



The DECam Plane Survey: Optical Photometry of Two Billion Objects in the Southern Galactic Plane

E. F. Schlafly^{1,8}, G. M. Green², D. Lang^{3,4}, T. Daylan⁵, D. P. Finkbeiner^{5,6}, A. Lee⁵, A. M. Meisner¹,
D. Schlegel¹, and F. Valdes⁷

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

² Kavli Institute for Particle Astrophysics and Cosmology, Physics and Astrophysics Building, 452 Lomita Mall, Stanford, CA 94305, USA

³ Department of Astronomy & Astrophysics and Dunlap Institute, University of Toronto, 50 Saint George Street, Toronto, ON, M5S 3H4, Canada

⁴ Department of Physics & Astronomy, University of Waterloo, 200 University Avenue West, Waterloo, ON, N2L 3G1, Canada

⁵ Department of Physics, Harvard University, 17 Oxford Street, Cambridge, MA 02138, USA

⁶ Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA

⁷ National Optical Astronomy Observatories, P.O. Box 26732, Tucson, AZ 85719, USA

Received 2017 October 3; revised 2017 December 12; accepted 2017 December 18; published 2018 February 9

Abstract

The DECam Plane Survey is a five-band optical and near-infrared survey of the southern Galactic plane with the Dark Energy Camera at Cerro Tololo. The survey is designed to reach past the main-sequence turn-off of old populations at the distance of the Galactic center through a reddening $E(B - V)$ of 1.5 mag. Typical single-exposure depths are 23.7, 22.8, 22.3, 21.9, and 21.0 mag (AB) in the *grizY* bands, with seeing around 1". The footprint covers the Galactic plane with $|b| \lesssim 4^\circ$, $5^\circ > l > -120^\circ$. The survey pipeline simultaneously solves for the positions and fluxes of tens of thousands of sources in each image, delivering positions and fluxes of roughly two billion stars with better than 10 mmag precision. Most of these objects are highly reddened and deep in the Galactic disk, probing the structure and properties of the Milky Way and its interstellar medium. The fully-processed images and derived catalogs are publicly available.

Key words: catalogs – surveys – techniques: photometric

1. Introduction

Many of the Milky Way stars and much of its gas and dust reside in a disk. Accordingly, observations of the Milky Way disk are critical to understanding the Milky Way—particularly observations toward the inner Galaxy, where most of the mass lies. At optical wavelengths, however, the interpretation of observations of the Milky Way disk can be challenging due to the tremendous number of stars and due to extinction by dust, motivating surveys of the disk at infrared wavelengths where extinction is greatly reduced.

However, optical observations have advantages over infrared observations. Critically, the intrinsic colors of typical stars depend more strongly on stellar type in the optical than in the infrared: optical colors are often more useful than infrared colors in determining the temperature of a star. Additionally, for studies of the Milky Way dust, the sensitivity of optical colors to extinction is valuable, allowing measurements with higher signal-to-noise ratio (S/N) of the column of dust to stars. Similarly, the optical extinction curve is more sensitive to variation in the properties of dust than is the infrared extinction curve, permitting the variation in the dust extinction curve to be studied throughout the Milky Way disk.

In this work, we present the DECam Plane Survey (DECaPS): *grizY* imaging and derived catalogs of one third of the Milky Way disk, covering roughly $|b| \lesssim 4^\circ$, $5^\circ > l > -120^\circ$, as shown in Figure 1. Observations were made with the Dark Energy Camera (DECam; Flaugher et al. 2015) with the 4 m Blanco telescope at Cerro Tololo, and were designed to reach the main-sequence turn-off of old populations at a distance of 8 kpc through 1.5 mag $E(B - V)$. The final survey catalog contains more than 20 billion

photometric measurements of two billion objects, a rich trove of data for understanding the Milky Way and its stars and dust.

DECaPS occupies a special niche among the many surveys targeting the Milky Way. The most comparable survey is the PS1 survey (Chambers et al. 2016). DECaPS and PS1 use a very similar set of filters and achieve similar depths (DECaPS is roughly 1 mag deeper in individual images), although PS1 covers the entire sky above $\delta = -30^\circ$ and contains more epochs than DECaPS does (12 versus 3). PS1 has been a valuable survey for understanding the structure of the Milky Way disk, both in stars (Slater et al. 2014; Morganson et al. 2016) and dust (Green et al. 2015; Schlafly et al. 2015; Hanson et al. 2016; Schlafly et al. 2016, 2017). DECaPS intends to duplicate the success of PS1 in the southern Galactic plane sky where no PS1 imaging is available, while obtaining greater depths to reach farther across the Milky Way disk.

Other optical surveys targeting the Milky Way disk include IPHAS (Drew et al. 2005) and VPHAS+ (Drew et al. 2014), which are roughly a magnitude shallower than DECaPS. In contrast to DECaPS, however, these surveys include observations in H α (and *u* in VPHAS+), critical for identifying ionized nebulae and unresolved pre- and post-main sequence stars.

Near-infrared surveys like 2MASS (Skrutskie et al. 2006) have long been used to study the Galactic plane, in part because of the reduced sensitivity in the near-infrared to extinction relative to the optical. Dedicated surveys of the Galactic plane reach dramatically fainter magnitudes, for example, the UKIDSS Galactic Plane (Lucas et al. 2008) and Vista Variables in the Via Lactea surveys (Minniti et al. 2010). Finally, space missions such as the all-sky WISE survey (Wright et al. 2010) and the plane-focused Spitzer GLIMPSE surveys (Benjamin et al. 2003; Churchwell et al.

⁸ Hubble Fellow.




Figure 1. DECaPS footprint. The figure shows the number density of sources that are detected in at least three bands and are brighter than 20th mag in r . Dark regions have the most stars. Major dust clouds are prominent, significantly reducing the number of stars observed behind them. Prominent dust clouds include the Pipe Nebula at $(1^\circ, 4^\circ)$ and the Vela Molecular Ridge cloud C at $(-95^\circ, 1^\circ 5')$; Murphy & May 1991). White corresponds to three thousand sources per deg^2 , while black corresponds to three million sources per deg^2 (log scale).

2009) have imaged the entire Galactic plane at longer wavelengths, where dust extinction has a still smaller effect.

The combination of surveys at different wavelengths offers the best approach to disentangling the rich phenomena at play in the inner Galaxy, providing long-wavelength observations that are relatively insensitive to dust and pierce deep into the enshrouded disk of the Milky Way, as well as optical observations that are much more sensitive to both dust and to intrinsic stellar properties, but can be effectively much shallower in regions of high dust column. DECaPS provides optical observations of the southern Galactic plane to complement observations at longer wavelengths to enable new maps of stars and dust in the southern Galactic plane.

This paper is organized into eight sections. In Sections 1 and 2 we introduce DECaPS and its strategy. In Section 3 and Section 4 we describe the pixel-level processing of the images and their reduction into catalogs. Next we describe the photometric calibration in Section 5. The characteristics of the survey are described in Section 6. Information about the data release is provided in Section 7, and we conclude in Section 8.

2. Survey Strategy

The DECam Plane Survey is a survey of the southern Galactic plane in five broadband filters from 400 nm to 1050 nm ($grizY$) in arcsecond seeing, to depths of 23.7, 22.8, 22.3, 21.9, and 21.0 mag, respectively, on the AB magnitude system (Oke & Gunn 1983). The survey covers $\delta < -30^\circ$, $|b| \lesssim 4^\circ$, about 1000 square degrees. More than two billion sources are detected; the spatial distribution of bright sources is shown in Figure 1. When combined with PS1, the two surveys cover the entire Galactic plane with arcsecond-resolution five-band optical photometry. The basic parameters of the survey are given in Table 1.

Table 1
DECam Plane Survey Parameters

Filter	λ	Time	6σ Depth (AB)	Num	FWHM	Sky
g	480	96	23.7	3	1.16	21.4
r	638	30	22.8	3	1.07	20.6
i	777	30	22.3	3	0.97	19.3
z	911	30	21.9	3	0.93	18.7
Y	985	30	21.0	3	0.91	18.0

Note. Basic parameters of DECaPS. The survey delivers photometry from about 400–1050 nm in five broad bands, to depths between 21st and 24th mag with seeings of around an arcsecond. Each part of the footprint was observed roughly three times in photometric conditions. The table gives the filters observed in, their central wavelengths in nm, the typical exposure times, the 6σ limiting magnitudes, the typical number of photometric observations of each part of the footprint, the FWHM in arcseconds ($0''.26$ pixels), and the sky brightness in mag per square arcsecond.

DECam is the best instrument available for this survey. The survey requires both a large aperture and a large field of view to efficiently survey the faint, extinguished stars in the Galactic plane. The Blanco 4 m telescope and DECam imager (2.2° diameter field of view, $0''.26/\text{pixel}$ plate scale) deliver the aperture, field of view, and image quality required to resolve the billions of stars in the often extremely crowded inner Galaxy.

