

The unWISE Catalog: Two Billion Infrared Sources from Five Years of WISE Imaging

EDWARD SCHLAFLY,^{1,2} AARON M. MEISNER,^{3,1,2} AND GREGORY M. GREEN^{4,5}

¹*Lawrence Berkeley National Laboratory
One Cyclotron Road
Berkeley, CA 94720, USA*

²*Hubble Fellow*

³*National Optical Astronomy Observatory
950 N. Cherry Avenue
Tucson, AZ 85719, USA*

⁴*Kavli Institute for Particle Astrophysics and Cosmology
Physics and Astrophysics Building, 452 Lomita Mall
Stanford, CA 94305, USA*

⁵*Porat Fellow*

ABSTRACT

We present the unWISE Catalog, containing the positions and fluxes of roughly two billion objects observed by the Wide-field Infrared Survey Explorer (WISE) over the full sky. The unWISE Catalog has two advantages over the existing WISE catalog (AllWISE): first, it is based on significantly deeper imaging, and second, it features improved modeling of crowded regions. The deeper imaging used in the unWISE Catalog comes from the coaddition of all publicly available 3–5 micron WISE imaging, including that from the ongoing NEOWISE-Reactivation mission, thereby increasing the total exposure time by a factor of 5 relative to AllWISE. At these depths, even at high Galactic latitudes many sources are blended with their neighbors; accordingly, the unWISE analysis simultaneously fits thousands of sources to obtain accurate photometry. Our new catalog detects sources at 5σ roughly 0.7 magnitudes fainter than the AllWISE catalog and more accurately models millions of faint sources in the Galactic plane, enabling a wealth of Galactic and extragalactic science. In particular, relative to AllWISE, unWISE doubles the number of galaxies detected between redshifts 0 and 1 and triples the number between redshifts 1 and 2, cataloging more than half a billion galaxies over the whole sky.

Keywords: surveys — catalogs — techniques: photometric

1. INTRODUCTION

Expanding interest in cool objects and in high redshifts has driven continual progress in infrared astronomy, as enabled by tremendous improvements in detector sensitivity (Low et al. 2007). Within the past decade, the Wide-field Infrared Survey Explorer (WISE; Wright et al. 2010) has provided an orders of magnitude leap forward relative to its predecessor IRAS (Wheelock et al. 1994), mapping the entire sky at 3–22 microns with unprecedented sensitivity. As next-generation infrared missions like JWST, Euclid and WFIRST move forward, maximizing the value of existing infrared surveys like WISE will be critical.

By now, over 80% of archival WISE data have been acquired via an ongoing asteroid-characterization mission called NEOWISE (Mainzer et al. 2014). We are leading a wide-ranging effort to repurpose NEOWISE

observations for astrophysics, starting by building deep full-sky coadds from tens of millions of 3.4 micron (W1) and 4.6 micron (W2) exposures. Through our resulting “unWISE” line of data products, we have already created the deepest ever full-sky maps at 3–5 microns (Meisner et al. 2017a,b, 2018a), generated a new class of time-domain WISE coadds (Meisner et al. 2018b,c), and performed forced photometry on these custom WISE coadds at the locations of more than a billion optically detected sources (Lang et al. 2014; Schlegel et al. 2015; Dey et al. 2018).

However, until now there has never been a WISE-selected catalog leveraging the enhanced depths achieved by incorporating NEOWISE data. Here we create and release such a catalog. Although WISE delivers exceptionally uniform and high quality imaging, its analysis requires careful application of appropriate computa-

tional techniques. At the WISE resolution of $\sim 6''$, many sources substantially overlap others in the images, even at high Galactic latitudes where the fewest sources are detected. This blending together of nearby sources renders many standard photometry codes unusable. The `crowdsourc` crowded field point source photometry code (Schlafly et al. 2018), recently developed for the DECam Galactic Plane Survey (DECaPS)¹, is well-suited to the task of modeling unWISE images, where nearly all objects are unresolved and blending is pervasive.

We have applied the `crowdsourc` photometry pipeline to deep unWISE coadds built from five years of publicly available WISE and NEOWISE imaging. The result is a catalog of ~ 2 billion unique objects detected in the W1 and/or W2 channels, reaching depths ~ 0.7 magnitudes fainter than those achieved by AllWISE (Cutri et al. 2013). Our “unWISE Catalog” can therefore be considered a deeper W1/W2 successor to AllWISE with more than twice as many securely detected objects. This new catalog will have far-reaching implications, from discovering previously overlooked brown dwarfs in the solar neighborhood (e.g., Kirkpatrick et al. 2011) to revealing quasars in the epoch of reionization (e.g., Bañados et al. 2018).

In §2 we recap the relevant history of the WISE and NEOWISE missions. In §3 we briefly highlight salient features of the unWISE coadds which form the basis of our unWISE Catalog. In §4 we describe our photometry pipeline. In §5 we provide an overview and evaluation of our resulting catalog. In §6 we discuss limitations of our current catalog processing and related avenues for future improvements. In §7 we describe the data release contents. We conclude in §8.

2. WISE OVERVIEW

Launched in late 2009, the WISE satellite resides in a ~ 95 minute period low-Earth orbit. During the first half of 2010, WISE completed its primary mission by mapping the entire sky once in all four of its available channels, labeled W1 (3.4 microns), W2 (4.6 microns), W3 (12 microns) and W4 (22 microns), with a point spread function (PSF) of full-width at half-maximum (FWHM) of 6.1, 6.4, 6.5, and $12''$. Over the following months, WISE ceased observations in W3 and W4 due to cryogen depletion, but nevertheless continued observing in W1 and W2 through early 2011 thanks to an asteroid-characterizing extension called NEOWISE (Mainzer et al. 2011). In 2011 February, WISE was placed into hibernation for nearly three years. In late 2013, how-

ever, it was reactivated, and continued observations in W1/W2 observations as the NEOWISE-Reactivation mission (NEOWISER; Mainzer et al. 2014). The ongoing NEOWISE mission has now obtained nearly five full years (10 full-sky mappings) of W1 and W2 imaging, and has publicly released single-exposure images and catalogs corresponding to the first four of those years (observations acquired between 2013 December and 2017 December; Cutri et al. 2015).

Because NEOWISE is an asteroid characterization and discovery project, the mission itself does not publish any coadded data products of the sort that would maximize the raw NEOWISE data’s value for Galactic and extragalactic astrophysics. AllWISE (Cutri et al. 2013) represents the most recent such set of coadded data products published by the WISE/NEOWISE teams, but was released at a time when only one fifth of the presently available W1/W2 data had been acquired. Because AllWISE already incorporates all available W3 and W4 imaging, we only construct catalogs in W1 and W2 in this work.

3. UNWISE COADD IMAGES

Ideally, our WISE cataloging would proceed by directly and jointly modeling all available W1/W2 exposures. However, doing so would be computationally intensive because these inputs represent ~ 175 terabytes of data spread across ~ 25 million single-exposure (“L1b”) images. As a computational convenience, our cataloging operates on a full-sky set of 36,480 coadded images totaling less than 1 terabyte in size. Specifically, we model deep unWISE coadds built from five years of single-exposure images in each of W1 and W2. These unWISE coadds uniformly incorporate all publicly available single-exposure images ever acquired in these bands, spanning the WISE and NEOWISE mission phases. The unWISE coaddition procedure is described in Lang (2014), Meisner et al. (2017a,b) and Meisner et al. (2018a). The five-year unWISE coadds used in this work are yet to be publicly released (Meisner et al. 2019, in prep.).

The unWISE coadds we employ feature a $5\times$ increase in total exposure time relative to the AllWISE coadds, so that we expect to achieve depths substantially beyond those attained by the AllWISE Source Catalog. Furthermore, the added redundancy of NEOWISE imaging allows our catalog to be relatively free of time-dependent systematics that were present in AllWISE, especially contamination from scattered moonlight.

The 36,480 unWISE coadd images are each 2048×2048 pixels in size with a pixel scale of $2.75''$, covering about 2.5 square degrees. Each image is identified by

¹ <http://decaps.skymaps.info>

one of 18,240 `coadd_id` values giving the location of the coadd on the sky, and a band (W1 or W2). The unWISE tile centers and footprints match those of the AllWISE Atlas stacks.

4. CROWDED FIELD PHOTOMETRY PIPELINE

Cursory inspection of even the highest-latitude, least crowded images reveals that many sources in the unWISE coadds are significantly blended with their neighbors. Accordingly, fully taking advantage of the unWISE coadds requires a crowded-field photometry pipeline that simultaneously models the many blended sources in each field.

The WISE images are in many ways ideal for crowded-field analyses. A substantial challenge in modeling crowded fields is accurate determination of the PSF and its wings. Fortunately, the WISE satellite has a very stable PSF owing to its location above the atmosphere. Accordingly, we can adopt a nearly constant model for the shape of the PSF, and tweak only the very central region as necessary to improve the fit.

The WISE images are also ideal for crowded-field analysis because of their relatively large $\sim 6''$ PSF FWHM. This large PSF means that most distant galaxies are not resolved and can be adequately modeled as point sources, introducing only a small bias. The unWISE Catalog analysis simply assumes that *all* sources are point sources. This tremendously simplifies the modeling relative to typical optical extragalactic surveys with $\sim 1''$ PSFs, where detailed modeling of the shapes of galaxies is required to match the observed images.

We use the `crowdsourc` analysis pipeline to model the unWISE images (Schlafly et al. 2018). This pipeline attempts to simultaneously model all of the sources in each 512×512 pixel region of an unWISE tile, optimizing the positions and fluxes of the sources as well as the background sky to minimize the difference between the observed image and the model. This pipeline was designed for ground-based optical images, but application to WISE images poses few additional problems.

Figure 1 shows examples of the `crowdsourc` modeling in three fields of very different source densities. The first column shows a portion of the COSMOS field, at high Galactic latitude; the second shows a portion of the Galactic anticenter; and the third shows the Galactic bulge. At high latitudes, and even directly in the Galactic plane, the `crowdsourc` model (second row) is an excellent description of the unWISE images (first row). The differences between the two (third row) clearly come substantially from the shot noise in the images, at least outside the cores of bright stars. Often inspection of the

residual image at the locations of the detected sources (fourth row) shows no coherent residual signatures.




Figure 1. `crowdsourc` modeling results in three fields with very different source densities: the high-latitude COSMOS field (left), the Galactic anticenter (middle), and the Galactic bulge (right). From top to bottom, the rows show the unWISE coadded images, the `crowdsourc` model, the residuals, and the residuals with the locations of cataloged sources overplotted. Except for in the bulge field, shot noise accounts for a substantial fraction of the residuals. On the other hand, in the densest regions, like the bulge, the residuals are completely dominated by unresolved sources and challenges in sky subtraction. The bulge field is stretched $10\times$ less hard than the other two fields. In all cases, the model images account for most of the flux in the real images, though clearly in the bulge case significant residuals remain. All images are W1.

Meanwhile, in the Galactic bulge (third column), the story is very different. The large WISE PSF coupled with the tremendous number of sources and insensitivity to dust extinction make this field very confused. While the bulge image (top row, right) and model (second row, right) are qualitatively in good agreement, the residu-

als (third row, right) are entirely dominated by coherent structures stemming from the incomplete identification of significant sources in the field. It may be possible to do better here by allowing `crowdsourc` to more aggressively identify unmodeled stars overlapping brighter neighbors, but even very slight errors in the PSF model will have substantial effects on the residuals in fields as dense as this one, rendering the results of a more aggressive source identification uncertain at best.

