for the entire sky, similar to our approach for the small number of unconstrained masked pixels in this work. One challenge with this approach is representing such a prior over large areas of sky at high resolution, but this is feasible with recently developed methods such as GraphGP (B. Dodge et al. 2026).

Uncertainty quantification—Our method is a simple conditional (on the basis vectors we learned) linear optimization and the formal uncertainties on the fit are only close to the truth in the faint-sky limit. We had to empirically calibrate our reported uncertainties in Section 6.1 for the bright-sky limit to account for model errors and the fact that we ignore correlated uncertainties. A full posterior that accounts for the covariance between map pixels and yields internally consistent and well-calibrated errors would be valuable for differentiating real structure from noise, propagating uncertainty into subsequent measurements made from these maps, and trusting that systematics and the instrument response are completely understood.

Future directions—Many of the aforementioned limitations can in principle be solved by adding additional hyperparameters to the model (e.g., PSF, basis vectors, stellar spectra, zodi), choosing appropriate priors

(e.g., Gaussian processes), and drawing correlated samples from the joint posterior. This problem is no longer linear in its parameters nor can it be solved independently for each of its billions of model pixels, presenting a major challenge for many techniques. Fortunately, algorithms such as Metric Gaussian Variational Inference (J. Knollmüller & T. A. Enßlin 2019) and Geometric Variational Inference (P. Frank et al. 2021) have been developed to solve such Bayesian inverse problems at scale with acceptable computational cost. These methods have been applied successfully to 4D maps with three spatial dimensions and one scalar value previously (Frank & Saydjari et al. 2026, in prep.), and we are working on a joint spatial-spectral application to SPHEREx (Dodge et al. 2026, in prep.).

8. DATA AVAILABILITY

We release the maps of every field in this paper in two forms. The primary products are HDF5 tables on the HEALPix $N_{\rm side} = 2^{15}$ grid. For convenience, we also provide FITS rasters of the main maps on a $3\farcs22$ grid in the frame of each field. The FITS files hold one image extension for each primary map product. These are the infilled 3.3 μm PAH, 3.4 μm PAH, and hydrogen recombination line maps, the latter only on fields where it was fit (Section 3.6). These maps include both the mean and two conditional draws of the infill along with their calibrated uncertainties. The files also include the continuum surface brightness at 3.29 μm, the reduced χ^2, the number of measurements per pixel, and the mask. The HDF5 tables carry additional fields that completely describe the fit. They give the amplitude of every basis vector and its formal uncertainty, the fit line values before infilling, the pixel coordinates, and the list of exposures that contributed to each field. For masked pixels, we include a bitmask indicating why they were masked (Section 3.5): bit 0 for insufficient coverage or a failed fit, bit 1 for a star disk, bit 2 for a residual outlier, bit 3 for a residual outlier near a star, bit 4 for the ratio-based growth of a bright star's disk, and bit 5 for pixels added by the final closure step.

All of these data products, the PAH basis vectors, the spectral response model, our list of detector pixels with anomalous residuals (Section 3.1), the third-party maps used in our comparisons, and the code necessary to produce the maps from the SPHEREx images and every figure and number in this paper from the maps are provided in a Zenodo repository at doi:10.5281/zenodo.24068379. The SPHEREx images themselves are not redistributed, but the archive includes the measurement-level inputs of the two figures made from individual measurements (Figures 3 and 4). The maps were made from the Quick Release 2 and 3 spectral images, which are available from the NASA/IPAC Infrared Science Archive (IRSA)[12] (SPHEREx Science Data Center 2026). The total Zenodo archive is ∼45 GB unzipped. The figures in this article are only provided at 300 dpi, but the full-resolution versions of every figure are also included in the repository.

[12] https://irsa.ipac.caltech.edu/Missions/spherex.html

ACKNOWLEDGMENTS

A.K.S. acknowledges that support for this work was provided by NASA through the NASA Hubble Fellowship grant HST-HF2-51564.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. C.H. is supported by the National Science Foundation Graduate Research Fellowship under Grant Number DGE-2234660. B.D. is supported by the National Science Foundation Graduate Research Fellowship under grant number DGE-2146755. S.E.C. acknowledges support from the U.S. National Science Foundation under NSF CAREER grant No. AST-2441452, and from an Alfred P. Sloan Research Fellowship. The Flatiron Institute is a division of the Simons Foundation.