2.1. Depth

DECaPS is designed to measure the fluxes of stars over a substantial fraction of the Milky Way disk. Specifically, the survey targeted a depth reaching below the main-sequence turn-off of a 10 Gyr, solar metallicity population at a distance of eight kiloparsecs. For unextinguished stars, this would correspond to a magnitude of 18.6 in r , but substantial extinction in the Galactic plane requires dramatically deeper




Figure 2. Distribution of individual image depths in the DECaPS survey. Distributions are tight, giving homogeneous data over the survey footprint. Means of the distributions are indicated by vertical dashed lines.

observations. Extinction in the Galactic midplane varies greatly; much of the outer Galaxy has $E(B - V) < 1$ mag, but, especially in the inner Galaxy, extinction can become very high. Through $E(B - V) = 1.5$ mag at 8 kpc, the target stars lie at 24.5, 22.3, 21.2, 20.6, and 20.3 mag (AB) in the *grizY* bands. DECam can reach these magnitudes at 5σ in less than 30 s exposures in the *rizY* bands. In the *g* band, much longer exposures are necessary to reach this depth due to extinction, but 96 s exposures were expected to reach a 5σ depth of 24.1 mag, corresponding to $E(B - V) = 1.4$ mag. Because overheads between exposures with DECam are close to 30 s, it is wasteful to take exposures much shorter than 30 s; accordingly, we adopt a 96 s exposure time in *g* and 30 s in all other bands. Given the conditions actually obtained at the telescope, the survey reached typical 6σ depths of 23.7, 22.8, 22.3, 21.9, and 21.0 mag (AB) in *grizY*. The extra depth in the redder bands improves our distance reach in extinguished regions where *g* photometry, in particular, is not practical. The survey is slightly shallow in individual exposures in *g*, but three images are available of each part of the sky, allowing improved average photometry.

Figure 2 shows the 6σ depths obtained by DECaPS over all exposures taken by the program. The distributions of depths in each filter are tight. A few shallower exposures in each filter are attributable to cloudy weather. The *i*-band depth distribution is bimodal due to significantly darker skies in the second half of the DECaPS observations than in the first.

2.2. Tiling

The survey aims for coverage of the entire $|b| \lesssim 4^\circ$, $\delta < -30^\circ$ sky. A small fraction of the DECam focal plane is not available for imaging, primarily due to gaps between the CCDs. To address this problem, DECaPS adopts a three-pass dither strategy, where each part of the footprint is observed with three different exposures, slightly offset from one another. The survey uses the three-pass tiling strategy developed for the DECam Legacy Survey (DESI Collaboration et al. 2016a, 2016b).

This tiling scheme uniformly covers the sky with pointings separated by distances closely matched to the DECam field of

Table 2
Observing Runs

Start	# Nights	Filters	Notes
2016 Mar 13	4	<i>gr</i>	2016 Mar 17 clouded out
2016 Mar 23	4	<i>izY</i>	2016 Mar 25 clouded out
2016 Aug 10	0.5	<i>izY</i>	Scattered clouds
2016 Aug 14	0.5	<i>izY</i>	Scattered clouds
2016 Aug 15	0.5	<i>izY</i>	
2016 Aug 22	5×0.5	<i>gr</i>	2016 Aug 23 and 25 clouded out
2017 Jan 16	8×0.5	<i>grizY</i>	2017 Jan 22 and 23 clouded out
2017 Apr 19	2	<i>grizY</i>	2017 Apr 19 marginal
2017 Apr 27	2	<i>gr</i>	...
2017 Apr 29	2×0.5	<i>gr</i>	...
2017 May 03	2×0.5	<i>grizY</i>	2016 May 04 cloudy

Note. DECaPS observing runs. The runs included 12 full nights and 20 half nights, most of which were clear. The number of nights in each run and the primary filters observed on those nights are listed, together with notes about any inclement conditions.

view. The field centers of the three passes are slightly rotated relative to one another to dither the exposures around each tile center, filling in the gaps in the DECam focal plane.

The DECaLS tiling scheme happens to contain dithers primarily in the right-ascension direction around right ascensions of 90° and 270° . This leads to a small amount of area without any DECaPS coverage within 10° of the Galactic center.

2.3. Cadence

The DECaPS cadence aims to observe each part of the footprint in photometric conditions on three separate nights. Typically, these three nights fall in one observing run and occur subsequently, but the vagaries of the weather led to a few regions in which two passes were observed immediately after one another on a single night, and others in which repeat observations of the same field were taken about one year apart from one another.

This cadence was designed to enable an internal photometric calibration based on repeated observations of the same stars (e.g., Padmanabhan et al. 2008; Schlafly et al. 2012; Burke et al. 2017). Observing in different passes on different nights then both filled in the gaps in the footprint that would be left by a single pass, and provided the repeat observations needed to constrain the atmospheric throughput as a function of time.

2.4. Observations

We obtained 22 nights on the Blanco 4 m through the NOAO. The observing runs are described in Table 2. Observations were made in *gr* when the Moon was down and *izY* when the Moon was up. The longer exposure time in *g* meant that the survey needed roughly equal amounts of Moon-up and Moon-down time. The efficiency of the DECam system was excellent, enabling observations to keep up with the footprint as it crossed the meridian. The southern location of the footprint and telescope, combined with favorable scheduling, allowed the survey to obtain very low airmasses, with a mean airmass of 1.15 and a standard deviation of 0.10.

Owing to a slight imbalance in the amount of dark and bright time available in good conditions, *r*-band observations in 2017 January were taken with longer 50 s exposures when the Moon

Table 3
Photometric Calibration Performance

Filter	σ_{zp} (mmag)	σ_{bright} (mmag)
g	10.3	7.6
r	10.2	6.9
i	8.6	6.7
z	9.4	6.4
Y	8.3	6.0

Note. The performance of the photometric calibration. The simple nightly zero-point and airmass term described the zero-points of individual images to better than about 1% (σ_{zp}). Bright stars on each image agreed with the mean fluxes of those stars over all observations with a scatter of better than 8 mmag in each band (σ_{bright}).

was up, and in 2017 April/May, some *izY* exposures were taken with the Moon down.

3. Pixel Processing

DECaPS constructs catalogs from calibrated images delivered by the NOAO Community Pipeline (Valdes et al. 2014, CP). The CP transforms raw images into calibrated, science-grade data, performing the following steps (among others):

1. bias subtraction,
2. crosstalk correction,
3. flat fielding,
4. sky subtraction,
5. artifact masking (cosmic, satellite, saturation, bleed trail),
6. weight image construction,
7. astrometric solution, and
8. initial photometric calibration.

The DECaPS pipeline uses the CP calibrated images (*InstCal*), as well as the associated weight maps and data-quality images. Any pixel marked in the CP data-quality image as being problematic, except for pixels marked as being “transient,” are given zero weight in the subsequent analysis.

The CP data products were well suited for analysis with the DECaPS pipeline. The DECaPS pipeline performed only a few additional pixel-level corrections and additions:

1. supplementary cosmetic masking,
2. removal of jump in background level on CCD S7, and
3. identification of nebulosity.

We describe these additional steps in some detail as they are unique to DECaPS, although they represent only a small amount of processing relative to the CP.

3.1. Cosmetic Masking

While analyzing CP images, we found a number of bad columns caused by charge traps in the CCDs. In the worst cases, these were very significant and could lead the DECaPS pipeline to detect hundreds of spurious sources along the bad columns. To identify these features, we high-pass filtered images by subtracting a 7×7 median filter. We then took the median value of each pixel in a stack of 100 high-pass filtered images to identify pixels that were consistently significantly high or low. This process identified 24,357 pixels that had not been identified as problematic by the CP and were consistently discrepant by more than eight counts. This is a vanishingly small fraction of the DECam focal plane ($\approx 0.005\%$), but in

low-density regions, these pixels could result in large numbers of spurious sources. This improved masking has since been incorporated directly into the CP.

3.2. S7 Background Level

The gain on amplifier B on DECam CCD S7 is unstable. However, the CCD is often well behaved, motivating preservation of this CCD if possible. Significantly discrepant gain, however, can lead to a large jump in sky level at the amplifier boundary, which the DECaPS pipeline sky algorithm does not track. The result can be large positive backgrounds that the pipeline will try to explain as a blend of a large number of stars.

To avoid this, on S7 only, the pipeline compares the 10 pixels on each side of the amp boundary with one another. A linear function of row number is robustly fit to the difference between pixels and their mirrored pixels across the amp boundary. The derived fit is then applied to all pixels in amplifier B on S7 to remove the jump in background level. This by no means fully accounts for potential problems on S7 (a multiplicative problem has been “corrected” with an additive offset), but it is adequate to prevent catastrophic failure and to enable downstream processing to identify and exclude S7 amplifier B. For precision photometry in the merged DECaPS catalogs, S7 amplifier B is always excluded.