A modest number of modifications and improvements were made to the `crowdsourc` pipeline to allow it to model WISE images. These changes included changes to the PSF modeling, the mosaicing strategy, the sky modeling, and the treatment of nebosity and large galaxies. Additionally, a new diagnostic field `spread_model` was added to the pipeline outputs, which can help determine the size of detected objects.

4.1. PSF modeling

The most significant addition to `crowdsourc` was a PSF modeling module specifically designed for WISE. The unWISE PSF module² provides a 325×325 pixel PSF, developed using bright isolated stars (Meisner & Finkbeiner 2014). This PSF extends far into the wings of the WISE PSF, which has a full-width at half-maximum of only about 2.5 pixels. It includes details of the PSF like diffraction spikes, optical ghosts, and the optical halo. This model was originally designed for application to the unWISE single-exposure images. The unWISE coadd images sum many single-exposure images together, necessitating changes to the original unWISE PSF model. Near the ecliptic poles, the single-exposure images contributing to a given coadd image span a wide range of different detector orientations relative to the sky, leading the final coadd image PSF to be blurred over a range of azimuth. This is modeled by transformation of the PSF to polar coordinates and convolution with a boxcar kernel in azimuth, with the width and position of the boxcar kernel dependent on ecliptic latitude and longitude. The PSF is then projected back to cartesian coordinates.

The resulting model gives an excellent description of the observed PSF at any particular point in the survey. The convolution process is somewhat expensive, however. To save time, a grid of these PSF models is generated for each unWISE tile. At low ecliptic latitudes, the coadd PSF is essentially constant over each unWISE tile, and a relatively coarse grid is adequate. At high ecliptic latitudes (within about 20° of the ecliptic poles), the mean spacecraft orientation and the width of

the range of orientations vary more significantly across each coadd tile. This requires a denser grid of PSFs to model. Within 5° of the north ecliptic pole, one PSF per 128×128 pixel region is generated. The `crowdsourc` pipeline linearly interpolates between these PSFs when generating a model for a source on a tile.

These PSFs are reasonably accurate within a hundred pixels of the PSF center. Beyond this region, subtleties with the world-coordinate system of the coadds versus the PSF become important. Moreover, some coherent residuals in the PSF cores are also apparent. The `crowdsourc` PSF module addresses these by, at each `crowdsourc` modeling iteration, finding the median residual over all bright, unsaturated stars in the inner 9×9 pixel PSF core, and adding it to the PSF model. Typical remaining residuals stem from imperfect subpixel interpolation of the PSF, as indicated by ringing in the residuals around bright stars, as seen in the anticenter residual image in Figure 1. Further effort could eliminate these artifacts; the huge number of stars in each image clearly provide more than enough information to constrain the subpixel shape of the PSF.

4.2. Mosaicing Strategy

Modeling crowded images can require substantial amounts of computer memory. Individual unWISE coadd images in the inner Galaxy contain more than 200,000 detected sources. `crowdsourc` constructs a sparse matrix containing the PSF for each of these sources. Bright sources can require PSFs extending out to 299×299 pixels, for $\sim 90,000$ values per source, so a naive approach could require ~ 200 GB of memory for the sparse matrix alone in this extreme case.

To make this more manageable, `crowdsourc` splits each unWISE tile into 512×512 pixel subimages, with an additional 150 pixel border on each side. In the original incarnation of `crowdsourc` for the DECam Plane Survey, these final catalogs for each subimage were created in turn. In order to prevent duplicate source detection in overlap regions, the analysis of later subimages included fixed sources at the locations of sources detected in earlier subimages.

In contrast, in unWISE source detection occurs for the entire coadd simultaneously. The parameters of these sources are then optimized on each image subregion. Finally, again over the entire coadd, the PSF is refined and source detection is repeated. This new approach preserves the computer-memory advantages of the former approach while allowing PSF modeling to be performed on the entire coadd and more gracefully handling sources in the overlap regions between subimages.

² https://github.com/legacysurvey/unwise_psf

4.3. Sky Modeling

A slightly different approach for sky modeling was taken for WISE than for the DECam Plane Survey. In the DECam Plane Survey, the sparse linear algebra solver was allowed to adjust the overall sky level simultaneously with the fluxes of the stars. In dense regions, this allows an an initial sky overestimate (due to the presence sources in the image) to be improved by simultaneously decreasing the sky and increasing the fluxes of the sources.

This global approach has the disadvantage that isolated, very poorly fitting regions of the image can significantly drive the sky estimate over the entire image. The WISE coadds feature a larger dynamic range than DECam images, making it easier for small residuals in the cores of bright stars to make outsize contributions to the likelihood in the fitting. This problem was addressed by eliminating the overall sky parameter from the sparse linear algebra fit. The sky, however, is still improved at each iteration by median filtering the residual image, as in (Schlafly et al. 2018). This area of the pipeline offers room for improvement; see §6.3.

4.4. Nebulosity and large galaxies

Reflected starlight and thermally emitted light from dust grains can add a diffuse component with rich small-scale structure to observed infrared images. The `crowdscore` modeling assumes that all flux not explained by a smooth sky model must be attributable to point sources. Similarly, large galaxies present in the WISE imaging will be split into many point sources unless preventative measures are taken.

To address both of these cases, we identify nebulosity and large galaxies ahead of time and mark these regions. During source detection, candidate sources found in regions marked as containing nebulosity are then required to be sharp; significantly blended objects in regions with nebulosity will not be modeled. In Schlafly et al. (2018), these sources were additionally required to not overlap any neighbors substantially, but for WISE we removed this constraint. In the optical, regions with nebulosity are usually sparsely populated with stars; the extinction associated with the gas and dust limits the depth of the survey in nebulous areas. In contrast, the infrared light observed by WISE is hardly extinguished by dust, causing application of a strong blending criterion to source detection to eliminate many real sources.

Sources found in regions containing large galaxies, on the other hand, are not required to be sharp, but *are* required to not overlap any neighboring source substantially. This is appropriate since peaks corresponding to large galaxies will naturally be extended and fail a

sharpness cut, but only a single source should be associated with these galaxies. Requiring new sources to not overlap existing ones significantly discourages `crowdscore` from splitting these galaxies into many point-source components. The delivered fluxes, however, will still be inaccurate, as these large galaxies are modeled as single point sources.

Regions of nebulosity are identified with the same machine learning approach as in Schlafly et al. (2018). A few minor changes were made, however. First, the convolutional neural network was trained with 256×256 pixel images, instead of the 512×512 images used in Schlafly et al. (2018). Second, a slightly shallower convolutional neural network (see Appendix A) was adopted. Finally, the neural network was retrained using a set of WISE images rather than optical images. The neural network nebulosity classifications are available as part of the mask images in the Data Release (§7).

The HyperLeda catalog (Makarov et al. 2014) of large galaxies was used to mark regions containing galaxies resolved in WISE ($D_{25} > 10''$). These regions are marked in the unWISE mask images; see Meisner et al. (2019, in prep.) for details of their implementation in unWISE. HyperLeda is also used to select large galaxies for further analysis in the Legacy Survey Large Galaxy Atlas³.

4.5. `spread_model`

The unWISE catalog modeling assumes that all sources are point sources; their shapes are completely described by the modeled PSF. Nevertheless, roughly half of all sources in the catalog are galaxies, which are not point sources. The size of a source can be a useful diagnostic for identifying galaxies or problems with the modeling; for example, a single catalog source that corresponds to two overlapping point sources should be slightly larger than the PSF. Additionally, instrumental effects, like trends in PSF shape with magnitude (for example, as caused by the brighter-fatter effect, Downing et al. 2006; Antilogus et al. 2014), or with color (e.g., PSF chromaticity, Cypriano et al. 2010) can be identified through measurements of source size. The `crowdscore` pipeline now follows the lead of `SExtractor`, computing `spread_model` as a measure of a source’s size (Bertin & Arnouts 1996; Desai et al. 2012)⁴. In combination with its uncertainty `dspread_model`, this term roughly

³ <https://github.com/mostakas/LSLGA>

⁴ The `crowdscore spread_model` is corrected from the formulation in Desai et al. (2012) following the `SExtractor` documentation so that in the absence of noise point sources have `spread_model` equal to 0.

indicates how significantly increasing the size of the PSF would improve the fit.

5. UNWISE CATALOG

The unWISE catalog uses the deep unWISE coadds to detect sources a factor of two fainter than possible in AllWISE, detecting a total of about 2 billion objects. It couples this with the `crowdsourc` source detection and characterization pipeline to model significantly blended sources. Figure 2 illustrates the improvement realized by this combination, showing the AllWISE and unWISE images with the corresponding catalogs, and finally the much deeper, higher-resolution S-COSMOS imaging from Spitzer (Sanders et al. 2007a).

The dramatic improvement in depth realized by the four additional years of NEOWISE-Reactivation imaging is immediately apparent comparing the upper left and middle left panels of Figure 2. Additionally, the importance of modeling blended sources is clear. Many sources within the unWISE coadd significantly overlap one another, even in this high-latitude field. Comparison of the unWISE coadd image with the corresponding catalog (middle right) suggests that `crowdsourc` has done a good job identifying sources in the field. The much deeper Spitzer imaging (bottom row) generally confirms this impression, though there are clear instances of cases where `crowdsourc` has split galaxies into multiple sources (for example, the large galaxy near (530, 1526)), as well as cases where it has failed to split overlapping, unresolved, faint sources into their components (for example, the pair of sources near (503, 1518)). The former case is not particularly troubling; there are not that many galaxies resolved by WISE, and the ones that exist are usually securely identified in higher-resolution imaging.

Unsurprisingly, the unWISE catalog contains many more objects than the AllWISE catalog. Figure 3 shows the number densities of 5σ sources per square degree detected in unWISE (top) as compared with AllWISE (middle), and the ratio of the two number density maps (bottom). In typical high-latitude locations, unWISE includes about $2\times$ as many objects in W1 and $2.5\times$ as many objects in W2. In the Galactic plane, the difference is more dramatic, with unWISE cataloging $3\text{--}3.5\times$ more sources than AllWISE, owing to the aggressive identification of blended stars in `crowdsourc`.

The ratio maps additionally show some isolated regions of elevated unWISE/AllWISE detections at high latitudes. These tend to surround the brightest ~ 2000 stars in the sky. The unWISE Catalog detects too many sources in these regions, incorrectly interpreting errors in the wings of the PSF model as faint sources. All-

WISE, on the other hand, detects too few sources in these regions, because it overestimates the sky in the vicinity of bright objects and because it relegates sources too near bright objects to the “reject” table.

Regions of substantial nebulosity also are evident in Figure 3 as having low source densities in unWISE and AllWISE; the clearest example is in the vicinity of the Orion Molecular Cloud near $(l, b) = (-150^\circ, -20^\circ)$. The unWISE Catalog imposes an additional criterion on candidate sources in these regions to limit the number of spurious objects detected (see §4.4). However, significant numbers of real sources are also excluded by this criterion.

Clear striping along lines of constant ecliptic longitude are especially evident in the number count ratios (Figure 3, bottom). These stripes usually correspond to periods of the survey where the Moon was near the part of the sky being observed, scattering light into the images and reducing their depth. These features are absent from the unWISE catalog because the additional years of NEOWISE observations provide moon-free imaging of these regions. The particularly prominent stripe in the W1 ratio map lies at ecliptic longitude $\lambda \approx 240^\circ$, which was observed at the beginning of the 3-band Cryo phase. Sources in AllWISE in this region may be reported as having zero uncertainty, leading them to be excluded from Figure 3; see §II.2.c.ii of Cutri et al. (2013) for details.