We acknowledge helpful discussions with Matt Ashby, Chris Clark, Kinsey Cronin, Julianne Dalcanton, Bruce Draine, Doug Finkbeiner, Gregory Green, Tom Hilder, Joe Hora, Howard Hui, Adam Leroy, Claire Murray, Hans-Walter Rix, Karin Sandstrom, Eddie Schlafly, and Liz Tarantino. A.K.S. acknowledges Sophia Sánchez-Maes for helpful discussions and much support. This work was supported in part by the Carnegie Theoretical Astrophysics Center (CTAC) through its Theory Scholar program. The authors acknowledge the Interstellar Institute and the Institut Pascal at Université Paris-Saclay for hosting discussions that nourished the development of the ideas behind this work. This work resulted from collaborations supported by the Kavli Institute for Particle Astrophysics and Cosmology through a Workshop Program grant to Philipp Frank for the Mapping the Diffuse Milky Way workshop held at Stanford on October 20–24, 2025, made possible by an endowment by the Kavli Foundation.

Above all, it is a pleasure to thank the NASA SPHEREx team for generously making these data public on a fast timescale.

AUTHOR CONTRIBUTIONS

A.K.S. co-conceived of the project, developed the methods, wrote and ran the code, produced the figures, co-ideated testbeds, and co-wrote the first draft of the manuscript. C.H. developed the database key to the speed of our method, managed data ingestion, and provided feedback on the manuscript. B.D. co-conceived of the project, contributed to understanding banding, tested alternative methods, and wrote parts of the manuscript. P.F. helped develop alterative methods to

overcome limitations of LVFit and provided feedback over the course of the project. D.W.H. co-conceived of the project, provided feedback during method development, and wrote parts of the manuscript. S.E.C. co-ideated testbeds, provided feedback over the course of the project, and co-wrote the first draft of the manuscript. All authors have read and agree to the final manuscript.

GENERATIVE AI DISCLOSURE

The original concept, methods, code base, and first draft of the manuscript were written by the authors. We used a Claude Code agent (Anthropic; Fable 5.1 and Opus 5.5) to refactor and automate the code base, execute and monitor some of the batch fits and calibration tests, polish and write some figure-production scripts, draft and edit some passages of the text at our direction, audit some of the methods description and numbers against the code and data products, and locate and format some bibliography entries. All code, text, figures, and references were reviewed by the authors, and all text was edited by them. No AI tool generated data, images, or scientific results beyond running the authors' pipeline, which is fully available as described in Section 8. The authors take full responsibility for the content of this work.

Facilities: SPHEREx, IRSA, WISE, Spitzer (IRAC), VLA, GBT, FLWO:2MASS, CTIO:2MASS

Software: Julia (J. Bezanson et al. 2017); Makie.jl (S. Danisch & J. Krumbiegel 2021); Healpix.jl (M. Tomasi & Z. Li 2021) and HEALPix (K. M. Górski et al. 2005); DataFrames.jl (M. Bouchet-Valat & B. Kamiński 2023); DuckDB (M. Raasveldt & H. Mühleisen 2019) and Apache Parquet (Apache Software Foundation 2026); HDF5.jl with HDF5 (The HDF Group 2025); JLD2.jl; FITSIO.jl with cfitsio (W. D. Pence 2010); WCS.jl with WCSLIB (E. W. Greisen & M. R. Calabretta 2002; M. R. Calabretta & E. W. Greisen 2002); SkyCoords.jl; SPICE.jl with the SPICE toolkit (C. H. Acton 1996; C. Acton et al. 2018); StatsBase.jl (D. Lin et al. 2026); Interpolations.jl (M. Kittisopikul et al. 2022); LsqFit.jl; CSV.jl (J. Quinn et al. 2026a); JSON3.jl (J. Quinn et al. 2025); HTTP.jl (J. Quinn et al. 2026b); CodecZlib.jl; StructArrays.jl; LaTeXStrings.jl; ColorSchemes.jl with colorcet colour maps (P. Kovesi 2015); NumPy (C. R. Harris et al. 2020); PyNeb (V. Luridiana et al. 2015)