3.3. Nebulosity

The DECaPS pipeline attempts to explain all flux in images as coming from a smooth background (i.e., the sky) and stars. This fails badly in the presence of galaxies and nebulosity, which the pipeline attempts to shred into many stars. For the most part, it seems an adequate simplification to shred galaxies into stars in the DECaPS footprint: the vast majority of the sources *are* stars, and both the algorithmic challenge and computational complexity of modeling galaxies is not justified. Nebulosity is more problematic. A small fraction ($\sim 0.1\%$) of the footprint features substantial diffuse emission, primarily in the form of $\text{H}\alpha$, O III, and scattered light from dust. This nebulosity can contain substantial fine structure that is not compatible with the DECaPS smooth-sky model, leading the pipeline to try to explain the nebulosity as the sum of thousands of carefully distributed stars. We address this problem by automatically flagging regions with nebulosity and adopting stricter source-detection criteria in flagged regions (see Section 4.2 for details of the source detection).

Convolutional neural networks (CNNs) are ideally suited to the task of recognizing nebulosity in an image. CNNs have found wide use in image, signal, and natural language processing, among other fields (see LeCun et al. 2015 for an accessible introduction). CNNs process input in a way that respects locality (e.g., nearby pixels in an image are more closely related than more distant pixels) and translation invariance (e.g., a given object might appear anywhere in an image). CNNs work by first identifying features on small scales (by sliding different convolutional filters across the input), and then connecting these features on ever larger scales to build up a rich representation of the input. In a typical classification problem, where one has labeled training data, one first defines the overall structure of the CNN (the number of layers in the network, the number of convolutional filters to include in each layer, etc.). Then, one feeds training data into the network,

compares the network output with the desired answer, and uses backpropagation to vary the weights in the network so that the residuals are reduced. After many training epochs, the network “learns” the convolutional filter weights that minimize the classification errors. CNNs have proven remarkably successful in solving image- and signal-processing problems that humans can intuitively solve by eye or by ear (e.g., facial recognition, speech recognition, or identifying cats in images). A human with little to no prior knowledge can be quickly trained to identify nebulosity in astronomical images, making CNNs an obvious approach to solving this problem.

We trained a CNN to identify 512×512 -pixel regions of images containing significant nebulosity, selected from all filters. After identification, these regions were flagged by adding a bit to the CP data-quality images (Table 5). Downstream in the analysis pipeline, putative stars in these regions were required to satisfy a sharpness criterion and be relatively uncontaminated by light from any nearby stars; otherwise, the putative stars were eliminated from the model. This prevents the pipeline from trying to explain the flux in large nebulous regions with thousands of point sources, at the expense of reducing the pipeline’s capability to model blended stars in these regions.

In order to provide training and validation data for our network, we hand-classified 512×512 -pixel images. We sorted the images into four categories:

1. NEBULOSITY—significant contamination by nebulosity.
2. NEBULOSITY_LIGHT—faint nebulosity.
3. NORMAL—no contamination.
4. SKY_ERROR—spurious fluctuations in sky level injected by upstream DECam pipeline.

We arrived at a data set of 2000 images labeled NORMAL, 1775 images labeled NEBULOSITY, 1058 images labeled NEBULOSITY_LIGHT and 629 images labeled SKY_ERROR. We used 80% of the images to train the network, and the remaining 20% to validate the network. This is a relatively small number of training images with which to train a CNN. We therefore augmented our data set by flipping the images vertically and horizontally, but did not perform other augmentation techniques (such as arbitrary rotation) that might alter the noise properties of the images or remove valuable information (e.g., the orientation of diffraction spikes). We histogram-equalized⁹ the images before feeding them to the neural network. This prevents the dynamic range of the image from being used up by a single saturated star.

On the validation data set, our trained network achieved 90% completeness and 90% purity in its NEBULOSITY classifications, with a vanishingly small percentage of NORMAL images being misclassified as NEBULOSITY. Our final validation loss (a measure of the accuracy of the classifications) was similar to our training loss, indicating that our network does not suffer from overfitting.

We applied our trained CNN to each 512×512 -pixel region of each survey image in order to flag areas with nebulosity. The neural network identified regions affected by significant nebulosity very accurately. A few corner cases were marked

Table 4
Individual-image Catalog Fields

Name	Description
x	x coordinate (pix)
y	y coordinate (pix)
flux	instrumental flux (ADU)
dx	x uncertainty (pix)
dy	y uncertainty (pix)
dflux	flux uncertainty (ADU)
fluxlbs	local-background-subtracted flux (ADU)
dfluxlbs	fluxlbs uncertainty (ADU)
qf	PSF-weighted fraction of good pixels
rchi2	PSF-weighted fraction average χ^2
fracflux	PSF-weighted fraction of flux from this source
flags	CP flag value at central pixel
fwhm	FWHM of PSF at source location (pix)
sky	sky at source location (ADU)
gain	gain at source location (e-/ADU)

Note. Fields in the DECaPS individual-image catalogs. A more complete description is available at the survey website. All fields are stored as 32-bit floating point numbers, except for x and y , which are stored as 64-bit floating point numbers, and $flags$, which is stored as a 32-bit integer.

as nebulous, such as artifacts near extremely bright stars (brighter than 6th mag), or ghosts associated with these stars. These regions are extremely rare, and from the perspective of the catalog, it is appropriate to flag them as nebulous anyway: the smooth-sky plus stars modeling is likewise inadequate here.

A more technical description of our CNN is provided in the [Appendix](#).

3.4. Failure Modes

The resulting images after CP processing are an extremely clean description of the sky, with most instrumental signatures accurately removed. A few remaining problems impact a small fraction of the images, however.

First, in regions featuring very significant nebulosity (primarily η Carinae), the CP sky subtraction has difficulties and introduces artifacts into the images. Fortunately, these artifacts are spatially smooth and are absorbed into the DECaPS pipeline sky estimates, and so these artifacts are not expected to influence the source detection or photometry in DECaPS. They do, however, lead to unsightly regions in DECaPS images in these regions.

Second, the CP has difficulties identifying cosmic rays in crowded regions. Moreover, cosmic rays that are identified are frequently only partially masked. This leads to occasional spurious sources in the DECaPS source catalogs. Fortunately, the relatively short DECaPS exposures and dense fields means that only a tiny number of sources are affected. Nevertheless, sources detected in only a single image, especially when partially masked ($QF < 1$, Table 4), are likely to be spurious.

Bleed trails associated with very bright stars occasionally introduce crosstalk artifacts in images. In rare cases on DECam CCD N15, the crosstalk correction for extremely saturated parts of the bleed trail is incorrect and leads to dramatic over-subtraction of the crosstalk on the adjacent amplifier. This case introduces spurious sources in the DECaPS processing, because the only mechanism the DECaPS modeling has to address regions with substantial negative flux is to reduce the

⁹ Histogram equalization maps the original image to a modified image where the number of pixels of each brightness is the same; i.e., a histogram of the image’s pixel values is flat. Histogram equalization is often used in image classification problems to make images more comparable to one another. Two images of the same scene taken with different exposure times have different pixel values, but histogram equalization would remove this difference.

sky level. This can leave significant positive sky on the image everywhere away from the crosstalk artifact, corrupting the flux measurements of stars throughout the image and potentially introducing numerous spurious stars that attempt to explain the high background. Fortunately, we have only observed this effect on a few individual CCD frames of the 600,000 composing the survey.

Related effects occur whenever processing errors leave large negative fluxes in the images. Very bright stars saturate the serial register, leading to large artifacts in the images. These artifacts are usually masked by the CP, but in rare cases, they are not masked and result in spurious sources and contaminated fluxes.

4. Catalog Construction

DECaPS catalogs were constructed by the `crowdsources` crowded field photometry pipeline, which is described in detail in a future paper. These catalogs were then organized and merged using the Large Survey Database package (Juric 2012).

Briefly, the `crowdsources` pipeline models images as the sum of a smooth-sky component plus the sum of a large number of sources that share a common point-spread function (PSF) description, similar to the approaches taken by DAOPHOT (Stetson 1987), PS1 (Magnier et al. 2016b), DOLPHOT (Dolphin 2000), and the Tractor (Lang et al. 2016). The pipeline takes advantage of the LSQR sparse least-squares routine (Paige & Saunders 1982) to simultaneously solve for the sky level as well as the fluxes and positions of thousands of sources. In DECaPS, `crowdsources` operates on 1024×1024 pixel subregions (roughly $4\frac{1}{2} \times 4\frac{1}{2}$) of each 4094×2046 pixel CCD, simultaneously solving for the sky and sources in those regions. The densest fields in the survey feature 30,000 sources per subregion, leading to 90,000 parameters in the linear least-squares fits (flux, as well as x and y coordinates for each of the 30,000 sources).