5.1. Completeness and Reliability

To assess the completeness and reliability of our catalog in a representative sky location, we compare against Spitzer data in the COSMOS region. COSMOS is at high Galactic latitude, and so serves as a typical extragalactic field. Moreover, it is at low ecliptic latitude, so the WISE imaging is typical in depth. For our purposes, S-COSMOS is the preferred Spitzer data set in the COSMOS region, as it is much deeper than our unWISE catalog and covers the entire 2 square degree COSMOS footprint. Our analyses compare WISE W1 and W2 against Spitzer ch1 and ch2, respectively, since these pairs of bandpasses are quite similar, although not identical. In both our completeness and reliability analyses, we use a Spitzer-WISE cross-match radius of $2''$. Portions of four unWISE `coadd_id` footprints contribute to the analysis: 1497p015, 1497p030, 1512p015, 1512p030. All completeness and reliability values presented are differential.

To measure the unWISE catalog’s completeness, we wish to compare against a highly reliable Spitzer catalog. For this reason, we take as a sample of “true” Spitzer sources the subset of S-COSMOS catalog en-




Figure 2. AllWISE and unWISE compared with much deeper, higher-resolution imaging from Spitzer-COSMOS, for a small portion of the COSMOS field. The three rows show $3.4\mu\text{m}$ imaging of the same small patch of high-latitude sky from AllWISE (top), unWISE (middle), and Spitzer-COSMOS (right, $3.6\mu\text{m}$). The left column shows only the images, while the right column overplots the 5σ catalog entries from AllWISE (top) and unWISE (middle, bottom). The deeper unWISE stacks clearly allow many more sources to be detected, and the **crowdsourcing** catalog well describes these. Nevertheless, comparison with the Spitzer imaging reveals clear examples of unidentified sources (for instance, near (503, 1518)) and resolved sources that are split into multiple point sources (for instance, near (530, 1526)). Axis units are WISE pixels on `coadd_id 1497p015`.

tries with `f1_c? = 0` and `flux_c?_4 > 0` in the relevant channel. One consequence of our decision to cut on the S-COSMOS quality flag `f1_c?` is that sources brighter than ~ 13 in `ch1` and `ch2` are rejected, preventing our completeness analysis from reaching bright magnitudes. We also remove Spitzer sources with locations marked by `flags_info` bit 1 as being associated with a WISE-resolved galaxy – this effectively restricts our analysis to pointlike sources. The top two panels of Figure 4 show the completeness as a function of Spitzer magnitude, for both our unWISE Catalog and the AllWISE Source Catalog. Because AllWISE performs forced photometry in every WISE band for sources detected in any band, we have restricted to sources with signal-to-noise ratio greater than 5 in the corresponding band. This provides a fair comparison against the unWISE Catalog, which does not perform forced photometry. In W1, we find

that the unWISE Catalog is 0.76 mags deeper than AllWISE, with the former (latter) reaching 50% completeness at `ch1 = 17.93 (17.17)`. In W2, we find that the unWISE Catalog is 0.67 mags deeper than AllWISE, with the former (latter) reaching 50% completeness at `ch2 = 16.72 (16.05)`. In AB, these unWISE catalog depths are 20.72 (19.97) in W1 (W2).

Given that our unWISE catalog benefits from $\sim 5\times$ enhanced integer coverage relative to AllWISE, one might naively expect that we should reach $2.5\log_{10}\sqrt{5} = 0.87$ mags deeper in both W1 and W2, assuming that the pre- and post-hibernation data are of identical quality. In detail, there are two main factors that result in a lesser depth enhancement. First, the WISE PSF is a few percent broader in both W1 and W2 in post-hibernation imaging relative to pre-hibernation imag-

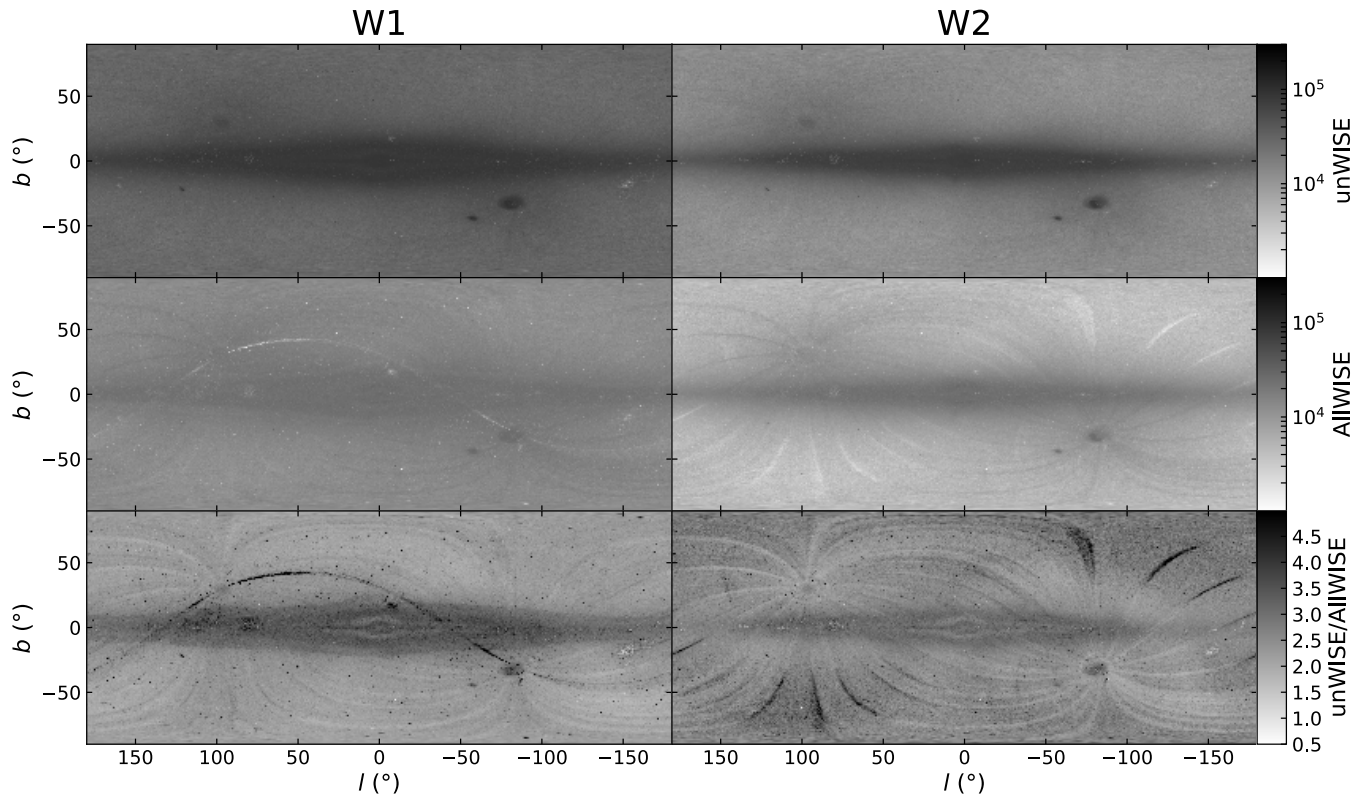


Figure 3. Number density of sources per square degree cataloged by unWISE and AllWISE. The top row shows the unWISE catalog densities; the middle row shows AllWISE; and the bottom row shows the ratio of the two. unWISE detects more than $2\times$ as many sources as AllWISE at high latitudes, and more than $3\times$ as many sources as AllWISE at low latitudes. Scattered light from the Moon leads to significant variations in the AllWISE number densities in W2 at high latitudes; these are mostly eliminated in unWISE owing to the greater number of observations, which allows all parts of the sky to be observed in conditions free from significant scattered moonlight. Isolated regions of elevated unWISE/AllWISE source count ratio usually correspond to the ~ 2000 brightest stars in the sky, where unWISE tends to detect too many sources (incorrectly identifying features in the wings of the PSF of very bright stars as separate sources) and AllWISE tends to detect too few sources (missing many objects in the wings of bright stars). Regions with significant nebulosity show a deficit of objects in both unWISE and AllWISE—for instance, in Orion near $(-150^\circ, -20^\circ)$.

ing⁵. This increases the number of effective noise pixels (n_{eff}) for point sources, which scales like FWHM^2 , while the limiting flux then scales as $\sqrt{n_{\text{eff}}} \propto \text{FWHM}$, so that the post-reactivation sensitivity is correspondingly decreased by a few hundredths of a mag. Furthermore, the W1 and W2 sensitivities have decreased slightly post-reactivation⁶. The post-reactivation sensitivity decrease ranges from 0.05–0.12 mag in W1 and 0.15–0.26 mag in W2 when measured in yearly intervals during the NEO-WISE mission. Combined, the increased PSF size and decreased sensitivity of post-reactivation data can account for the discrepancy between our achieved depths

⁵ http://wise2.ipac.caltech.edu/docs/release/neowise/expsup/sec4_2bi.html

⁶ http://wise2.ipac.caltech.edu/docs/release/neowise/expsup/sec2_1e.html, see Figure 1.

and the simplistic projection of 0.87 mag improvement. Confusion noise is an additional factor that would tend to reduce the depth improvement achieved relative to estimates based purely on reduced statistical pixel noise.

To measure the unWISE Catalog’s reliability, we wish to compare against a highly complete Spitzer catalog. Therefore, in our reliability analysis, we compare against the entire S-COSMOS catalog without making any quality or flux cuts. We again require AllWISE $w?\text{snr} \geq 5$ and unWISE $\text{flux}/\text{dflux} \geq 5$ in the band under consideration. For unWISE we required $\text{flags_unwise} = 0$, and that none of flags_info bits 1, 6, 7 be set (see Table 5). For AllWISE we required $w?\text{cc_map} = 0$. We find that the unWISE Catalog has roughly comparable reliability to AllWISE until reaching sufficiently faint magnitudes that AllWISE no longer contains any sources in the sky region analyzed. In both W1 and W2, AllWISE appears to have very high reliability up until the point




Figure 4. Summary of our completeness and reliability analysis based on a comparison against the Spitzer S-COSMOS data set over ~ 2 square degrees of extragalactic, low ecliptic latitude sky. Top: differential completeness as a function of Spitzer magnitude for AllWISE (green squares) and the unWISE Catalog (black plus marks). The unWISE Catalog hits 50% completeness ~ 0.7 magnitudes fainter than AllWISE in both bands. Bottom: differential reliability as a function of WISE magnitude, with the same marker symbols/colors as above. Vertical dashed green lines indicate magnitude bins with < 10 AllWISE $w?snr \geq 5$ sources in the relevant band. Vertical dashed black lines indicate magnitude bins containing < 10 sources with signal-to-noise ≥ 5 in either AllWISE or unWISE. Note that although the numerical values along the horizontal axes of the upper and lower panels are aligned, their units are different (Spitzer magnitudes in the upper panels and WISE magnitudes in the lower panels).

that it contains no more sources. We interpret this behavior as arising from the fact that AllWISE catalog construction and particularly its artifact flagging were engineered for high reliability, with the AllWISE Reject Table being used as a repository for lower confidence sources. The unWISE Catalog has no “reject table”, and so its reliability rolls off toward faint magnitudes until reaching 0.62 (0.70) in W1 (W2), at which point there are no fainter $\geq 5\sigma$ sources in this sky region. The vast majority of AllWISE sources are detected at $\geq 5\sigma$ in W1, so that the W1 AllWISE sample reflected in the

bottom left panel of Figure 4 can be considered roughly W1-selected. The unWISE sample in that same subplot is strictly W1-selected, making this a fair comparison. The W1 unWISE reliability is superior to the AllWISE reliability at the faintest AllWISE magnitudes, which makes sense given that the unWISE Catalog has benefited from a larger amount of W1 imaging. On the other hand, in W2, the AllWISE reliability appears to be slightly higher than the unWISE reliability at the faintest AllWISE magnitudes. We attribute this to the fact that AllWISE W2 detections are benefiting from

W1 information (e.g., during the source identification and centroiding) as a result of being jointly fit across all bands, while the unWISE Catalog fits the two bands entirely independently. The relatively high AllWISE reliability around $W2 \sim 16$ also coincides with AllWISE completeness which is much lower than that of the unWISE Catalog at the same magnitudes.