REFERENCES

Acton, C., Bachman, N., Semenov, B., & Wright, E. 2018, A look towards the future in the handling of space science mission geometry, Planet. Space Sci., 150, 9, doi: 10.1016/j.pss.2017.02.013

Acton, C. H. 1996, Ancillary data services of NASA's Navigation and Ancillary Information Facility, Planet. Space Sci., 44, 65, doi: 10.1016/0032-0633(95)00107-7

Aftab, N., Zhang, X., Walker, S., et al. 2025, The MDW Hα Sky Survey: Data Release 1, arXiv e-prints, arXiv:2510.22900. https://arxiv.org/abs/2510.22900

Akeson, R., Dubois-Felsmann, G. P., Crill, B. P., et al. 2026, The SPHEREx Image and Spectrophotometry Processing Pipeline, ApJS, 285, 67, doi: 10.3847/1538-4365/ae8843

Allamandola, L. J., Tielens, A. G. G. M., & Barker, J. R. 1985, Polycyclic aromatic hydrocarbons and the unidentified infrared emission bands: auto exhaust along the milky way., ApJL, 290, L25, doi: 10.1086/184435

André, P., Men'shchikov, A., Bontemps, S., et al. 2010, From filamentary clouds to prestellar cores to the stellar IMF: Initial highlights from the Herschel Gould Belt Survey, A&A, 518, L102, doi: 10.1051/0004-6361/201014666

Apache Software Foundation. 2026, Apache Parquet,, https://parquet.apache.org

Barker, J. R., Allamandola, L. J., & Tielens, A. G. G. M. 1987, Anharmonicity and the Interstellar Polycyclic Aromatic Hydrocarbon Infrared Emission Spectrum, ApJL, 315, L61, doi: 10.1086/184861

Baron, D., Sandstrom, K. M., Sutter, J., et al. 2025, PHANGS-ML: The Universal Relation between PAH Band and Optical Line Ratios across Nearby Star-forming Galaxies, ApJ, 978, 135, doi: 10.3847/1538-4357/ad972a

Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results, ApJS, 208, 20, doi: 10.1088/0067-0049/208/2/20

Berné, O., Habart, É., Peeters, E., et al. 2022, PDRs4All: A JWST Early Release Science Program on Radiative Feedback from Massive Stars, PASP, 134, 054301, doi: 10.1088/1538-3873/ac604c

Bertero, M., & Boccacci, P. 1998, Introduction toa Inverse Problems in Imaging (Bristol and Philadelphia: Institute of Physics Publishing)

Bezanson, J., Edelman, A., Karpinski, S., & Shah, V. B. 2017, Julia: A Fresh Approach to Numerical Computing, SIAM Review, 59, 65, doi: 10.1137/141000671

Bish, H. V., Peek, J. E. G., Zucker, C., et al. 2026, LightCube: A Parsec-Resolution 3D Model of the Local Far-Ultraviolet Interstellar Radiation Field, arXiv e-prints, arXiv:2608.27389, doi: 10.48550/arXiv.2608.27389

Blitz, L., & Thaddeus, P. 1980, Giant molecular complexes and OB associations. I. The Rosette molecular complex, ApJ, 241, 676, doi: 10.1086/158379

Bock, J. J., Aboobaker, A. M., Adamo, J., et al. 2026, The SPHEREx Satellite Mission, ApJ, 999, 139, doi: 10.3847/1538-4357/ae2be2

Boersma, C., Maragkoudakis, A., Allamandola, L. J., et al. 2026, SPHEREx: Aromatics, Aliphatics, and PAH Size across the Iris Nebula, ApJ, 997, 239, doi: 10.3847/1538-4357/ae2a33

Bouchet-Valat, M., & Kamiński, B. 2023, DataFrames.jl: Flexible and Fast Tabular Data in Julia, Journal of Statistical Software, 107, 1, doi: 10.18637/jss.v107.i04

Calabretta, M. R., & Greisen, E. W. 2002, Representations of celestial coordinates in FITS, A&A, 395, 1077, doi: 10.1051/0004-6361:20021327

Cardelli, J. A., Clayton, G. C., & Mathis, J. S. 1989, The Relationship between Infrared, Optical, and Ultraviolet Extinction, ApJ, 345, 245, doi: 10.1086/167900