`Crowdsources` iteratively finds fainter, more blended sources, and improves its PSF and sky estimates. Each iteration consists of the following steps:

1. Smooth sky subtraction.
2. Source detection on residual image.
3. Fit fluxes, positions, and global sky level.
4. Improve positions via centroid computation.
5. Improve PSF.

Each iteration begins with the current best model image and PSF. In the initial iteration, the best model image is simply constant and equal to zero, and the PSF shape is determined from the CP-measured FWHM.

After each iteration, a new PSF model, source list, and improved source locations are available, which are used as input to the next iteration. The pipeline continues iterating until at least four iterations have been performed. Afterward, the pipeline will cease iterating if fewer than 100 new sources have been detected, or if 10 iterations have occurred. The 10-iteration limit is reached in very dense regions like Baade’s window, and very rarely elsewhere. Finally, a last pass is made without detecting new sources or allowing positions to vary; only forced photometry at the existing list of stellar positions is performed.

In general, this procedure is very successful. Figure 3 examines a 100×100 -pixel region from a z -band image in the vicinity of $(l, b) = (-9^\circ.7, 3^\circ.3)$: toward the inner Galaxy and

somewhat off the midplane where the highest extinction lies. The first panel shows the image with an arcsinh stretch; the second panel shows the detected sources; the third panel shows the model image; and the fourth panel shows the residual image after additional stretching. Saturated regions flagged by the CP are set to zero. Imperfections in the PSF model lead to significant coherent residuals around all bright stars. However, the model identifies and attributes fluxes to blended sources very well; many complicated asterisms are neatly resolved into individual stars. For example, the stars near (395, 315) are neatly identified and photometered, although one faint star goes unidentified. The pipeline tries to be conservative when deblending sources. A number of close blends are not deblended and are readily recognized in the residual image as pairs of dark, positive residuals flanking a central light, negative residual. For instance, one such case is evident near (345, 375).

Ultimately, `crowdsources` relies on simple, rigid heuristics to identify a fixed set of objects in a field. More principled, probabilistic techniques are possible, and can better accommodate the highly covariant fluxes of different sources, albeit at greater computational expense. Some such techniques are explored in Brewer et al. (2013), Daylan et al. (2017), and Portillo et al. (2017).

4.1. Smooth-sky Subtraction

In the initial smooth-sky subtraction, the median of each 20×20 pixel subregion of the residual image (image minus best-fit model image) is computed. The pipeline linearly interpolates between these median values to construct a smooth-sky image, which is subtracted from the residual image. This sky estimate will remove most nebosity in addition to the sky, and will be biased high by any undetected sources in the image.

4.2. Source Detection

Next, sources are detected. The pipeline convolves the residual image times the CP inverse-variance image with the PSF to construct a significance image. This image gives the statistical significance of an isolated point source being present at each point in the image, neglecting covariance with any other sources. Following standard least-squares analysis for an image with varying noise, the significance image S_I is given by Equation (1),

$$s(\mathbf{x})^2 = \int d\mathbf{u} \, \sigma(\mathbf{x} + \mathbf{u})^{-2} P^2(\mathbf{u})$$

$$S_I(\mathbf{x}) = \frac{\int d\mathbf{u} \, I(\mathbf{x} + \mathbf{u}) \, \sigma(\mathbf{x} + \mathbf{u})^{-2} P(\mathbf{u})}{s(\mathbf{x})}, \quad (1)$$

where \mathbf{x} gives the pixel coordinates in the image, I is the residual image, σ is the CP statistical uncertainty image, and P is the PSF.

Peaks in S_I of greater than 5σ are potential newly detected stars. This value corresponds to 5σ for isolated stars including only photon and detector noise, neglecting difficulties with blending. Accordingly, if peaks are too blended with stars in the current best-fit model image M , they are excluded. To assess this, the pipeline defines a “blending threshold” B of 0.2 and constructs an analogous “model significance image,” S_M , according to Equation (1), with the current model image




Figure 3. Comparison between a 100×100 -pixel portion of a z -band image in a dense region and the `crowdsource` model for that image. The first panel shows the image; the second panel the image with detected sources overlaid; the third panel the model image; and the fourth panel the residual image. The red lines are centered five pixels above each detected source. The pipeline detects the sources in the image very well and also provides excellent photometry, resulting in a model that well describes the observation. The image shown represents about 2×10^{-9} of the total DECaPS imaging.

taking the place of the residual image. Peaks are kept only if one of the following two conditions are satisfied:

1. $S_I/S_M > 2B$: the ratio of the peak in the significance image to the model significance image is more than twice the blending threshold, or
2. $(S_I/S_M > B)$ and $(I/M > B)$: the same ratio is more than the blending threshold, and the ratio of the image to the model (without smoothing to construct the significance image) is also more than the blending threshold.

Initially, a criterion based only on I/M was desired, but since even isolated 5σ stars may have central pixels with negative fluxes due to noise, additional criteria based on the significance images were adopted.

In regions featuring substantial nebulosity, the above procedure identifies substructures in the nebulosity as stars. In order to prevent the pipeline from generating thousands of spurious sources in these regions, this procedure is altered in the roughly 0.1% of the footprint flagged as nebulous. First, the blending threshold is increased to 100 (from 0.2) in these regions, effectively preventing the pipeline from finding any blended sources. Second, the sharpness of the peak in the significance image is compared with expectations from the PSF model. The drop in significance from the peak pixel to any neighboring pixel must be more than 70% of the expectation to be kept (or 50% if the central pixel is saturated). Otherwise, peaks are rejected as unlikely to be stars and not included in the image modeling in subsequent steps. This is the only special treatment the pipeline performs in regions with nebulosity. No attempt beyond the normal sky subtraction is made to improve the fluxes of stars in these regions.

4.3. Flux and Position Solution

Any stars found are added to the list of currently detected stars in the image. These new stars may influence the fluxes and positions of other stars, at least through their influence on the sky, as well as through blending. `Crowdsource` simultaneously solves for an overall sky level and the fluxes and positions of all of the currently identified stars using the linear least-squares routine LSQR. A source's image is not a linear function of its position. Accordingly, for use with the LSQR, we linearize the problem by approximating changes in the position of the source by the first term in the Taylor series expansion. Fitting an overall sky level allows the fit to partially account for the initially high sky estimate.

4.4. Centroiding

The positions estimated in Section 4.3 are linearized approximations to the positions, and are accurate only for small changes in the positions of the source. Better performance can be obtained by computing the sources' centroids, but this is challenging in a crowded field, where the flux of neighboring stars is blended with the flux of a target star. `Crowdsource` attempts to overcome this by using its new best-fit model of the image to eliminate contaminating neighboring stars from the centroid computation. The pipeline visits each source in turn, and subtracts its model of all other sources from the image. The source's centroid is then measured on the neighbor-subtracted image. These centroids are used as input positions in the next iteration.

4.5. PSF Modeling

Finally, the PSF model is improved. The PSF is modeled as an ideal-seeing PSF, convolved with an elliptical Moffat PSF (Moffat 1969) representing the current seeing, plus a 9×9 residual image in the core of the PSF. The parameters of the Moffat PSF and the value of each pixel of the 9×9 residual image are allowed to change linearly over each 1024×1024 subregion of a CCD. The pipeline takes the brightest 200 unsaturated sources in each subregion, which are chosen to be unmasked, relatively unblended, and to have positions that did not move by more than 1 pixel in the previous iteration. The images of these stars, after neighbors have been subtracted according to the best model so far, are used to fit the parameters of the PSF model for the next iteration.

A primary goal of the ideal-seeing PSF is to allow the pipeline to describe diffraction spikes and other features in the wings of the PSF. These features may not be constant over the focal plane, but treating them as constant is an improvement over neglecting them altogether. The ideal-seeing PSF is derived by averaging the PSF of many stars in the focal plane over all CCDs and several images on a night with very good seeing. This PSF is then deconvolved with a Moffat PSF using Richardson-Lucy deconvolution to obtain a PSF with better seeing than obtained on any night. The wings of the deconvolved PSF are then fit as the sum of six Moffat PSFs plus power-law diffraction spikes and horizontal and vertical features. The resulting noise-free model PSF wings are then blended with the central high S/N core of the deconvolved PSF to produce the final ideal-seeing PSF.

4.6. Failure Modes

The `crowdsources` analysis can fail in a number of ways. Because `crowdsources` attempts to explain all flux in the image as due to stars and a smooth background, violations of this assumption and imperfections in the PSF can introduce spurious sources into the catalog.

By far the most common cause of spurious sources is very bright stars. The `crowdsources` PSF wings are based on a single ideal-seeing model PSF that is fixed for all images in each band. This ideal-seeing PSF is then convolved with a Moffat profile, whose parameters may vary linearly in each image. This limited allowed variation, however, is not adequate to describe all of the features in the wings of the PSF. For instance, the relative amount of flux in the different diffraction spikes varies over the focal plane and cannot be captured in the `crowdsources` model. Accordingly, some spurious sources are present near most bright stars (brighter than roughly 12th mag). Extremely bright stars, like α Cen, can produce thousands of spurious sources. Demanding that a source be detected in multiple filters, or at least detected multiple times, can significantly reduce the number of this kind of spurious source.