The depth of the unWISE Catalog varies across the sky, generally becoming deeper at higher ecliptic latitude and shallower at lower Galactic latitude. Especially in W2, zodiacal light further reduces the depth at low ecliptic latitude. The COSMOS field is at an ecliptic latitude of $\sim 10^\circ$, so the COSMOS field corresponds to a relatively shallow region in unWISE.

5.2. Astrometry

The simplest test of the unWISE Catalog astrometry is to compare it with Gaia (Gaia Collaboration et al. 2016). The comparison shows systematic offsets in declination of about -1 and 0 mas in right ascension (α) and 28 and 36 mas in declination (δ), in W1 and W2, respectively. These offsets vary spatially over the sky, dominated by signals related to Galactic rotation; the rms scatter in the offset over the sky is 51 and 51 mas in α and 38 and 39 mas in δ in W1 and W2, respectively. After removing these offsets over the sky, the rms scatter in α and δ for individual bright stars ($G < 15$) between Gaia and unWISE is 39 mas.

That said, a number of features appear in the comparison of the astrometry to Gaia unrelated to Galactic rotation. To inspect these in more detail, we compare the unWISE astrometry to AllWISE. In analyzing the unWISE coadds, we attempt to directly propagate the underlying WISE single-exposure astrometry forward without adjustment. The unWISE Catalog astrometry should therefore agree quite well with the AllWISE astrometry. Figure 5 shows the mean and root-mean-square (rms) difference between the measured unWISE and AllWISE right ascension (top) and declination (bottom) across the sky in W1 (left) and W2 (right) for bright stars ($W1 < 14$ or $W2 < 14$). Overall, agreement is excellent. There are spatially correlated differences of 25 mas in size, but this is only a hundredth of a WISE pixel. Individual bright stars agree in position in AllWISE and unWISE with an rms difference of 70 mas. We note that since AllWISE simultaneously fits the W1 and W2 bands, but W1 provides substantially more signal-to-noise for typical sources, Figure 5 primarily compares unWISE W1 and W2 to AllWISE W1. In particular, differences between the W1 and W2 comparisons stem from differences in the unWISE W1 and W2 astrometry; the AllWISE astrometry is identical.

Despite this great overall agreement, Figure 5 shows interesting structures in the differences between the unWISE and AllWISE astrometry. Most prominently, the differences in α and δ show stripes along lines of constant ecliptic longitude. This is due to a difference between the definitions of the “center” of the PSF in unWISE and AllWISE. AllWISE adopted a PSF with a flux-weighted centroid slightly offset from the origin, while unWISE uses a PSF with a centroid at the origin (Meisner et al. 2018b). However, unWISE adopts the published WISE WCS solutions without modification, leading unWISE to report slightly different world coordinates for sources than AllWISE. This difference is most pronounced in locations where, due to Moon avoidance maneuvers, more observations have been made in one WISE scan direction than in the other, leading to structures at constant ecliptic longitude. The larger amplitude of the striping in W1 than W2 is due to the greater PSF asymmetry in W1 than W2 (Meisner et al. 2018b).

Another source of low-level disagreement between unWISE Catalog positions and AllWISE positions are the “MFPRex” astrometric corrections (Cutri et al. 2013, §V.2) implemented in the AllWISE pipeline but not reflected in the single-exposure WCS used by unWISE. These corrections presumably are the source of the coherent small scale features in Figure 5, most prominent near the Galactic Center in the δ offset map (second row). These corrections additionally account for the proper motions of the stars in the astrometric reference catalog, leading to some small differences between AllWISE and unWISE on large angular scales. Comparison to Gaia confirms that the AllWISE astrometry is superior to unWISE, owing to these corrections, so users seeking the best possible astrometry are encouraged to use AllWISE until the unWISE astrometry can be improved (§6.6).

The δ offset maps (second row of Figure 5) show significantly smaller residuals in the celestial south than in the north; we do not understand the cause of this north-south discrepancy.

The rms differences in α and δ are very uniform over the high-latitude sky. In the inner Galaxy, however, they increase dramatically, presumably owing to the more aggressive unWISE identification of stars in the wings of neighboring stars and the corresponding impact on the positions of the stars and their neighbors. Similarly, the Magellanic Clouds and M31 show enhanced rms. Two small regions of slightly enhanced residuals and scatter in δ fall near $(113^\circ, 17^\circ)$ and $(148^\circ, -18^\circ)$ and are not understood, though the modeling in these regions looks accurate, leading us to suspect problems with the WCS.




Figure 5. Astrometry comparison between unWISE and AllWISE for bright, relatively isolated stars. Each panel is a map of the sky; the rows show the average difference in right ascension α and declination δ , as well as the rms difference in α and δ over the sky. All units are in milliarcseconds (mas). The left column is W1, while the right column is W2. The mean μ and standard deviation σ of each map above $|b| = 15^\circ$ is shown. Agreement is excellent; typical correlated residuals are only 25 mas, less than a hundredth of a WISE pixel. Individual bright stars agree in position to an rms difference of 70 mas. Nevertheless, many WISE survey and processing related structures are evident; see text for details. The color scale covers 200 mas in all panels.

5.3. Photometry

To assess the accuracy of the unWISE photometry, we compare it with AllWISE photometry in Figure 6. Agreement is excellent, with overall offsets of 7 and 35 mmag in W1 and W2 and sub-percent level fluctuations over the high-latitude sky. The rms difference between unWISE and AllWISE photometry of bright stars is 15 mmag in W1 and 17 mmag in W2.

However, a number of structures are evident in the comparison. Most prominently, the comparison fea-

tures a number of streaks running between the ecliptic poles at constant ecliptic longitude. These appear to stem from the AllWISE photometry, as they correlate with the moon-avoidance maneuvers made during the pre-hibernation phase of the WISE mission; the post-reactivation WISE imaging has now covered the sky more uniformly. Structures are apparent at $(l, b) = (145^\circ, 45^\circ)$ and $(-40^\circ, -40^\circ)$, which stem from the WISE spacecraft's dumping angular momentum using its magnetic torque rods (Cutri et al. 2013, IV.2.c.i.).

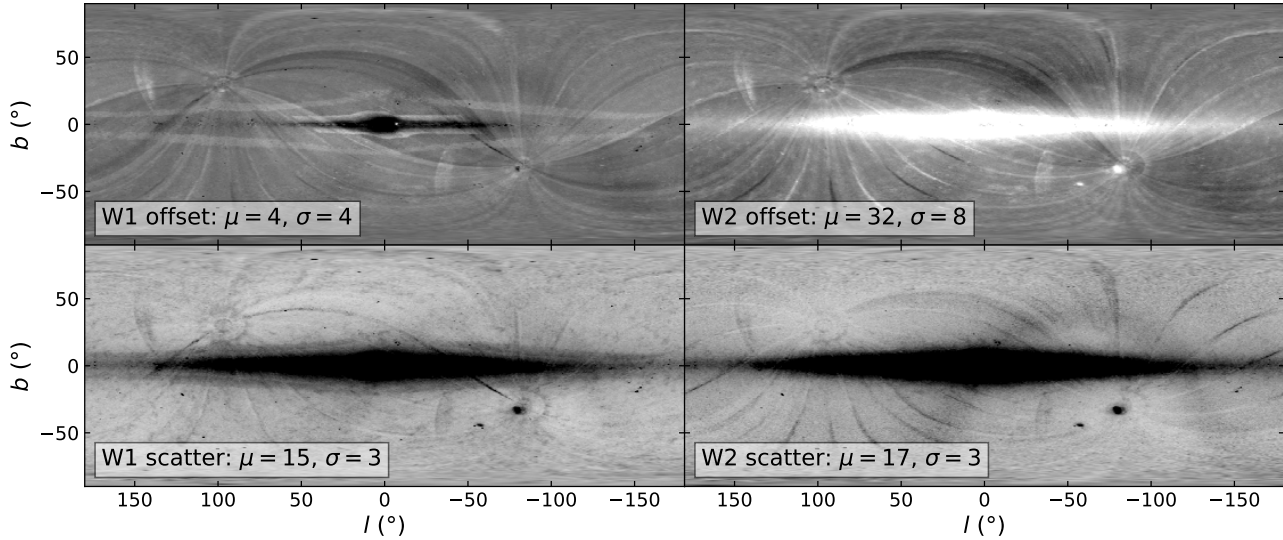


Figure 6. Photometry comparison between unWISE and AllWISE for bright (W1 or W2 brighter than 14th mag), relatively isolated stars. Each panel is a map of the sky; the first row shows the average difference in magnitude between unWISE and AllWISE, while the second row shows the rms scatter in the magnitude differences. The left column shows W1, while the right column shows W2. The mean μ and standard deviation σ of each map above $|b| = 15^\circ$ is shown, with units of mmag. Agreement is excellent, with percent-level offsets and uniformity, and scatter of 1%, though many WISE survey and processing related structures are evident; see text for details. The color scale covers 100 mmag in the top two panels and 50 mmag in the bottom two panels.

A dark band of higher-than-average scatter is apparent in the W2 rms map at $\delta \approx -30^\circ$, due to the South Atlantic Anomaly. The improved cosmic-ray rejection made possible by the many NEOWISE epochs (Meisner et al. 2017b) allows the unWISE photometry to improve on the AllWISE photometry in this region. Very crowded regions and regions with significant nebulosity appear as areas of poor AllWISE-unWISE agreement: for example, the Galactic plane, the Large and Small Magellanic Clouds, large globular clusters, and the Orion nebula. The W1 offset map additionally shows some peculiar regions around the Galactic plane and inner bulge where stars are preferentially a couple hundredths brighter in unWISE than in AllWISE. We do not understand these features, but comparison with 2MASS indicates that they are present in AllWISE and absent from unWISE.

Figure 6 addresses the spatial uniformity of the photometry of bright stars, but is insensitive to the catalog accuracy of faint stars. Figure 7 shows the differences between unWISE and AllWISE photometry as a function of magnitude for point sources identified by Gaia in a 25 square degree region around the COSMOS field. Agreement is again good; for bright stars the rms difference is a few hundredths, as anticipated in Figure 6, and near the AllWISE faint limit the uncertainties increase to the expected ≈ 0.2 magnitudes. In W2, the unWISE

and AllWISE fluxes have a tight linear relationship between the saturation limit at about 8th magnitude to the faint limit at about 16th magnitude. However, in W1 there is a clear trend in the magnitude difference with magnitude; unWISE sources are 0.03 mag fainter than AllWISE sources at 8th mag, while they are 0.01 mag brighter than AllWISE sources at 17th mag. It is unclear where this nonlinearity comes from; one hint is that the measured sizes of point sources likewise depends on magnitude, albeit in both W1 and W2; see §6.1 for further discussion. At about 8th mag in both W1 and W2, the unWISE magnitudes depart sharply from the AllWISE magnitudes. This is due to the onset of saturation. Due to the way the unWISE patches the saturated cores of stars in the construction of the deep unWISE coadds (Lang 2014), the inner 7×7 pixel regions of saturated stars are unreliable. Flux estimates of saturated stars are then made entirely on the basis of the wings of the PSF in the unWISE Catalog, outside the 7×7 pixel region where roughly 90% of the flux resides. The onset of saturation is tracked in the unWISE Catalog columns `flags_unwise` and `qf`, allowing easy identification of saturated sources (§7).