Cepeda-Arroita, R., Dickinson, C., Rubiño-Martín, J. A., et al. 2026, A Full-Sky SPHEREx Map of Integrated 3.3 and 3.4 μm PAH Emission, arXiv e-prints, arXiv:2609.37878, doi: 10.48550/arXiv.2609.37878

Córdova Rosado, R., Hensley, B. S., Clark, S. E., et al. 2024, The Atacama Cosmology Telescope: Galactic Dust Structure and the Cosmic PAH Background in Cross-correlation with WISE, ApJ, 960, 96, doi: 10.3847/1538-4357/ad05cd

Cramér, H. 1946, Mathematical Methods of Statistics, Princeton Mathematical Series No. 9 (Princeton, NJ: Princeton University Press)

Cukierman, A. J., Chen, S.-S., Kang, J. H., et al. 2026, Spectral Map Making with SPHEREx, ApJS, 286, 22, doi: 10.3847/1538-4365/ae91df

Danisch, S., & Krumbiegel, J. 2021, Makie.jl: Flexible high-performance data visualization for Julia, Journal of Open Source Software, 6, 3349, doi: 10.21105/joss.03349

De Looze, I., Baes, M., Parkin, T. J., et al. 2012, Herschel and JCMT observations of the early-type dwarf galaxy NGC 205, MNRAS, 423, 2359, doi: 10.1111/j.1365-2966.2012.21044.x

Deans, S. R. 1983, The Radon Transform and Some of Its Applications (New York: John Wiley & Sons)

Dodge, B., Frank, P., & Clark, S. E. 2026, GraphGP: Scalable Gaussian Processes with Vecchia's Approximation, arXiv e-prints, arXiv:2606.11402, doi: 10.48550/arXiv.2606.11402

Doi, M., Tanaka, M., Fukugita, M., et al. 2010, Photometric Response Functions of the Sloan Digital Sky Survey Imager, AJ, 139, 1628, doi: 10.1088/0004-6256/139/4/1628

Draine, B. T., & Li, A. 2007a, Infrared Emission from Interstellar Dust. IV. The Silicate-Graphite-PAH Model in the Post-Spitzer Era, ApJ, 657, 810, doi: 10.1086/511055

Draine, B. T., & Li, A. 2007b, Infrared Emission from Interstellar Dust. IV. The Silicate-Graphite-PAH Model in the Post-Spitzer Era, ApJ, 657, 810, doi: 10.1086/511055

Draine, B. T., Li, A., Hensley, B. S., et al. 2021a, Excitation of Polycyclic Aromatic Hydrocarbon Emission: Dependence on Size Distribution, Ionization, and Starlight Spectrum and Intensity, ApJ, 917, 3, doi: 10.3847/1538-4357/abff51

Draine, B. T., Li, A., Hensley, B. S., et al. 2021b, PAH emission: Dependence on Starlight Spectrum, Intensity, PAH Size Distribution, and PAH Ionization: A Library of Model Emission Spectra,, Harvard Dataverse dataset doi: 10.7910/DVN/LPUHIQ

Draine, B. T., Aniano, G., Krause, O., et al. 2014, Andromeda's Dust, ApJ, 780, 172, doi: 10.1088/0004-637X/780/2/172

Dwek, E., & Arendt, R. G. 2008, Infrared Echoes Reveal the Shock Breakout of the Cas A Supernova, ApJ, 685, 976, doi: 10.1086/589988

Edenhofer, G., Zucker, C., Frank, P., et al. 2024, A parsec-scale Galactic 3D dust map out to 1.25 kpc from the Sun, A&A, 685, A82, doi: 10.1051/0004-6361/202347628

Fazio, G. G., Hora, J. L., Allen, L. E., et al. 2004, The Infrared Array Camera (IRAC) for the Spitzer Space Telescope, ApJS, 154, 10, doi: 10.1086/422843

Frank, P., Leike, R., & Enßlin, T. A. 2021, Geometric Variational Inference, Entropy, 23, 853, doi: 10.3390/e23070853

Fruchter, A. S., & Hook, R. N. 2002, Drizzle: A Method for the Linear Reconstruction of Undersampled Images, PASP, 114, 144–152, doi: 10.1086/338393

Glenn, J., Meixner, M., Bradford, C. M., et al. 2025, PRIMA mission concept, Journal of Astronomical Telescopes, Instruments, and Systems, 11, 031628, doi: 10.1117/1.JATIS.11.3.031628