Bright stars are especially problematic in Y band because of the extremely long PSF features extending vertically and horizontally from the PSF center in Y band. The `crowdsources` PSF model describes these features accurately, but for sufficiently bright stars, they are highly significant even several hundred pixels from the star's center, beyond the range that the `crowdsources` PSF extends. Accordingly, very bright stars in Y band can lead to clusters of spurious sources 150 pixels away in the horizontal and vertical directions, where the PSF model ends and flux is left on the image. These sources will not be detected in any other bands, but will slightly contaminate the Y -band photometry in their vicinity.

Because the `crowdsources` PSF model is expected to be imperfect, no attempt is made to find sources too near existing, brighter sources. This means that many close blends are not identified. Blends are usually not deblended if the two stars involved are separated by less than one half FWHM, although the flux ratio of the blended sources also matters. More aggressive deblending is possible, but must be balanced against avoiding incorrect deblending of single sources due to PSF-misestimation.

In extremely dense regions like Baade's window, the `crowdsources` algorithm never completely converges. A very small fraction of faint but clearly significant sources are never identified, and extremely complicated blends of hundreds of stars are not fully worked out. It is not clear how much better one can do in these regions, however: they are extremely blended, and the derived model images account for most of the flux.

In dense regions, the sky tends to be overestimated. The algorithm may move stars to regions where the sky has been overestimated and assign them negative fluxes to improve the fit of the model image to the observations, introducing spurious sources into the catalog. These objects can lead to sources with `fracflux` outside the range 0–1; see Section 7.2.1. A more sophisticated sky estimation algorithm might eliminate these spurious sources. However, even in the densest regions, less than half a percent of sources have negative fluxes. These are also easily identified and are rarely badly blended with other sources.

For simplicity, the `crowdsources` pipeline adopts a fixed size, $5'' \times 5''$ PSF stamp independent of the seeing of the image. In the worst seeing conditions, this PSF stamp size does not enable the pipeline to track flux far enough into the PSF wings. This can lead to spurious sources $2''.5$ from true sources, but only in images with $\text{FWHM} > 2''$, roughly 0.4% of all images. Most such spurious sources should be detected in only one of the three observations of each part of the sky, facilitating their rejection from scientific analyses.

The DECaPS pipeline aims to characterize all sources in individual images down to a point-source depth of 5σ . In practice, DECaPS achieved something closer to 6σ in typical images, owing to crosstalk between the sky subtraction routine and the source detection routine. Sky subtraction was performed by median-filtering the image in a 20×20 pixel region. Sources in this region will bias the sky high. Having been detected, the sources are subtracted and the sky is refit. This process is iterated several times, so the effect on the final fluxes is small. However, this means that once sources are detected at 5σ after an initial sky subtraction has been performed, they are later detected at apparently higher significance once the sky has been lowered. We could have eliminated this effect by using larger regions for the median filtering or considering a local background subtraction when defining the initial source significance, but both of these approaches also have downsides. Accordingly, the final survey does not quite reach 5σ .

5. Photometric Calibration

The DECaPS photometric calibration follows the photometric calibration from PS1 (Schlafly et al. 2012; Magnier et al. 2013, 2016a) closely, which was modeled on the calibration of the SDSS (Padmanabhan et al. 2008). The basic idea is to solve for the throughput of the atmosphere and detector as a function of time, airmass, and location of the observed source in the focal plane, in order to minimize the variance in repeated measurements of the same sources. Since DECaPS was designed to observe most stars in its survey area three times, and since billions of stars were observed, precise measurements of the system throughput and its variation can be made.

In PS1, we adopted a simple model for the system throughput: a single constant throughput on each night (i.e., a system zero-point), an airmass term on each night, and a constant-in-time flat-field correction. In DECaPS, the same model was used, with an additional simplification: the airmass term was held constant in each filter over all nights of the survey. This simplification was necessary because of the small range of airmass observed on each night, coupled with imperfect aperture corrections (see Section 5.1). Each filter is calibrated independently. The flat-field correction treats each 256×256 pixel portion of each detector independently.

The performance of this simple model is excellent. Table 3 gives the accuracy of the calibration model in predicting the zero-points of individual images, as compared to a model where the zero-point is allowed to float in each image (σ_{zp}), and the average scatter in differences in fluxes of bright stars on an image from their mean fluxes over all observations (σ_{bright}). On a typical image, fluxes of bright stars agreed with the mean fluxes over all observations of those stars to better than 8 mmag in all filters. Y band showed the best performance, with 6 mmag scatter. This agreement is superior to the results we achieved in




Figure 4. Distribution of residuals as a function of hours since noon UTC on the night of 2017 April 29. Each column shows the distribution of residuals at a particular time, relative to the average magnitude of all detections of those sources. Solid lines show the 16th, 50th, and 84th percentile of the residuals. At any given time, the residuals are consistent to better than about 8 mmag, and overall zero-points are predicted to better than 1%. The dominant systematic in the overall zero-points is the aperture correction of the PSF.

PS1, where 15 mmag scatter in the flux of bright stars was typical. Figure 4 shows the residuals of the calibration on a typical night of the survey.

The flat-field correction derived in the calibration is shown in Figure 5. The flat fields show a large number of features. Most prominent is amplifier B of CCD S7, shown as a saturated black region at the left of the *gri* flat-field correction maps. No photometry from S7 amplifier B is used in the mean magnitudes we compute for DECaPS sources. Potentially, photometry from this amplifier could have been separately calibrated, but this was not implemented for DECaPS.

The *g*- and *Y*-band flat fields feature significant CCD-to-CCD offsets. This may be caused by the combination of varying throughput as a function of wavelength between the different DECam CCDs, and the significantly redder typical spectrum of DECaPS sources than the spectrum used to determine the DECam flat field. The strong radial gradient in *i* is thought to be due to variation in the *i*-filter bandpass as a function of angle. The *z* band shows a significant “pupil ghost” pattern, which is related to the technique used to determine the initial DECam sky flats; to a lesser extent, *g* and *i* show similar effects. The <0.01 mag *z*-band pattern is a small residual remaining after the much larger pupil corrections made in the CP. A hint of tree rings (Plazas et al. 2014) is noticeable in the upper right CCD in *r* and *i* bands. The DECaPS flat field does not have the resolution to fully resolve the tree rings, and some mmag leak into the flat-field corrections. In any case, the CP flat fields are quite good: the flat-field corrections have less than 9 mmag scatter in all filters, and the correction for the worst filter, *g*, is largely caused by varying CCD throughputs with wavelength, which no single flat field could correct.

5.1. Limitations of the Photometric Calibration

The calibration model could be improved in several respects. It ignores “systematic chromatic errors” (Li et al. 2016), essentially neglecting the fact that changes in system throughput have a wavelength dependence. This causes stars to require different corrections into a “standard” system according to their spectra, but the DECaPS calibration performs the same correction for all objects. More physics-based analyses can readily accommodate these effects, and are already being performed for DECam (e.g., Burke et al. 2017).

The photometric calibration model is only able to predict the zero-point of a given exposure to within about 9 mmag (σ_{zp}).




Figure 5. DECaPS flat-field correction in the *grizY* bands. The corrections are in all cases small, with typical values of about 6 mmag. The features that are detected, however, are detected with extremely high significance, and relate to real features in the DECam CCDs and their processing. The mean and standard deviation of the flat-field correction are given in the lower left and right of each panel, respectively.

While not bad, this is substantially worse than achieved by PS1, which had 3 mmag precision. Evidently, the simple single nightly system zero-point and airmass term provided a much better description of the zero-points in PS1 images than the similar model adopted for DECaPS. The dominant unmodeled source of variation in the zero-points of DECaPS images turns out to be the “aperture correction” of the crowded DECaPS photometry. If the pipeline PSF were perfectly accurate, then the measured PSF flux should equal the true flux on average. With an imperfect PSF, the measured PSF fluxes become systematically biased. The bias can be addressed by measuring aperture magnitudes for bright stars in very large apertures, and comparing them with the measured PSF fluxes of those stars. The difference is related to the aperture correction and can be applied to all PSF fluxes to correct them to total fluxes. This procedure is challenging in crowded fields, so DECaPS did not perform an aperture correction. Comparison of different PSF fitting routines indicates that the pipeline was too inflexible in fitting the PSF roughly $2''$ from the center of the PSF, leading to variation in the shape of the PSF at this radius to cause apparent ~ 7 mmag variations in measured PSF flux—that would have been absorbed by the aperture correction, had one been made. Future reductions of the DECaPS imaging can improve on these results by using a more flexible PSF model, or alternatively, `crowdsource` could be adapted to compute an aperture correction on neighbor-subtracted images.