The unWISE Catalog absolute photometric calibration derives from the photometric calibration of the unWISE coadds (Meisner et al. 2017a), which is tied to the original WISE zero points through aperture fluxes in a

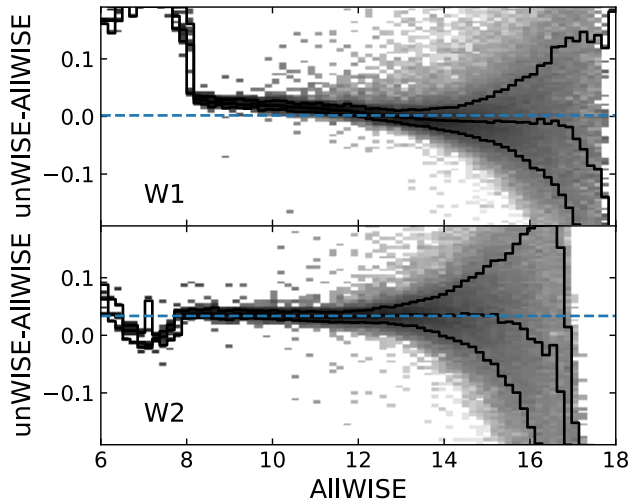


Figure 7. Photometry comparison between unWISE and AllWISE as a function of magnitude. The grayscale shows the number of stars at each unWISE-AllWISE magnitude difference as a function of the magnitude of the stars. The solid lines show the 16th, 50th, and 84th percentiles of the distribution at each magnitude. There is good agreement between AllWISE and unWISE at bright magnitudes fainter than the WISE saturation limit of 8, with about 15 mmag scatter. W2 shows excellent linearity, but in W1 there is a trend of 40 mmag from 8th to 16th mag, with unWISE being fainter than AllWISE for bright stars and brighter than unWISE for faint stars.

27.5'' radius. The unWISE Catalog fluxes are defined in the context of a PSF that is normalized to unity in a 19×19 pixel box, corresponding to $52.25''$. Meanwhile the AllWISE PSFs are normalized to unity in a $220''$ box. Because the PSF is extremely stable, these different conventions lead one to expect slight offsets (~ 30 mmag) in the absolute calibration of unWISE and AllWISE, consistent with Figure 6. We recommend subtracting 4 mmag and 32 mmag in W1 and W2 from the unWISE magnitudes to better match the AllWISE absolute flux calibration, with the caveat that the W1 correction is particularly uncertain because of the W1 nonlinearity (Figure 7).

5.4. Example Uses

The unWISE Catalog should prove a valuable resource for a range of astronomical applications. This section presents simple examples of potential uses in Galactic, extragalactic, and high-redshift science.

5.4.1. Galactic

A wide variety of Galactic science relies on WISE infrared fluxes, ranging from nearby stars (Kirkpatrick et al. 2011) to dust-extinguished stars throughout the

Galaxy (Schlafly et al. 2016, 2017). Figure 8 illustrates the improvement of the unWISE fluxes relative to AllWISE in three fields of different source densities: the high-latitude COSMOS field (top row), the Galactic anticenter (middle row), and the Galactic bulge (bottom row). The figure also shows high-resolution data from Spitzer, taken from the ultra deep S-COSMOS program on the COSMOS field (Sanders et al. 2007a), from GLIMPSE360 in the Galactic anticenter, and from GLIMPSE3D in the Galactic bulge (Benjamin et al. 2003; Churchwell et al. 2009). In each case, the greater depth of the unWISE imaging (left) than the AllWISE imaging (center) is apparent; the sources in the color-magnitude diagrams extend roughly 0.7 mag fainter in the unWISE diagrams than the AllWISE diagrams.

In the COSMOS field and toward the Galactic anticenter, the unWISE measurements are very similar to the AllWISE measurements and extend them naturally to fainter magnitudes. In the bulge, however, there is a pronounced difference between the unWISE and AllWISE color-magnitude diagrams; the typical star in AllWISE becomes bluer at fainter magnitudes, while it becomes redder in unWISE. The higher resolution Spitzer observations better match the unWISE measurements than the AllWISE measurements, giving us confidence that unWISE is improving on AllWISE in dense regions like the bulge. That said, in these regions the Spitzer data are substantially superior to the unWISE data—unsurprisingly, given the extreme crowding and $\sim 4\times$ better resolution of Spitzer.

5.4.2. Extragalactic

At $z = 0$, the WISE W1 and W2 bands sample a steeply falling portion of the typical early-type galaxy’s spectral energy distribution. At higher redshifts, more and more of a galaxy’s light redshifts into the WISE bands. This makes WISE an effective tool for detecting galaxies at redshifts $0 < z < 2$. To illustrate this, Figure 9 shows the redshift distribution of galaxies detected by AllWISE and unWISE in the COSMOS field (Sanders et al. 2007b), with redshifts taken from the photometric redshift catalog of Laigle et al. (2016). WISE sources were matched to COSMOS sources with a match radius of $2.75''$, considering only COSMOS sources with Spitzer $3.6\mu\text{m}$ or $4.5\mu\text{m}$ fluxes bright enough that the sources could conceivably be detected in WISE. The resulting redshift distribution peaks at $z \approx 1$. The bulk of the distribution falls between $0 < z < 2$, with a tail to higher redshifts.

Figure 9 further indicates that unWISE roughly doubles the number of $0 < z < 1$ galaxies detected, while increasing the number of galaxies with $1 < z < 2$ by

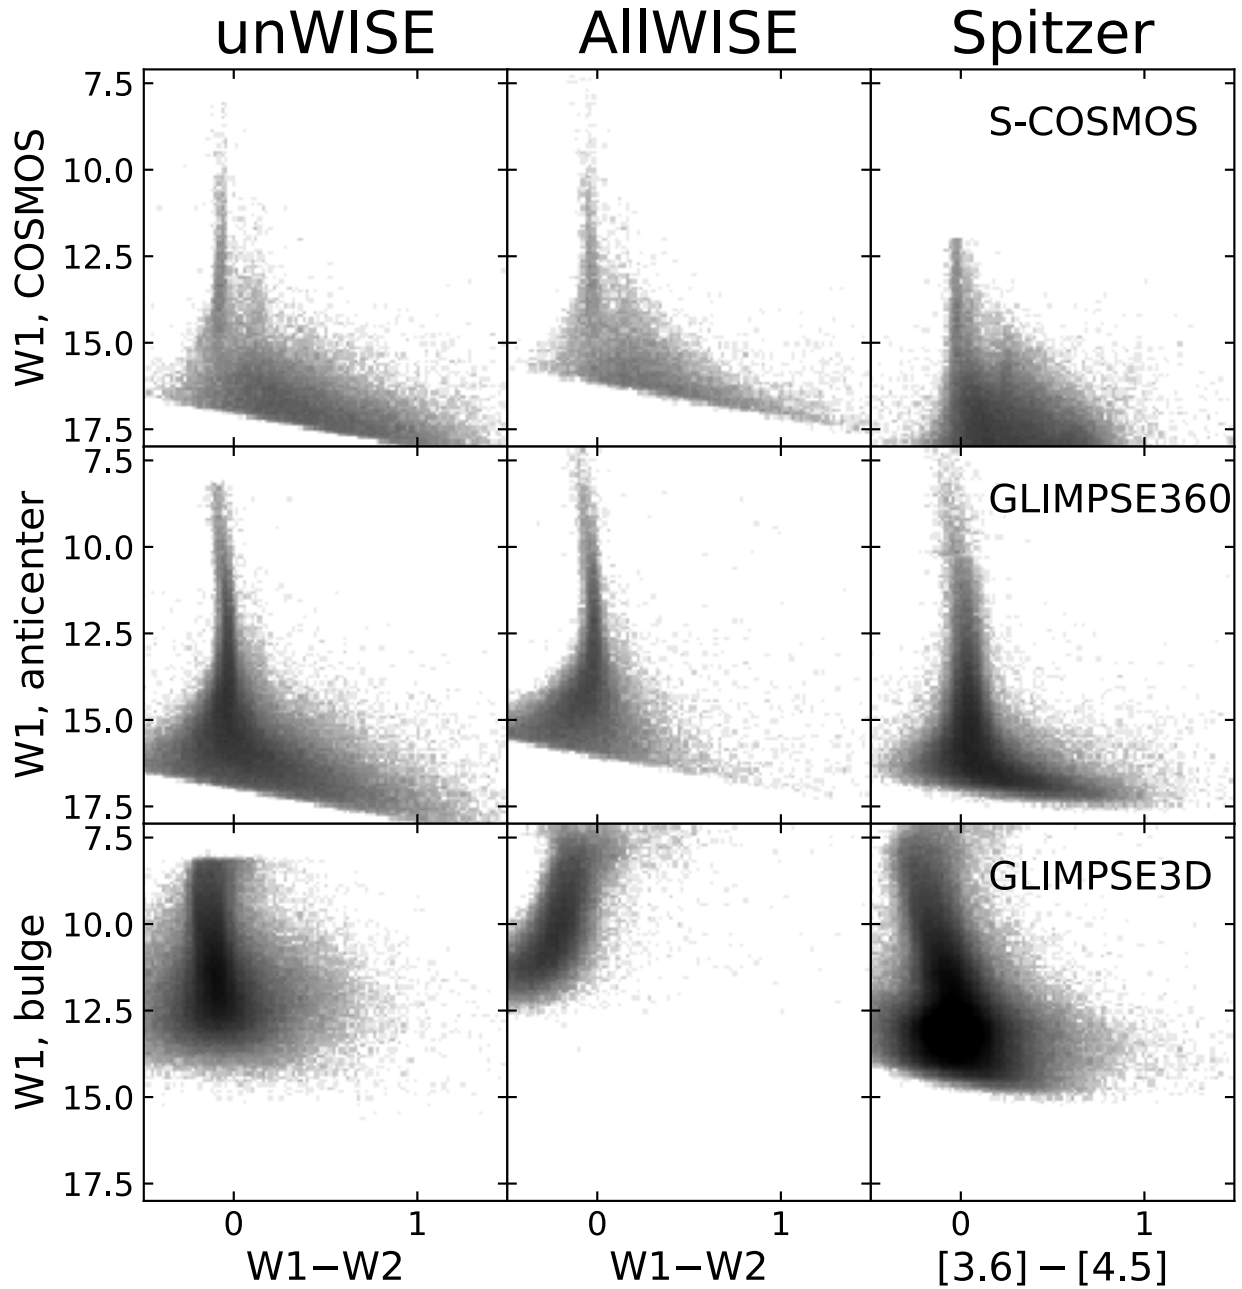


Figure 8. Color-magnitude diagram of 5σ sources from unWISE, AllWISE, and Spitzer in three different fields: the high-latitude COSMOS field, the Galactic anticenter, and the Galactic bulge. The greater depth of the unWISE catalog relative to AllWISE is immediately apparent. In relatively uncrowded high-latitude fields, unWISE extends roughly 0.7 mag fainter than AllWISE, as expected from §5.1. This continues to be the case in the Galactic anticenter. In the Galactic bulge, both the unWISE and AllWISE color-magnitude diagrams clearly suffer from crowding, with broad color distributions even at bright magnitudes. Comparison with much higher resolution Spitzer observations from S-COSMOS, GLIMPSE360, and GLIMPSE3D show good agreement with the unWISE magnitudes. The deep Spitzer observations on the COSMOS field are naturally much deeper than unWISE, but unWISE is competitive with GLIMPSE in depth in the outer Galaxy, while providing a tighter stellar locus. In the inner Galaxy, the higher-resolution GLIMPSE observations allow for deeper catalogs.