Gonçalves, D. R., Magrini, L., Teodorescu, A. M., & Carneiro, C. M. 2014, Planetary nebulae: the universal mass-metallicity relation for Local Group dwarf galaxies and the chemistry of NGC 205, MNRAS, 444, 1705, doi: 10.1093/mnras/stu1464

Gordon, K. D., Bailin, J., Engelbracht, C. W., et al. 2006, Spitzer MIPS Infrared Imaging of M31: Further Evidence for a Spiral-Ring Composite Structure, ApJL, 638, L87, doi: 10.1086/501046

Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere, ApJ, 622, 759, doi: 10.1086/427976

Green, P. J. 1984, Iteratively Reweighted Least Squares for Maximum Likelihood Estimation, and some Robust and Resistant Alternatives, Journal of the Royal Statistical Society: Series B (Methodological), 46, 149, doi: 10.1111/j.2517-6161.1984.tb01288.x

Greisen, E. W., & Calabretta, M. R. 2002, Representations of world coordinates in FITS, A&A, 395, 1061, doi: 10.1051/0004-6361:20021326

Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, Array programming with NumPy, Nature, 585, 357, doi: 10.1038/s41586-020-2649-2

Helou, G., Roussel, H., Appleton, P., et al. 2004, The Anatomy of Star Formation in NGC 300, ApJS, 154, 253, doi: 10.1086/422640

HGBS Team. 2020, Herschel Gould Belt Survey,, NASA IPAC DataSet, IRSA72 doi: 10.26131/IRSA72

Hogg, D. W., & Casey, A. R. 2024, Frizzle: Combining spectra or images by forward modeling, arXiv e-prints, arXiv:2403.11011, doi: 10.48550/arXiv.2403.11011

Holland, P. W., & Welsch, R. E. 1977, Robust regression using iteratively reweighted least-squares, Communications in Statistics - Theory and Methods, 6, 813, doi: 10.1080/03610927708827533

Hora, J. L., Noh, J. K., Melnick, G. J., et al. 2026, SPHEREx Widefield Infrared Spectral Mapping of Interstellar Ices and Polycyclic Aromatic Hydrocarbons, ApJ, 1001, 165, doi: 10.3847/1538-4357/ae5180

Huber, P. J. 1964, Robust Estimation of a Location Parameter, The Annals of Mathematical Statistics, 35, 73, doi: 10.1214/aoms/1177703732

Hui, H., Bock, J. J., Condon, S., et al. 2026, Spectral Response of SPHEREx, ApJS, 284, 10. https://arxiv.org/abs/2602.09139

Hummer, D. G., & Storey, P. J. 1987, Recombination-line intensities for hydrogenic ions - I. Case B calculations for H I and He II, MNRAS, 224, 801, doi: 10.1093/mnras/224.3.801

Jarrett, T. H., Cluver, M. E., Brown, M. J. I., et al. 2019, The WISE Extended Source Catalog (WXSC). I. The 100 Largest Galaxies, ApJS, 245, 25, doi: 10.3847/1538-4365/ab521a

Joblin, C., Tielens, A. G. G. M., Allamandola, L. J., & Geballe, T. R. 1996, Spatial Variation of the 3.29 and 3.40 Micron Emission Bands within Reflection Nebulae and the Photochemical Evolution of Methylated Polycyclic Aromatic Hydrocarbons, ApJ, 458, 610, doi: 10.1086/176843

Kelsall, T., Weiland, J. L., Franz, B. A., et al. 1998, The COBE Diffuse Infrared Background Experiment Search for the Cosmic Infrared Background. II. Model of the Interplanetary Dust Cloud, ApJ, 508, 44, doi: 10.1086/306380

Kittisopikul, M., Holy, T. E., & Aschan, T. 2022, Interpolations.jl, Zenodo, doi: 10.5281/zenodo.593143

Knollmüller, J., & Enßlin, T. A. 2019, Metric Gaussian Variational Inference, arXiv e-prints, arXiv:1901.11033, doi: 10.48550/arXiv.1901.11033

Koch, E. W., Leroy, A. K., Rosolowsky, E. W., et al. 2025, The Karl G. Jansky Very Large Array Local Group L-band Survey (LGLBS), ApJS, 279, 35, doi: 10.3847/1538-4365/ade0ad