Other improvements on the DECaPS photometric calibration are possible. We neglect in this photometric calibration the “tree ring” pixel-area effects discussed in Plazas et al. (2014), which are expected to contribute a few mmag of photometric noise. Eliminating systematic chromatic errors is difficult without knowing the spectral energy distributions of the stars within the filter bandpasses, but upcoming *Gaia* low-resolution spectrophotometry will eliminate this problem for bright stars. By addressing these shortcomings, we expect that significantly better than 5 mmag photometry for typical stars is possible over large areas with DECam, even in crowded regions.

5.2. Absolute Zero-point

The absolute zero-point was obtained by reference to PS1, in combination with predicted colors in the DECam bands. The procedure is laid out in Scolnic et al. (2015): stellar energy distributions are taken from calibrated *Hubble Space Telescope* (*HST*) measurements and integrated over the PS1 and DECam filter bandpasses to predict the color differences between the surveys. Absolutely calibrated PS1 photometry is then transformed to absolutely calibrated DECam photometry using the derived color transformations for blue stars ($0.25 < g_{\text{PS1}} - i_{\text{PS1}} < 1$) where the transformations are most accurate. The resulting absolutely calibrated, synthesized DECam photometry is then used to fix the absolute zero-points for DECaPS. Ultimately, this calibration derives from *HST* measurements from Bohlin (2014). A calibration field at $(\alpha, \delta) = (100^\circ, -27^\circ)$ is used to link the two surveys together.

The DECaPS absolute calibration is then no better than the PS1 absolute calibration. DECaPS adopted the original PS1 absolute calibration from Tonry et al. (2012) to remain consistent with the original PS1 photometry. The work of Scolnic et al. (2015) finds that this absolute calibration is somewhat offset from AB. To shift DECaPS to the absolute calibration of Scolnic et al. (2015), offsets of 0.020, 0.033,

0.024, 0.028, and 0.011 mag must be added to the DECaPS *grizY* magnitudes, respectively.

6. Survey Characteristics

6.1. Photometry

The accuracy of DECaPS photometry can be tested by comparison with PS1 over the limited area of overlap. The DECaPS survey extends slightly north of $\delta = -30^\circ$ toward the Galactic center and anticenter to enable this comparison, and also contains a few calibration fields off the plane above $\delta = -30^\circ$. The results of this comparison for the *r* band are shown in Figure 6. The first panel shows the comparison on a calibration field at $(l, b) = (236^\circ, -14^\circ)$, the second panel shows the comparison at $(l, b) = (-115^\circ, 0^\circ)$, and the third panel shows the comparison at $(l, b) = (0^\circ, 0^\circ)$.

The two surveys show good agreement in their photometry; the bright-end root mean square (rms) difference is 0.02 mag, and smallest in the unreddened calibration field. There is a modest trend in the photometry with magnitude in the anticentral field (0.03 mag over from 15th to 21st mag). The trend may be again connected to reddening (fainter stars are systematically more reddened), but we see the opposite trend, if any, in the most reddened Galactic center field. No two fields have quite the same overall zero-point; the DECam photometry becomes systematically brighter relative to the PS1 photometry in the more reddened fields. The Galactic center field is offset by -0.05 mag relative to the calibration field. We attribute this to the imperfect color transformations in the presence of reddening, and conclude that the calibration is more uniform than about 0.05 mag, although expectations from the internal repeatability suggest that the calibration should be significantly better than this. Other filters show qualitatively similar trends.

Another comparison between PS1 and DECaPS is shown in Figure 7, an *i* versus *i* − *z* color–magnitude diagram (CMD) in the PS1 and DECam bands. The superior depth and deblending of the DECaPS pipeline makes for a much more obvious bulge red clump at $i = 21.5$, $i - z = 2.0$; the DECaPS measurements clearly improve on the PS1 measurements in the small overlap region between the two surveys. On the other hand, a few blue sources around $i = 22.0$, $i - z = 0.5$ are presumably largely spurious deblending errors in DECaPS; they are absent from the PS1 CMD, and can be largely eliminated by removing the most blended sources (`fracflux`, see Table 6).

6.2. Color Transformations

PS1 photometry is of qualitatively similar depth to DECaPS photometry. The PS1 filter bandpasses are moreover very similar to the DECaPS bandpasses. Accordingly, it is useful to have a set of color transformations to allow comparison between DECam *grizY* filters used in DECaPS and the PS1 *grizy* filters. To develop these color transformations and improve the stability of the photometric calibration, DECaPS observed a calibration field at $(\alpha, \delta) = (100^\circ, -27^\circ)$ on many different nights of the survey. This field has low reddening (≈ 0.15 mag $E(B - V)$) but is at low latitude ($b = -14^\circ$), meaning that it contains many stars suitable for deriving color transformations.

Color transformations between DECam and PS1 were fit with a cubic polynomial in the PS1 *g* − *i* color. The color transformations are given in Equation (2). Only point sources brighter than 19th mag in a nearby filter that was neither *g*, *i*, nor the filter of




Figure 6. Comparison between DECaPS and PS1 photometry in the r band. Other bands are similar. The left panel shows a calibration field at $(l, b) = (236^\circ, -14^\circ)$; the middle panel shows a field at $(l, b) = (-115^\circ, 0^\circ)$; and the right panel shows a field at $(l, b) = (0^\circ, 0^\circ)$. In general, agreement is good, although there is a 0.05 mag zero-point offset between the Galactic center field and the calibration field. This may be due to the color correction, which was developed in a field of little reddening, while the Galactic center field (and, to a lesser extent, the other midplane field) is heavily reddened. The mean photometric offset and the scatter between the magnitudes in the two surveys for stars greater than 19th mag is shown; typically, the two surveys agree to better than 2%. The background gray scale shows the density of stars at a given PS1 magnitude and magnitude difference. The mean and standard deviation for bright stars ($PS\ r < 17$) are given by μ and σ . The black lines show the 16th, 50th, and 84th percentiles of the magnitude difference at each PS1 r magnitude, and the dashed horizontal line shows the mean offset μ .




Figure 7. Comparison between DECaPS and PS1 $i - z, z$ CMD for a 0.1° radius beam around $(l, b) = (0.3, 0.7)$. The region is modestly crowded. The DECaPS CMD is significantly deeper and clearly shows the red clump in the bulge, which is hard to separate in the PS1 diagram.

interest were used to determine the transformations. Furthermore, only stars with $0 < g_{PS1} - i_{PS1} < 2.9$ were used; the color transformation outside this broad range requires more parameters and more stars to constrain. Figure 8 shows the derived color

transformations. After color correction, the rms difference between DECaPS and PS1 photometry is about 0.015 mag.

$$\begin{aligned}
 c &= g_{PS1} - i_{PS1} \\
 g_{DECam} - g_{PS1} &= 0.00062 + 0.03604c \\
 &\quad + 0.01028c^2 - 0.00613c^3 \\
 r_{DECam} - r_{PS1} &= 0.00495 - 0.08435c \\
 &\quad + 0.03222c^2 - 0.01140c^3 \\
 i_{DECam} - i_{PS1} &= 0.00904 - 0.04171c \\
 &\quad + 0.00566c^2 - 0.00829c^3 \\
 z_{DECam} - z_{PS1} &= 0.02583 - 0.07690c \\
 &\quad + 0.02824c^2 - 0.00898c^3 \\
 Y_{DECam} - Y_{PS1} &= 0.02332 - 0.05992c \\
 &\quad + 0.02840c^2 - 0.00572c^3. \quad (2)
 \end{aligned}$$

6.3. Sample CMDs

The DECaPS footprint includes open clusters, globular clusters, nearby molecular clouds, distant molecular clouds, and regions with a wide range of stellar densities. Accordingly, CMDs from DECaPS show rich diversity.

To demonstrate the precision of the DECaPS photometry, Figure 9 shows the $g, g - r$ CMD of stars within 0.03° of the center of NGC 2660, an approximately 1 Gyr old open cluster at a distance of 2.7 kpc (Sandrelli et al. 1999), together with a control annulus of equal area 0.1° from the cluster center. Reddening of the field stars in the region leads them to have a very similar CMD to that of NGC 2660, but nevertheless, the sharp NGC 2660 sequence clearly stands out from the background. A secondary sequence roughly 0.7 mag brighter than the primary sequence may indicate a substantial fraction of binaries (Sandrelli et al. 1999).

The NGC 2660 CMD demonstrates the value of DECaPS photometry to constrain stellar populations in the Milky Way. However, typical CMDs in the footprint are much more complicated than that in the vicinity of NGC 2660. Figure 10 shows a selection of CMDs in different filters and Galactic longitudes from DECaPS.