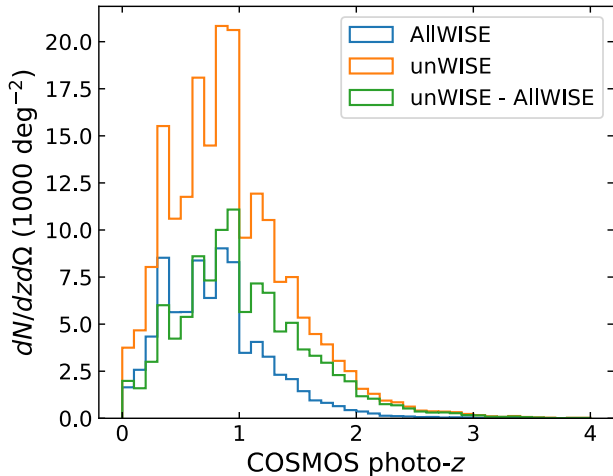


Figure 9. Redshift distribution of galaxies in the COSMOS field detected in AllWISE (blue), unWISE (orange), and the difference (green). The unWISE catalog roughly doubles the number of galaxies detected with $z < 1$, while tripling it at $z > 1$. Extrapolating to the entire sky, the unWISE catalog should contain > 500 million galaxies broadly distributed over $0 < z < 2$.

Table 1. Number density of objects at different redshifts

Catalog	stars		z range				
	Gaia	-Gaia	0, 0.5	0.5, 1	1, 1.5	1.5, 2	> 2
AllWISE	2977	411	2151	3760	1501	411	126
unWISE	4562	1297	3941	8457	4598	1876	773

NOTE—Number of objects per square degree for stars and galaxies of different redshifts, based on comparison to objects with photometric redshifts in COSMOS (Laigle et al. 2016). unWISE increases the number of galaxies detected by a factor of 2–4. Stars are marked as having been identified by Gaia (Gaia), or not (-Gaia). Note that counts are given only for objects matching objects in (Laigle et al. 2016), but roughly 6% of unWISE objects have no match, primarily due to masked regions near bright stars in COSMOS and large galaxies split into multiple PSF components in unWISE.

a factor of three or more. Extrapolated over the whole sky, the unWISE catalog should contain $> 500 \times 10^6$ galaxies with $0 < z < 2$. Table 1 summarizes the number densities of sources of different types and redshifts in the COSMOS field for AllWISE and unWISE.

For extragalactic purposes, the presence of stars in the catalog can be a nuisance. Often, however, these stars can be identified by their pointlike morphology in Gaia imaging, which like WISE, is available for the entire sky.

Table 1 indicates the number density of stars detected by Gaia, and the number density not detected by Gaia, usually due to faint magnitudes and red colors.

From the unWISE Catalog alone, the only information available about a typical galaxy is its flux in the W1 and W2 bands. This makes efforts to estimate a galaxy’s redshift from its unWISE Catalog entry challenging. Nevertheless, there is a good correlation between the WISE color of a galaxy and its redshift. Figure 10 shows the redshift distribution of unWISE Catalog galaxies in the COSMOS field satisfying four simple color cuts. The galaxies passing these cuts have mean redshifts steadily increasing from $z = 0.4$ to $z = 1.5$, with a typical rms of 0.4, as detailed in Table 2.

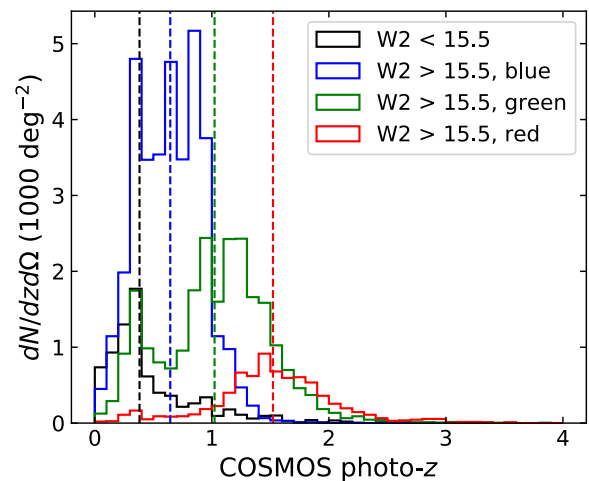


Figure 10. Redshift distribution of galaxies in the COSMOS field detected in unWISE, satisfying four different color and magnitude selections. Judicious cuts on galaxies’ WISE colors can produce samples with mean redshifts ranging from 0.4 to 1.5. Vertical lines give the mean redshifts of the different selections.

5.4.3. High-redshift

Mid-infrared colors provide an efficient means of selecting quasars, making them effective for detecting objects at high redshifts (Wang et al. 2016). By providing deep mid-infrared photometry, the unWISE Catalog should prove valuable in searches for the highest redshift quasars.

Consistent with this expectation, the unWISE catalog contains detections of more $z > 5$ quasars than AllWISE. Among the 453 quasars currently known (Ross & Cross, in prep.), 268 are detected in W1 and 183 are detected in W2 in AllWISE, where “detection” means that the catalog contains a source within $2.75''$ of the

Table 2. Color cuts for different WISE galaxy samples

W2	W1–W2 > x	W1–W2 < x	\bar{z}	δz
< 15.5			0.4	0.3
> 15.5		$(17 - W2)/4 + 0.3$	0.6	0.3
> 15.5	$(17 - W2)/4 + 0.3$	$(17 - W2)/4 + 0.8$	1.0	0.4
> 15.5	$(17 - W2)/4 + 0.8$		1.5	0.4

NOTE—Color and magnitude cuts for selecting galaxies of different redshifts, together with the mean redshift \bar{z} and the width of the redshift distribution δz for the selections, as measured by matching to objects with photometric redshifts on the COSMOS field (Laigle et al. 2016).

QSO with $> 5\sigma$ significance. Meanwhile, 355 and 307 are detected in unWISE in W1 and W2, respectively. Roughly half of all high-redshift quasars formerly undetected in WISE now have secure detections.

6. LIMITATIONS AND FUTURE DIRECTIONS

The unWISE Catalog is a first attempt to measure the fluxes and positions of all of the sources in the unWISE coadds. During the construction of the catalog, a number of issues were identified that could be addressed in future unWISE catalogs. We list some limitations of the catalog here, and in many cases describe improvements that could eliminate these limitations in future releases.

6.1. Non-linearity

Comparison of AllWISE and unWISE magnitudes reveals a slight non-linearity in W1, as evident in Figure 7. It is not clear what the source of this trend is. Variation in the non-linearity of the WISE detector pixels between the mission phases could account for some of the effect. Inspection of the shape of the PSF as a function of magnitude (as measured by `spread_model`) shows a significant trend in both W1 and W2, large enough to account for the trend in W1. However, because the shape-dependent trend is present in both W1 and W2, the obvious avenues for eliminating it introduce a significant trend between AllWISE and unWISE in W2. No correction for non-linearity is made in the unWISE catalog, incurring systematic uncertainties of ≈ 0.02 mag. More work is needed to identify the source of the non-linearity and to develop effective mitigation strategies.

6.2. Saturation

The performance of the unWISE Catalog rapidly degrades at the onset of saturation, as indicated in Figure 7. The unWISE coadd images (§3) use a 7×7 Lanczos kernel to resample the native WISE images onto the

unWISE coadds. Before resampling, any masked pixels are first “patched” by replacing the masked value with the mean of the surrounding pixels. This procedure tends to make the peaks of stars flatter than they would otherwise have been, changing the shape of the PSF. This changed shape is then propagated out to the full 7×7 pixel surrounding neighborhood, albeit with decreasing influence toward the edges.

Consistent with this, saturated stars show significantly worse residuals than unsaturated stars in the unWISE imaging, even outside the saturated center. To mitigate this effect, all pixels within a three pixel radius of a potentially saturated pixel in unWISE are masked in the analysis. Because roughly 90% of the flux of a star lands within 3 pixels of a star’s center, the unWISE catalog fluxes of these bright stars are highly uncertain and dependent on the amount of flux in the wings of the PSF.

Significant improvement here could come from improving the “patching” process for saturated pixels in the unWISE coadds by incorporating knowledge of the WISE PSF.

6.3. Sky Subtraction

The `crowdsourc` sky subtraction analysis removes the median residual in each 19×19 pixel region of the image during each iteration of its fitting process. In very dense regions like the Galactic bulge this process tends to be biased high, and many potential sources are not discovered and analyzed.

The sky could instead be fit as an additional set of linear parameters in the main `crowdsourc` optimization. In the DECam Plane Survey (Schlafly et al. 2018), `crowdsourc` was run in a mode in which a single global sky was simultaneously optimized, but in unWISE we found that this tended to propagate small residuals around bright stars into global sky errors that interacted with the median sky filtering to slow convergence. The locality of the sky fit could be preserved with simultaneous fitting of the sky by adopting a small-scale cardinal basis spline approach to the sky modeling. This approach would require somewhat more memory than the existing analysis, but should improve the speed of convergence and the accuracy of the sky fits.

6.4. Inconsistency of W1 & W2 Modeling

Sources in the unWISE Catalog are modeled completely independently in different bands, in contrast to AllWISE. This means that images that are modeled with three stars in W1 may be modeled with two stars in W2, confusing the linking of W1 and W2 catalogs into multiband object lists. Similarly, a single, isolated object will have slightly different positions in its W1 and

W2 analysis in the unWISE Catalog, again in contrast to AllWISE.

The AllWISE approach has many advantages over the approach taken for unWISE; the primary motivation for the unWISE approach was computational and algorithmic convenience. Future releases could straightforwardly adopt an AllWISE-like simultaneous modeling scheme. Nevertheless, color-magnitude diagrams like Figure 8 indicate that the negative effects of the inconsistent modeling are not large for typical stars.

6.5. No source motions or variability

The unWISE Catalog is strictly a static-sky catalog; it does not fit for the motions of sources or for their variability. Time-resolved unWISE coadds (Meisner et al. 2018b) preserve almost all of the information present in the WISE System. In principle, analysis of these images could recover proper motions for many stars. Like §6.4, this would require simultaneously modeling several images, but beyond this, the generally very small, subpixel motions would be naturally accommodated in the sparse linear algebra analysis at the core of `crowdsourc`.

Similarly, the unWISE Catalog contains no variability information; this has likewise been lost in the construction of the deep static-sky unWISE coadds we have analyzed. Variability on time scales of 0.5–8 years, however, would be accessible through analysis of the time-resolved unWISE coadds.

6.6. Astrometry

The unWISE coadds adopt the world-coordinate system of the underlying WISE single-exposure images without modification. The WISE single-exposure world-coordinate system is correct in the context of a slightly asymmetric PSF model⁷ that is different from the PSF model adopted for the unWISE Catalog. The resulting inconsistency leads to the ≈ 25 milliarcsecond residuals throughout the upper two rows of Figure 5.

The unWISE coadds were made from single-epoch images that did not include MFPRex WCS improvements. Inclusion of these improvements and correction for the different AllWISE and unWISE PSF conventions would remove the dominant sources of coherent astrometric residuals in unWISE.

Another option for improving the unWISE Catalog’s astrometric accuracy would be to recalibrate either the

single-exposure WISE image astrometry or the unWISE coadd astrometry, using a procedure similar to that described in Meisner et al. (2018b). The release of unprecedentedly accurate astrometry and proper motions from Gaia should enable substantial improvements to the native WISE single-exposure astrometry (currently tied to 2MASS, in some cases using UCAC4 proper motions).