Korngut, P. M., Bock, J. J., Akeson, R., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10698, Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave, 106981U, doi: 10.1117/12.2312860

Kovesi, P. 2015, Good Colour Maps: How to Design Them, arXiv e-prints, arXiv:1509.03700, doi: 10.48550/arXiv.1509.03700

Koziol, H. B., Sandstrom, K., Baron, D., et al. 2026, Variations in the 3.3 μm Polycyclic Aromatic Hydrocarbon Feature Across Nearby Galaxies Driven by Metallicity and Radiation Field Spectrum, arXiv e-prints, arXiv:2608.05286, doi: 10.48550/arXiv.2608.05286

Lang, D. 2014, unWISE: Unblurred Coadds of the WISE Imaging, AJ, 147, 108, doi: 10.1088/0004-6256/147/5/108

Lewis, A. R., Dolphin, A. E., Dalcanton, J. J., et al. 2015, The Panchromatic Hubble Andromeda Treasury. XI. The Spatially Resolved Recent Star Formation History of M31, ApJ, 805, 183, doi: 10.1088/0004-637X/805/2/183

Li, C. 2026, SPHEREx 3.3 Micron Aromatic Emission in the Magellanic Clouds: Separating Dust Column, Excitation, and Environmental Suppression, arXiv e-prints, arXiv:2609.20295, doi: 10.48550/arXiv.2609.20295

Li, T. S., DePoy, D. L., Marshall, J. L., et al. 2016, Assessment of Systematic Chromatic Errors that Impact Sub-1% Photometric Precision in Large-area Sky Surveys, AJ, 151, 157, doi: 10.3847/0004-6256/151/6/157

Lim, B., Nazé, Y., Hong, J., et al. 2021, A Kinematic Perspective on the Formation Process of the Stellar Groups in the Rosette Nebula, AJ, 162, 56, doi: 10.3847/1538-3881/abffd8

Lin, D., Bouchet-Valat, M., Noack, A., et al. 2026, JuliaStats/StatsBase.jl: v0.34.10,, v0.34.10 Zenodo, doi: 10.5281/zenodo.18225767

Lissberger, P. H., & Wilcock, W. L. 1959, Properties of All-Dielectric Interference Filters. II. Filters in Parallel Beams of Light Incident Obliquely and in Convergent Beams, Journal of the Optical Society of America, 49, 126, doi: 10.1364/JOSA.49.000126

Lomeli, D., Silsbee, K., Ivlev, A., & Obolentseva, M. 2026, Map of the Interstellar Continuum FUV Field within 800 Parsecs, arXiv e-prints, arXiv:2608.27400, doi: 10.48550/arXiv.2608.27400

Luridiana, V., Morisset, C., & Shaw, R. A. 2015, PyNeb: a new tool for analyzing emission lines. I. Code description and validation of recombination and collisionally excited lines, A&A, 573, A42, doi: 10.1051/0004-6361/201323152

Maragkoudakis, A., Peeters, E., & Ricca, A. 2020, Probing the size and charge of polycyclic aromatic hydrocarbons, MNRAS, 494, 642, doi: 10.1093/mnras/staa681

Meisner, A. M., & Finkbeiner, D. P. 2014, A Full-sky, High-resolution Atlas of Galactic 12 μm Dust Emission with WISE, ApJ, 781, 5, doi: 10.1088/0004-637X/781/1/5

Miville-Deschênes, M.-A., Duc, P.-A., Marleau, F., et al. 2016, Probing interstellar turbulence in cirrus with deep optical imaging: no sign of energy dissipation at 0.01 pc scale, A&A, 593, A4, doi: 10.1051/0004-6361/201628503

Monaco, L., Saviane, I., Perina, S., et al. 2009, The young stellar population at the center of NGC 205, A&A, 502, L9, doi: 10.1051/0004-6361/200912412

Mori, T. I., Onaka, T., Sakon, I., et al. 2014, Observational Studies on the Near-infrared Unidentified Emission Bands in Galactic H II Regions, ApJ, 784, 53, doi: 10.1088/0004-637X/784/1/53