Figure 8. Color transformations between DECam and PS1. The difference between DECam and PS1 photometry in the *grizY* bands is shown as a function of PS1 $g - i$ color for bright stars in the vicinity of $(\alpha, \delta) = (100^\circ, -27^\circ)$. A cubic fit to the difference is shown by the blue dashed lines; these are the color transformations adopted in DECaPS. Photometric agreement is good, about 0.015 mag. The background gray scale shows the number density of stars at each PS1 $g - i$ color and magnitude difference. The three solid black lines show the 16th, 50th, and 84th percentile of the magnitude difference at each PS1 $g - i$ color, while the blue dashed line shows the fit cubic polynomial transformation.




Figure 9. DECaPS $g, g - r$ CMD of NGC 2660 (upper panel), together with the same CMD for an annulus of equal area centered on the cluster but excluding all stars within 0.1° of the cluster (lower panel). The upper panel shows a sharp main sequence superimposed on a background of field stars whose significant and increasing reddening with distance causes them to overlap significantly with the NGC 2660 sequence. A second sequence roughly 0.7 mag above the NGC 2660 sequence is also visible. The lower panel shows only the background of field stars from the control region.

Each row of Figure 10 shows the stars in a one-degree radius beam in the Galactic plane ($b = 0^\circ$) at different Galactic longitudes, stepping out in 30° increments from $l = 0^\circ$ to $l = -120^\circ$. Each diagram shows a strong sequence of blue stars that are reddened as one proceeds to fainter magnitudes: FGK main-sequence stars in the presence of increasing extinction along the line of sight. The Galactic center field (top row) additionally features a fainter, redder sequence, coming from red clump stars in the Galactic bulge. Finally, there is a redder sequence visible at all magnitudes in the inner galaxy fields (top three rows), produced by the red clump in the Galactic disk. These sequences are marked in the Galactic center i versus $i - z$ CMD. The blue main-sequence, bulge red clump, and disk red clump features are marked with offset solid, dashed, and dotted red lines, respectively. The outer Galaxy fields show relatively unreddened nearby stars (e.g., the unreddened M-dwarf sequence with $g - r = 1.5$ mag at $l = -90^\circ$), on top of fainter, redder populations.

The CMDs in Figure 10 show stars in 1° radius beams. Within these regions, the reddening varies dramatically, complicating interpretation of the diagrams. For instance, despite the prominent bulge red clump sequence in the top row of Figure 10, bulge red clump stars are not detected directly at $(l, b) = (0^\circ, 0^\circ)$; most of the stars in this sequence come from the edge of the 1° field.

6.4. Photometric Uniformity

The internal consistency obtained in the DECaPS photometric calibration (Section 5) suggests that spatial systematic variations in the photometric calibration over the footprint should be no larger than 0.01 mag in size. This estimate can be tested by measuring the consistency of stellar colors over the DECaPS footprint (e.g., High et al. 2009; Schlafly et al. 2010). Unfortunately, because the DECaPS footprint features very large extinction, the typical colors of stars are completely dominated by reddening. Nevertheless, Figure 11 shows the median $r - i$ color of stars with Y -band magnitudes brighter than 18 over the footprint. White to black spans 0.2–2.5 mag.

The figure is clearly dominated by extinction in the Galactic plane. Well-known molecular clouds like the Vela Molecular Ridge around $l = -90^\circ$ and the Pipe Nebula at




Figure 10. DECaPS CMDs of 1° radius beams in the Galactic plane. Each row shows a different field in the Galactic midplane; from top to bottom, the rows correspond to $l = 0^\circ, -30^\circ, -60^\circ, -90^\circ$ and -120° . The diagrams show a diverse array of stellar populations and extensive reddening. Three features are marked with offset red lines in the i vs. $i - z$ CMD toward the Galactic center: blue main-sequence stars (solid line), bulge red clump stars (dashed line), and disk red clump stars (dotted line). Extensive differential reddening in the 1° fields complicates interpretation of the diagrams; for instance, the bulge red clump stars are not detected at $(l, b) = (0^\circ, 0^\circ)$; the red clump stars in plot come from close to the edge of the 1° radius beam.

$(l, b) = (1^\circ, 4^\circ)$ are apparent. In the very inner Galaxy ($|l| < 5^\circ$), the median $r - i$ color of stars becomes steadily redder as $b \rightarrow 0^\circ$, until $b \sim 1^\circ$, where the $r - i$ color suddenly becomes more blue. This occurs because behind dense molecular clouds, only few stars are detected, leading the median-colored star to become a star in the foreground of those molecular clouds, rather than a star in their background. No clear signs of photometric errors due to the photometric calibration are present in Figure 11, but given the wide range of colors shown and the large variations in color due to dust reddening, this test does not strongly constrain the quality of the photometric calibration.

6.5. Astrometry

Precise astrometry was not a major objective of DECaPS. The crowdsourcing pipeline computes best-fit x and y positions for the sources it detects, and transforms these into celestial coordinates using the CP-computed World Coordinate System (WCS) with no additional steps or calibration.

This is a simple but problematic approach. The original CP WCS may have a different notion for the “center” of a source than the crowdsourcing pipeline. Ultimately, both pipelines attempt to use the centroids of the PSF as the center, but the centroid may have varying definitions depending on how much weight is put on flux far out in the wings of the PSF. This is a




Figure 11. Median $r - i$ color of stars detected in DECaPS with Y magnitude brighter than 18. The $r - i$ color is dominated by reddening from dust and its distribution along the line of sight; major dust clouds like the Vela Molecular Ridge at $l = -90^\circ$ are prominent. There are no clear signs of errors in the photometric calibration. The discontinuity in $r - i$ color at $|b| = \pm 1^\circ$, $|l| < 5^\circ$ is the result of increasingly dense molecular clouds, leading the median star to be in the foreground of these clouds, rather than in the background.




Figure 12. Mean difference in position between DECaPS and *Gaia* as a function of position over the DECaPS footprint. White to black spans $0''.0$ – $0''.5$. Regions of the footprint which were processed only using CP version 4 or later have typical differences close to zero, while regions processed using earlier versions of the CP have typical differences of $0''.1$.

small effect, but it would motivate directly calibrating the `crowdsourc` positions to external catalogs rather than using a separate set of centroids in the calibration.

Additionally, the CP WCS calibration changed over the course of DECaPS. Starting with CP version 4, WCSs were computed using *Gaia* as the astrometric reference catalog, while before that point, 2MASS was used. This is unfortunate for DECaPS, because it means that the astrometry from the first part of the survey is slightly inconsistent with the last part of

the survey. Still, the 2MASS and *Gaia* astrometry are very similar, and DECaPS only requires that the astrometry be similar enough to one another so that most stars can be merged with an association distance of $0''.5$. However, other applications of the DECaPS observations, such as the construction of coadded stacks, may require greater fidelity.

Figure 12 shows the comparison between DECaPS and *Gaia* astrometry for stars brighter than 19th mag in r . The mean difference in position between the DECaPS stars in each part of

the sky is shown; white to black is $0''$ – $0''.5$. Typical values are $0''.1$ in the part of the footprint that was primarily observed using CP versions earlier than 4, while the CP 4-only fraction of the footprint shows much lower residuals. This level of difference between DECaPS and *Gaia* is unproblematic from a DECaPS perspective; the astrometry is much better than the $0''.5$ necessary to merge the catalogs from the individual images.

6.6. Spurious Sources

The `crowdsource` pipeline iteratively identifies new sources after having subtracted models of other sources, later fitting the positions and fluxes of both newly identified and preexisting sources simultaneously. In the presence of an imperfect PSF and noise, this procedure can generate large numbers of spurious sources, for instance, around very bright stars (Section 4.6).

Even away from bright stars and without noise, in very dense regions, the heuristics `crowdsource` uses to avoid overfitting (Section 4.2) prevents the pipeline from finding the true solution, instead forcing it to find an approximate description that may include stars with fluxes and positions very different from those of the actual stars in the image. The prevalence of such spurious sources can be estimated using simulated images.

Images of the densest areas in the footprint are simulated using the observed PSFs from specific DECam CCDs. The simulation uses a luminosity function that is chosen to match the bright end of the observed luminosity function in dense Galactic bulge fields. 10^5 sources are injected into a 1000×1000 pixel region with a spatially varying PSF taken from the `crowdsource` model of one of the DECaPS images. `Crowdsource` then analyzes these images with the PSF from another part of the same CCD—a PSF that is not exactly right, but is similar to the correct PSF. These tests are performed using a PSF with better-than-average seeing ($0''.7$) as well as worse-than-average seeing ($1''.3$). No noise is added to the images. The no-noise limit is appropriate because here the focus is on the densest regions, where the pipeline does not approach the 5σ photon-noise significance limit, and is instead limited by deblending heuristics and PSF systematics.