6.7. Bright Stars

The WISE PSF extends very far from the source; flux is readily detected 2° from the centers of the brightest stars. The PSF used for the unWISE analysis extends “only” ~ 150 pixels ($\sim 0.1^\circ$) from the centers of stars. Moreover, the wings of the unWISE PSF are very uncertain. Accordingly, there are significant residuals $\approx 0.1^\circ$ from very bright stars that may present as spurious sources in the unWISE Catalog. Closer to the center of extremely bright stars, large saturated regions likewise occasionally lead to the generation of spurious sources in the catalog.

Diffraction spikes around very bright stars pose an additional challenge. The interpolation process from the original images onto new tile centers makes the detailed path a diffraction spike should follow in the coadd images subtle. Far from the center of a star, the unWISE PSF diffraction spikes tend to fall slightly off the true diffraction spikes, leading to significant residuals and affecting the flux of stars in the vicinity. Spurious sources may also be detected in these locations.

The unWISE images contain a number of flags indicating pixels affected by bright stars; these are included in the unWISE catalog in the `flags_unwise` column (§7). Improved modeling of the wings of very bright stars is also possible, but would require significant effort while improving the analysis in only a tiny fraction of the sky.

6.8. Extended Sources

All sources in the unWISE analysis are modeled as if they were point sources. Fluxes and locations of extended sources will be correspondingly biased and sub-optimal. However, the broad WISE PSF of $6''$ FWHM means that most extended sources can be treated as point sources. For example, a $1''$ FWHM galaxy would only broaden the PSF by 1.4%, corresponding to a flux estimate 1.4% too small in the unWISE catalog.

Some galaxies, however, are much larger, and are readily resolved by WISE. Absent intervention, `crowdsourc` will split these galaxies into several point-source components. Many of these large galaxies have already been identified in other surveys, however, making it easy to modify the behavior of `crowdsourc` in these regions. The unWISE mask images include a bit

⁷ See <http://wise2.ipac.caltech.edu/docs/release/neowise/expsup/sec3.2.html#astrom> and §4.1 of Meisner et al. (2018b) for in-depth discussion of this single-exposure PSF model asymmetry.

indicating that a pixel is contained in a large galaxy, as tabulated in the HyperLeda catalog (Makarov et al. 2014). In such pixels, `crowdsourc` rejects candidate new sources if they significantly overlap with a neighboring source.

The effectiveness of this mitigation strategy depends on the underlying catalog of large galaxies. The HyperLeda catalog is rather heterogeneous; over the Sloan Digital Sky Survey (York et al. 2000) footprint it contains many more large galaxies than in the southern sky, for instance. The unWISE catalog will correspondingly spuriously split varying numbers of large galaxies into multiple point sources depending on sky location.

This strategy discourages `crowdsourc` from splitting large galaxies, but significant residuals are still left on the images around large galaxies because the point source model is a poor fit. As part of the optimization process, `crowdsourc` may still decide that the image is best modeled by slowly moving nearby stars into the large galaxy to reduce these residuals, leading to occasional splitting even of large galaxies identified in HyperLeda.

Explicitly modeling all of the sources as potentially resolved galaxies or unresolved stars would be the best approach, but presents substantial conceptual, algorithmic, and computational challenges. Alternatively, full galaxy fitting could be enabled exclusively for objects marked in external catalogs as being resolved galaxies; this would improve the fits of these galaxies and their neighbors. Galaxy models tend to be less well approximated by linear models, however, so their optimization does not fit neatly into the `crowdsourc` framework.

6.9. *Ecliptic Poles*

The unWISE coadd images are qualitatively different at high ecliptic latitudes than at low ecliptic latitudes. At low ecliptic latitudes, the typical number of WISE observations contributing to a given part of the sky in the unWISE coadds is ~ 120 per band. At the ecliptic poles, however, the number is greater than 23000, $\sim 200\times$ larger. Moreover, due to the WISE scan strategy, at low ecliptic latitudes, the position angle of the WISE focal plane on the sky is nearly constant, while the ecliptic poles were observed at all position angles.

The huge number of WISE observations at the ecliptic pole potentially make for images ~ 3 magnitudes deeper than typical at low latitudes. However, at this depth, large-scale residuals in the unWISE coadd sky become significant. Absent intervention, the `crowdsourc` analysis would incorrectly attempt to explain what remains of these residuals after sky subtraction with many sources. To mitigate this, a flux uncertainty floor was

added to the unWISE uncertainty images such that the 5σ sources correspond to at most ~ 19.8 mag in W1 and ~ 18.2 mag in W2. These are roughly 1.5 mag fainter than the typical depths of the unWISE coadd images (§5.1), but much brighter than the nominal depth achievable at the ecliptic poles.

The varying position angle of the WISE images taken near the ecliptic pole leads the PSF there to be different from the PSF at lower ecliptic latitudes. As discussed in §4.1, the `crowdsourc` analysis uses a dense grid of rotated and azimuthally smoothed PSF models to capture much of the PSF variation in the unWISE coadds in these areas. However, our flagging of diffraction spikes and ghosts in the unWISE coadd masks was not modified to account for the azimuthal smoothing of the PSF near the ecliptic poles. Accordingly, these features in the unWISE masks near the ecliptic pole are narrower than they should be.

This problem primarily affects regions within about 5° of the ecliptic poles, $\sim 0.4\%$ of the footprint. Moreover, because sharp features like diffraction spikes spread over a wider area due to this effect, their amplitude is lower and their masking is less important. Directly on the pole, for instance, it is not clear if a diffraction spike mask would be desired; the diffraction spikes have been fully blurred out into the wings of the PSF.

6.10. *Nebulosity*

The `crowdsourc` analysis identifies regions believed to contain significant nebulosity and requires that new sources detected in these regions be relatively sharp and PSF-like; sources significantly broadened by neighboring sources will not be modeled. This process occasionally incorrectly identifies features in the wings of bright stars as nebulous, leading to fewer detected sources in these regions. Only roughly 0.7% of the footprint is identified as being affected by nebulosity, however, so this is a small effect. Likewise, some areas containing significant nebulosity are not marked and may contain many spurious sources.

6.11. *Uncertainty estimates*

The unWISE Catalog analysis adopts the unWISE coadd uncertainty images with little alteration, except near the ecliptic poles (§6.9). Analysis of the coadd residual images after subtraction of the best fit models suggests that the coadd uncertainty images overestimate the actual uncertainty by 15% in W1 and 20% in W2. On the other hand, the residual images show significant correlated uncertainties, which are neglected in the analysis. If the correlated errors in the residuals could be eliminated or modeled out, and the uncertainty images modified to better describe the true dispersion in

the values, WISE catalogs could reliably detect objects ~ 0.2 mag fainter than at present.

7. DATA RELEASE

The unWISE coadds comprise 18240 tiles in two bands. The corresponding 36480 catalog FITS files (Pence et al. 2010) composing the unWISE Catalog are available at the catalog web site, <http://catalog.unwise.me>. The Catalog consists of 2.03 billion detections of objects in the “primary” regions of their unWISE coadds that have at least 5σ significance in W1 or W2. Some basic numbers describing the catalog are given in Table 3.

The content of the FITS files is essentially identical to the files of the DECam Plane Survey (Schlafly et al. 2018), with the addition of a few metadata columns and `spread_model` (§4.5). The catalog columns are listed in Table 4.

Fluxes and corresponding uncertainties are given in linear flux units, specifically, in Vega nanomaggies (nMgy). The corresponding Vega magnitudes are given by $m_{\text{Vega}} = 22.5 - 2.5 \log_{10} \text{flux}$. The following equations give the corresponding AB magnitudes:

$$\begin{aligned} m_{\text{W1, AB}} &= m_{\text{W1, Vega}} + 2.699 \\ m_{\text{W2, AB}} &= m_{\text{W2, Vega}} + 3.339. \end{aligned}$$

As noted in §5.3, the agreement between unWISE and ALLWISE magnitudes can be improved by subtracting 4 mmag and 32 mmag from W1 and W2.

Additional files give the `crowdsourcing` model image and sky image for each unWISE tile. The PSF flux inverse variance image, mask image, and PSF model are also available for each tile.

The column `fwhm` is intended to give a sense of the size of the model PSF for a particular detection. Given a PSF model, it is computed as the FWHM a Gaussian PSF with equal n_{eff} would have. Typical values of $7.2''$ in W1 and $7.8''$ in W2 are larger than the true WISE FWHMs because the WISE PSF has more flux in its wings than a Gaussian, increasing n_{eff} .

The column `primary` marks whether a particular source is located in the “primary” region of its coadd. The unWISE coadds overlap one another by roughly 60 pixels, so that sources residing on the edges of an unWISE coadd will be detected in multiple coadds. By selecting only “primary” sources, duplicate sources can be eliminated. Determining whether a source is primary in a given tile is purely a geometric operation, and does not involve any cross-matching of detections on neighboring tiles. For each source in a given tile’s catalog, we compute the source’s minimum distance from any edge of that tile’s footprint. Using the source’s (RA, Dec)

coordinates, we perform the same minimum edge distance computation for all neighboring tiles. The source is labeled primary in the tile under consideration if that tile’s footprint provides a larger minimum distance to any edge than do all other neighboring tile footprints.

Finally, merged catalogs linking W1 detections and W2 detections into multiband objects are also available. Each W1 source is matched to the nearest W2 source within $2.4''$; W2 sources within $2.4''$ that are not the closest source to a W1 source are considered unmatched. The merged catalogs include the same columns as in the individual catalogs (Table 4), but each column now contains a two element vector for the W1 and W2 quantities. The canonical right ascension and declination are taken to be the W1 quantities, when available, and otherwise the W2 quantities. Likewise an object is considered “primary” when its W1 detection is considered “primary.” Finally, a unique `unwise_objid` is assigned to each entry in the merged catalog. Unmatched detections contain zeros in all columns corresponding to the missing band; these are easily identified, for example, by the empty `unwise_detid`.

In addition to the catalogs themselves, the unWISE Catalog release contains a few items intended to facilitate the catalog’s use. These are

- model images,
- model sky images,
- PSF depth images,
- mask images, and
- PSF images.

A sense for the plausibility of the modeling of a particular object in the unWISE Catalog can be obtained by comparing the unWISE coadds with the model images. The unWISE model sky images and images of the PSF are potentially useful for users seeking to do their own photometry on the unWISE coadds; estimating the sky and PSF in crowded fields can be challenging. Users wondering about spatial variations in the depth of the survey may find the PSF depth images valuable.

We also provide mask images that replicate some of the elements in the unWISE coadd mask images, but also add a few elements specific to the unWISE Catalog processing. The values of the mask image are given in Table 5. In particular, the mask images indicate sky regions in which candidate sources significantly overlapping other sources are not modeled, due to the presence of a nearby large galaxy (§4.4). They also indicate which parts of the sky the convolutional neural network indicates to be affected by significant nebulosity (§4.4).

Table 3. Number of Sources in the unWISE Catalog

type	W1 & W2	W1 only	W2 only	total
all	1,063,569,639	1,032,419,887	118,744,698	2,214,734,224
primary	979,399,857	949,169,093	108,858,251	2,037,427,201
> 5 σ	1,062,021,630	1,027,533,192	116,683,978	2,206,238,800
primary & > 5 σ	978,015,749	944,811,706	106,974,415	2,029,801,870

NOTE—Number of objects in the unWISE Catalog satisfying various criteria. “primary” sources discard duplicate sources in sky regions included in multiple coadds. > 5 σ sources are detected with at least 5 σ significance.