Murgia, G., Cukierman, A. J., Hensley, B. S., et al. 2026, SPHEREx Mapping of Diffuse PAH and H II Emission in the Galactic Plane, ApJ, 1004, 176, doi: 10.3847/1538-4357/ae734b

Naess, S., Aiola, S., Austermann, J. E., et al. 2020, The Atacama Cosmology Telescope: arcminute-resolution maps of 18 000 square degrees of the microwave sky from ACT 2008–2018 data combined with Planck, JCAP, 2020, 046–046, doi: 10.1088/1475-7516/2020/12/046

Nir, G., Zackay, B., & Ofek, E. O. 2018, Optimal and Efficient Streak Detection in Astronomical Images, AJ, 156, 229, doi: 10.3847/1538-3881/aaddff

Park, R. S., Folkner, W. M., Williams, J. G., & Boggs, D. H. 2021, The JPL Planetary and Lunar Ephemerides DE440 and DE441, AJ, 161, 105, doi: 10.3847/1538-3881/abd414

Peeters, E., Habart, E., Berné, O., et al. 2024, PDRs4All. III. JWST's NIR spectroscopic view of the Orion Bar, A&A, 685, A74, doi: 10.1051/0004-6361/202348244

Pence, W. D. 2010, CFITSIO: A FITS File Subroutine Library,, Astrophysics Source Code Library, record ascl:1010.001 http://ascl.net/1010.001

Phelps, R. L., & Lada, E. A. 1997, Spatial Distribution of Embedded Clusters in the Rosette Molecular Cloud: Implications for Cluster Formation, ApJ, 477, 176, doi: 10.1086/303713

Piecka, M., Posch, L., Meingast, S., et al. 2025, Direct measurement of ISM proper motion with image registration, A&A, 702, L1, doi: 10.1051/0004-6361/202556622

Planck Collaboration, Aghanim, N., Akrami, Y., et al. 2020, Planck 2018 results. III. High Frequency Instrument data processing and frequency maps, A&A, 641, A3, doi: 10.1051/0004-6361/201832909

Quinn, J., Bouchet-Valat, M., Robinson, N., et al. 2026a, JuliaData/CSV.jl: v0.10.16,, v0.10.16 Zenodo, doi: 10.5281/zenodo.18681799

Quinn, J., McGrath, M., Aluthge, D., et al. 2025, quinnj/JSON3.jl: v1.14.3,, v1.14.3 Zenodo, doi: 10.5281/zenodo.15497156

Quinn, J., O'Connor, S., Malmaud, J., et al. 2026b, HTTP.jl, Zenodo, doi: 10.5281/zenodo.4527806

Raasveldt, M., & Mühleisen, H. 2019, in Proceedings of the 2019 International Conference on Management of Data (SIGMOD '19) (ACM), 1981–1984, doi: 10.1145/3299869.3320212

Rao, C. R. 1945, Information and the accuracy attainable in the estimation of statistical parameters, Bulletin of the Calcutta Mathematical Society, 37, 81

Rogers, N. S. J., Skillman, E. D., Pogge, R. W., et al. 2022, CHAOS VII: A Large-scale Direct Abundance Study in M33, ApJ, 939, 44, doi: 10.3847/1538-4357/ac947d

Román-Zúñiga, C. G., Elston, R., Ferreira, B., & Lada, E. A. 2008, A FLAMINGOS Deep Near-Infrared Imaging Survey of the Rosette Complex. I. Identification and Distribution of the Embedded Population, ApJ, 672, 861, doi: 10.1086/523785

Sandstrom, K. M., Chastenet, J., Sutter, J., et al. 2023, PHANGS-JWST First Results: Mapping the 3.3 μm Polycyclic Aromatic Hydrocarbon Vibrational Band in Nearby Galaxies with NIRCam Medium Bands, ApJL, 944, L7, doi: 10.3847/2041-8213/acb0cf

Saydjari, A. K., & Finkbeiner, D. P. 2022, Photometry on Structured Backgrounds: Local Pixel-wise Infilling by Regression, ApJ, 933, 155, doi: 10.3847/1538-4357/ac6875

Saydjari, A. K., Uzsoy, A. S. M., Zucker, C., Peek, J. E. G., & Finkbeiner, D. P. 2023, Measuring the 8621 A Diffuse Interstellar Band in Gaia DR3 RVS Spectra: Obtaining a Clean Catalog by Marginalizing over Stellar Types, ApJ, 954, 141, doi: 10.3847/1538-4357/acd454

Schlafly, E. F., Meisner, A. M., & Green, G. M. 2019, The unWISE Catalog: Two Billion Infrared Sources from Five Years of WISE Imaging, ApJS, 240, 30, doi: 10.3847/1538-4365/aafbea

Sellgren, K. 1984, The near-infrared continuum emission of visual reflection nebulae, ApJ, 277, 623, doi: 10.1086/161733

Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, The Two Micron All Sky Survey (2MASS), AJ, 131, 1163, doi: 10.1086/498708

SPHEREx Science Data Center. 2026, SPHEREx Explanatory Supplement, SSDC–DP–017, Version 2.0,, https://irsa.ipac.caltech.edu/data/SPHEREx/docs/SPHEREx_Expsupp_QR.pdf Caltech/IPAC

SSC and IRSA. 2020, Spitzer Enhanced Imaging Products (SEIP),, IPAC/NASA Infrared Science Archive doi: 10.26131/IRSA433

Storey, P. J., & Hummer, D. G. 1995, Recombination line intensities for hydrogenic ions - IV. Total recombination coefficients and machine-readable tables for Z=1 to 8, MNRAS, 272, 41, doi: 10.1093/mnras/272.1.41

The HDF Group. 2025, Hierarchical Data Format, version 5, Zenodo, doi: 10.5281/zenodo.17808558

Tielens, A. G. G. M. 2008, Interstellar polycyclic aromatic hydrocarbon molecules., ARA&A, 46, 289, doi: 10.1146/annurev.astro.46.060407.145211

Tomasi, M., & Li, Z. 2021, Healpix.jl: Julia-only port of the HEALPix library,, Astrophysics Source Code Library, record ascl:2109.028 http://ascl.net/2109.028

Tonry, J. L., Stubbs, C. W., Lykke, K. R., et al. 2012, The Pan-STARRS1 Photometric System, ApJ, 750, 99, doi: 10.1088/0004-637X/750/2/99

Vigoureux, G., Flagey, N., Boulanger, F., et al. 2026, JWST imaging of the Pleiades: Anisotropy of turbulence in the cold neutral medium, Astronomy and Astrophysics, 707, A374, doi: 10.1051/0004-6361/202557000

Villa-Durango, M. A., Barrera-Ballesteros, J., Román-Zúñiga, C. G., et al. 2025, SDSS-V Local Volume Mapper (LVM): revealing the structure of the Rosette Nebula, MNRAS, 543, 1196, doi: 10.1093/mnras/staf1530

Werner, M. W., Roellig, T. L., Low, F. J., et al. 2004, The Spitzer Space Telescope Mission, ApJS, 154, 1, doi: 10.1086/422992

Whidden, P. J., Kalmbach, J. B., Connolly, A. J., et al. 2019, Fast Algorithms for Slow Moving Asteroids: Constraints on the Distribution of Kuiper Belt Objects, AJ, 157, 119, doi: 10.3847/1538-3881/aafd2d

Wright, E. L., Hinshaw, G., & Bennett, C. L. 1996, Producing Megapixel Cosmic Microwave Background from Differential Radiometer Data, ApJL, 458, L53, doi: 10.1086/309927

Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, The Wide-field Infrared Survey Explorer (WISE): Mission Description and Initial On-orbit Performance, AJ, 140, 1868, doi: 10.1088/0004-6256/140/6/1868

Yang, X. J., Li, A., Glaser, R., & Zhong, J. X. 2016, The C-H Stretching Features at 3.2–3.5 μm of Polycyclic Aromatic Hydrocarbons with Aliphatic Sidegroups, ApJ, 825, 22, doi: 10.3847/0004-637X/825/1/22

Zucker, C., Speagle, J. S., Schlafly, E. F., et al. 2020, A compendium of distances to molecular clouds in the Star Formation Handbook, A&A, 633, A51, doi: 10.1051/0004-6361/201936145

Zucker, C., Saydjari, A. K., Speagle, J. S., et al. 2025, A Deep, High-angular-resolution 3D Dust Map of the Southern Galactic Plane, ApJ, 992, 39, doi: 10.3847/1538-4357/adfbe6