Table 4. Catalog Columns

Name	Description
ra	right ascension (deg)
dec	declination (deg)
x	x coordinate (pix)
y	y coordinate (pix)
flux	instrumental flux (nMgy)
dx	x uncertainty (pix)
dy	y uncertainty (pix)
dflux	flux uncertainty (nMgy)
fluxlbs	local-background-subtracted flux (nMgy)
dfluxlbs	fluxlbs uncertainty (nMgy)
qf	PSF-weighted fraction of good pixels
rchi2	PSF-weighted average χ^2
fracflux	PSF-weighted fraction of flux from this source
spread_model	SExtractor-like source size parameter
dspread_model	uncertainty in spread_model
fwhm	FWHM of PSF at source location (pix)
sky	residual sky at source location (nMgy)
nm	number of images in coadd at source
primary	source located in primary region of coadd
flags_unwise	unWISE flags at source location
flags_info	additional flags at source location
coadd_id	unWISE coadd_id of source
band	1 for W1, 2 for W2
unwise_detid	detection ID, unique in catalog

NOTE—Columns in the unWISE catalogs. A more complete description is available at the survey web site.

7.1. Source Designations

We prescribe that unWISE Catalog source designations contain the prefix “WISEU”. For example, WISEU J112234.53+122954.3 refers to the object with UNWISE_OBJID = 1699p121o0017067.

Table 5. unWISE Catalog Info Flags

Name	Bit	Description
bright_off_edge	2 ⁰	bright source off coadd edge
resolved_galaxy	2 ¹	in large galaxy in HyperLeda
big_object	2 ²	in M31 or Magellanic Cloud
bright_star_cen	2 ³	may contain bright star center
crowdsat	2 ⁴	may be affected by saturation
nebulousity	2 ⁵	nebulousity may be present
nodeblend	2 ⁶	deblending discouraged here
sharp	2 ⁷	only “sharp” sources here

NOTE—Informational flags in the unWISE Catalogs. A more complete description is available at the survey web site.

8. CONCLUSION

The continuing NEOWISE-Reactivation mission has provided more than four years of imaging beyond the initial year of WISE data that was available for AllWISE processing. The unWISE project has combined the ~ 25 million single-frame WISE images into coadds reaching $2\times$ deeper than AllWISE. The unWISE Catalog is the result of the analysis of these coadds. It contains ~ 2 billion sources, roughly $3\times$ as many as cataloged in AllWISE. Because of the broad WISE PSF, only a small fraction of extragalactic sources are resolved in WISE, making the analysis ideally suited to crowded-field pipelines that aggressively model images as sums of many overlapping point sources. Application of the **crowdsat** crowded-field image analysis to the unWISE coadds provides accurate measurements for the unWISE Catalog, both in extragalactic fields where the improved depth of unWISE is critical, and also in the Galactic bulge where the analysis is limited by crowding.

Comparison between bright sources in the unWISE and AllWISE catalogs shows good agreement between

the two catalogs. For faint sources, comparison with deeper Spitzer imaging confirms that the catalog reaches $2\times$ deeper than AllWISE. The unWISE Catalog reaches stars in the Milky Way at greater distances, detects hundreds of millions of new galaxies over $0 < z < 2$, and finds half of high-redshift quasars undetected in AllWISE.

We have outlined possible ways in which future versions of the unWISE Catalog may be able to augment or improve upon the present data products. Incorporating future NEOWISE data releases would enable the unWISE Catalog to push yet deeper. In combination with the complementary WISE-based proper motions that will be supplied by CatWISE (PI: Eisenhardt), the unWISE Catalog realizes much of the NEOWISE data set’s tremendous potential for Galactic and extragalactic astrophysics.

unWISE coadd images, the derived unWISE catalog, the corresponding model PSF, sky images, and depth maps are publicly available at the unWISE web site, <http://unwise.me>.

It is a pleasure to thank Roc Cutri for detailed feedback and John Moustakas for help with the HyperLeda catalog and the Legacy Survey Large Galaxy Atlas. David Schlegel provided invaluable guidance and motivation, and Dustin Lang built much of the framework on which this catalog stands.

This work has been supported in part by NASA ADAP grant NNH17AE75I. ES and AMM acknowledge

support for this work provided by NASA through Hubble Fellowship grants HST-HF2-51367.001-A and HST-HF2-51415.001-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 NAS 5-26555.

We acknowledge the usage of the HyperLeda database (<http://leda.univ-lyon1.fr>).

This publication makes use of data products from the Wide-field Infrared Survey Explorer, which is a joint project of the University of California, Los Angeles, and the Jet Propulsion Laboratory/California Institute of Technology, and NEOWISE, which is a project of the Jet Propulsion Laboratory/California Institute of Technology. WISE and NEOWISE are funded by the National Aeronautics and Space Administration.

The unWISE Catalog analysis was run on the Odyssey cluster supported by the FAS Division of Science, Research Computing Group at Harvard University, and on the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. The unWISE nebulosity CNN was trained on the XStream computational resource, supported by the National Science Foundation Major Research Instrumentation program (ACI-1429830).

APPENDIX

A. NEBULOSITY NEURAL NETWORK STRUCTURE

Table 6 summarizes the architecture of the convolutional neural network that detects nebulosity in images. Because the input images are a quarter the size of those used in Schlafly et al. (2018), this network is slightly shallower.

We trained the network on three classes: `normal`, `nebulosity` and `nebulosity_light`. We labeled 10,389 images by hand, using 80% for training and the remaining 20% for validation. After training, our network was able to separate cleanly between images labeled `normal` and `nebulosity`, with less than 0.1% of images labeled `normal` being labeled `nebulosity` and vice versa. As the distinction between `nebulosity` and `nebulosity_light` is more subtle (and likely less consistent in our hand classifications), our network misclassified 12% of the images we considered `nebulosity` as `nebulosity_light`.

Facilities: WISE

Software: `astropy` (Astropy Collaboration et al. 2013)

REFERENCES

- Antilogus, P., Astier, P., Doherty, P., Guyonnet, A., & Regnault, N. 2014, *Journal of Instrumentation*, 9, C03048
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33

Table 6. Network architecture

layer	output shape	details
conv2d_1	256 × 256, 12	$\begin{pmatrix} 3 \times 3, \text{ same padded} \\ 3 \times 3, \text{ same padded} \\ 1 \times 1 \end{pmatrix}$
maxpool2d_1	128 × 128, 12	2 × 2
conv2d_2	128 × 128, 16	$\begin{pmatrix} 3 \times 3, \text{ same padded} \\ 3 \times 3, \text{ same padded} \\ 1 \times 1 \end{pmatrix}$
maxpool2d_2	64 × 64, 16	2 × 2
conv2d_3	64 × 64, 24	$\begin{pmatrix} 3 \times 3, \text{ same padded} \\ 3 \times 3, \text{ same padded} \\ 1 \times 1 \end{pmatrix}$
maxpool2d_3	32 × 32, 24	2 × 2
conv2d_4	32 × 32, 32	$\begin{pmatrix} 3 \times 3, \text{ same padded} \\ 3 \times 3, \text{ same padded} \\ 1 \times 1 \end{pmatrix}$
maxpool2d_4	16 × 16, 32	2 × 2
conv2d_5	16 × 16, 32	$\begin{pmatrix} 3 \times 3, \text{ same padded} \\ 3 \times 3, \text{ same padded} \\ 1 \times 1 \end{pmatrix}$
maxpool2d_5	8 × 8, 32	2 × 2
global_avg_pool2d	32	
dense_1	9	20% dropout
dense_2	3	10% dropout
softmax	3	

NOTE—All convolutional and dense layers use ReLU activation, and have an L2 weight penalty of 10^{-4} . The final output one-hot encodes the class, and the cross-entropy loss function is used.

- Bañados, E., Venemans, B. P., Mazzucchelli, C., et al. 2018, *Nature*, 553, 473
- Benjamin, R. A., Churchwell, E., Babler, B. L., et al. 2003, *PASP*, 115, 953
- Bertin, E., & Arnouts, S. 1996, *A&AS*, 117, 393
- Churchwell, E., Babler, B. L., Meade, M. R., et al. 2009, *PASP*, 121, 213
- Cutri, R. M., Wright, E. L., Conrow, T., et al. 2013, Explanatory Supplement to the AllWISE Data Release Products, Tech. rep.
- Cutri, R. M., Mainzer, A., Conrow, T., et al. 2015, Explanatory Supplement to the NEOWISE Data Release Products, Tech. rep.
- Cypriano, E. S., Amara, A., Voigt, L. M., et al. 2010, *Monthly Notices of the Royal Astronomical Society*, 405, 494. <http://dx.doi.org/10.1111/j.1365-2966.2010.16461.x>
- Desai, S., Armstrong, R., Mohr, J. J., et al. 2012, *ApJ*, 757, 83
- Dey, A., Schlegel, D. J., Lang, D., et al. 2018, ArXiv e-prints, arXiv:1804.08657
- Downing, M., Baade, D., Sinclair, P., Deiries, S., & Christen, F. 2006, in *Proc. SPIE, Vol. 6276, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, 627609
- Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, *A&A*, 595, A1
- Kirkpatrick, J. D., Cushing, M. C., Gelino, C. R., et al. 2011, *ApJS*, 197, 19
- Laigle, C., McCracken, H. J., Ilbert, O., et al. 2016, *ApJS*, 224, 24
- Lang, D. 2014, *AJ*, 147, 108
- Lang, D., Hogg, D. W., & Schlegel, D. J. 2014, ArXiv e-prints, arXiv:1410.7397

- Low, F. J., Rieke, G. H., & Gehrz, R. D. 2007, *ARA&A*, 45, 43
- Mainzer, A., Bauer, J., Grav, T., et al. 2011, *ApJ*, 731, 53
- Mainzer, A., Bauer, J., Cutri, R. M., et al. 2014, *ApJ*, 792, 30
- Makarov, D., Prugniel, P., Terekhova, N., Courtois, H., & Vauglin, I. 2014, *A&A*, 570, A13
- Meisner, A. M., & Finkbeiner, D. P. 2014, *ApJ*, 781, 5
- Meisner, A. M., Lang, D., & Schlegel, D. J. 2017a, *AJ*, 153, 38
- . 2017b, *AJ*, 154, 161
- . 2018a, *Research Notes of the American Astronomical Society*, 2, 1
- . 2018b, *AJ*, 156, 69
- Meisner, A. M., Lang, D. A., & Schlegel, D. J. 2018c, *Research Notes of the American Astronomical Society*, 2, 202
- Pence, W. D., Chiappetti, L., Page, C. G., Shaw, R. A., & Stobie, E. 2010, *A&A*, 524, A42
- Sanders, D. B., Salvato, M., Aussel, H., et al. 2007a, *ApJS*, 172, 86
- . 2007b, *ApJS*, 172, 86
- Schlafly, E. F., Peek, J. E. G., Finkbeiner, D. P., & Green, G. M. 2017, *ApJ*, 838, 36
- Schlafly, E. F., Meisner, A. M., Stutz, A. M., et al. 2016, *ApJ*, 821, 78
- Schlafly, E. F., Green, G. M., Lang, D., et al. 2018, *ApJS*, 234, 39
- Schlegel, D. J., Blum, R. D., Castander, F. J., et al. 2015, in *American Astronomical Society Meeting Abstracts*, Vol. 225, American Astronomical Society Meeting Abstracts #225, 336.07
- Wang, F., Wu, X.-B., Fan, X., et al. 2016, *ApJ*, 819, 24
- Wheelock, S. L., Gautier, T. N., Chillemi, J., et al. 1994, *NASA STI/Recon Technical Report N*, 95
- Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *AJ*, 140, 1868
- York, D. G., Adelman, J., Anderson, Jr., J. E., et al. 2000, *AJ*, 120, 1579