For these tests, the true source catalog is known, so the number of spurious sources can be assessed. A catalog source is considered correct if it lies within one half FWHM of a true source with flux within a factor of two of the catalog flux. Any source not satisfying this condition is considered spurious. With this definition, the simulations find that roughly 12% of sources are spurious in the densest regions, for both good seeing and poor seeing images. This fraction is unchanged when the true PSF is replaced with a modestly different PSF from a different part of the same CCD; it stems from the deblending heuristics not allowing the pipeline to fully deblend sources, and not from PSF errors. The completeness rate is much higher for the good seeing image than the poor seeing image: roughly 28,000 sources are detected in the good seeing image, while 12,000 are detected in the poor seeing image.

These tests give a general sense for the rate of spurious sources in dense regions in the DECaPS individual epoch catalogs (see Section 7.2.1). In some ways, the tests underestimate the spurious source prevalence: they assume an absence of artifacts or nebulousity, that the noise model is known, and that the PSF has smooth wings. In other ways, the tests will overestimate the spurious detection prevalence:




Figure 13. DECam Legacy Survey Viewer view of a small part of NGC 6188. The three-color image is made of DECaPS *g*, *r*, and *z* stacked images.

because they have zero noise, every bit of positive residual flux corresponds to a much greater than 5σ source. This forces the pipeline to deblend even the faintest, most blended sources, where the spurious source rate is highest, subject only to the deblending heuristics of Section 4.2. We defer to future work a more complete analysis of how the `crowdsource` pipeline behaves in different conditions.

DECaPS also includes band-merged catalogs, where an attempt is made to link together detections in different images of the same objects (Section 7.2.2). Linking detections together into objects can be challenging in dense regions where depth is determined largely by image quality. The number of spurious sources in band-merged catalogs depends on the variation in image qualities of the full set of images taken in a given part of the sky.

7. Data Release

All DECaPS images and catalogs are now publicly available. Raw and CP processed images are available through the NOAO Archive (NOAO Prop. ID 2014A-0429, 2016A-0327, and 2016B-0279). For both individual and merged images, catalogs are available from the survey website, <http://decaps.skymaps.info>. The catalogs contain more than twenty billion detections of more than two billion objects.

7.1. Viewer

DECaPS images have been imported into the DECam Legacy Survey viewer (DESI Collaboration et al. 2016a, 2016b). Zero-points from the photometric calibration (see Section 5) are applied to each image, and a single constant sky per CCD is removed, according to the median sky estimate of the `crowdsource` analysis. Images are then projected onto a common frame and coadded with inverse-variance weighting to produce stacks. Stacks in the *g*, *r*, and *z* bands are then combined to produce color images and are made available online in a zoomable and pannable viewer. Figure 13 shows an example image from the viewer.

Dramatic dust clouds are evident in the top half of the image, where nearly all stars are extremely red due to foreground extinction. A reflection nebula around one bright star lights up a large region with relatively blue reflected light. Several foreground stars are immediately identifiable from their bright, white colors, as compared with the red background stars. The greenish continuum emission at the bottom of the image is H α and reflected light.

In order to preserve many of these continuum features in the viewer, a single background value is subtracted from each DECam CCD before coaddition. This well preserves features with spatial scales much smaller than a CCD, but corrupts structures with a larger spatial scale. Accordingly, some of the brightest and largest objects in the sky become corrupted in their display in the viewer. For instance, near α Centauri and η Carinae, the per-CCD sky subtraction removes a significant amount of flux from the objects themselves and leads to a multicolored checkerboard pattern around them in the viewer. A more advanced sky-subtraction routine that matches the sky in overlapping images would solve this problem. However, since this problem only affects a small number of very bright objects and since stellar photometry is not affected, this issue has not been addressed.

The viewer contains analogous stacks for the `crowdsource` model images, and the resulting residual images. These allow for rapid comparison between the data and the model, to assess the reliability of the modeling in any region of the sky. In general, the model is an excellent description of the data.

7.2. Catalogs

DECaPS provides catalogs from individual images analyzed by `crowdsource` as well as derived, band-merged catalogs giving average photometry in each band for each object.

7.2.1. Individual-image Catalogs

Catalogs generated by `crowdsource` for each individual DECaPS image are available from the DECaPS website. These catalogs contain the position of each source, its flux, and the associated uncertainties. These uncertainties are those that would be obtained for isolated stars given photon and detector noise and a perfect model, and they therefore underestimate the true uncertainties. A few other entries assess the reliability of the detection, giving the χ^2 of the detection, and any flags at the location of the central pixel of this source in the CP data-quality image. A “quality factor” is also computed; it is the PSF-weighted fraction of pixels that have non-zero weight. A detection centered on the edge of a detector would have a quality factor of 0.5, while a good detection has a quality factor close to 1.

Blending is assessed via `fracflux`,

$$\text{fracflux} = \frac{\int du N(u)P(u)(\sigma^{-2}(u) > 0)}{\int du I(u)P(u)(\sigma^{-2}(u) > 0)}, \quad (3)$$

where P is the PSF and σ^{-2} is the CP inverse-variance image (0 for problematic pixels like saturated pixels). In this equation, N is a $5'' \times 5''$ postage stamp centered on the star, after subtracting the best-fit model of all other stars, and I is the

same image without subtracting other stars. The value of `fracflux` usually ranges from 0 for very blended objects to 1 for isolated objects. Spurious sources with negative fluxes can lead to other values in rare cases (see Section 4.6).

Additional fields give further metadata: the FWHM of the PSF at the location of each source, and the value of the sky and gain. These fields, and short descriptions, are provided in Table 4. A total of slightly over 20 billion records are included in the individual-image catalogs.

The catalogs contain flux estimates and “local-background-subtracted” flux estimates, where an additional constant background was left free to vary around each source. To compute these fluxes, each star is considered one by one, and the best model of all of the other stars is subtracted. The $5'' \times 5''$ region containing the now isolated star is fit with a two-parameter model: a flux, and a constant background. The resulting flux is recorded as the local-background-subtracted flux. In general, there is little difference between the default fluxes and the local-background-subtracted fluxes. This is expected because the usual sky subtraction operates on a $5''.2$ scale, very similar to the scale of the background subtraction performed here. The primary difference is that the sky subtraction here is simultaneously fit and constant over the region. Nevertheless, given the similarity between the local-background-subtracted and the ordinary fluxes, significant differences between the two can be used to indicate problems with the photometry.

7.2.2. Merged Catalogs

DECaPS also provides merged catalogs that join all detections of a single object on individual images into a single record. Photometry is then averaged by band, and an overall average position is measured. The resulting merged catalogs contain the best estimates of the average flux of the object. This merging process averages together images taken at different times, and *gr* observations are usually taken at significantly different times from *izY* observations. This makes the average magnitudes and colors somewhat ill-defined for variable stars; projects focusing on variable stars should make use of the individual image catalogs.

The merge is performed by considering exposures one by one. All detections in each image are compared with the list of currently existing objects and matched at a $0''.5$ radius. Matching detections are considered new detections of existing objects, and non-matching detections are considered the first detections of new objects.

With an assignment of individual detections to objects, the mean photometry of each object in the different bands is computed. The fluxes of the different detections are calibrated using the throughput model of Section 5. DECaPS then computes mean and median calibrated fluxes in each band for each object. Mean fluxes are weighted by the pipeline inverse-variance estimates, with a 1% floor. Additionally, 25% and 75% quantiles of the fluxes in each band are computed. Given that the typical object has only three photometric measurements in each band, this set of summary statistics is unnecessarily extensive, but this pipeline was originally developed for PS1, where more epochs were available in each filter. An average blendedness (`fracflux`) is also computed.

Table 5
DECaPS Flags

Bit	Meaning	Exclude?
1	Bad pixel	Y
3	Saturated	Y
4	Bleed trail	Y
5	Cosmic ray	Y
6	Low weight	Y
8	Long streak	Y
20	Additional bad pixel	Y
21	Nebulosity	N
22	S7 amplifier B	Y

Note. Flag bits used in DECaPS. Flag bits 1–8 are inherited directly from the DECam CP. Flag bits 20–22 are added by DECaPS. All flags prevent affected sources from having their fluxes included in the object average fluxes, except for the nebulosity flag (21).

Table 6
Merged Catalog Fields

Name	Description
ra	R.A. (deg)
dec	Decl. (deg)
posstddev	standard deviation in position (arcsec)
ra_ok	R.A. (ok detections only) (deg)
dec_ok	Decl. (ok detections only) (deg)
posstddev_ok	standard deviation (ok detections only) (arcsec)
ndet	number of detections
ndet_ok	number of ok detections
nmag	number of detections in each band
nmag_ok	number of okay detections in each band
mean	mean flux (units of 3631 Jy)
stdev	flux standard deviation (units of 3631 Jy)
err	uncertainty in mean flux (units of 3631 Jy)
median	median flux (units of 3631 Jy)
q25	25th percentile flux (units of 3631 Jy)
q75	75th percentile flux (units of 3631 Jy)
epochrange	MJD between first and last detections
epochrange_ok	MJD between first and last ok detections
epochmean	average MJD
epochmean_ok	average MJD (ok only)
maglimit	5 σ magnitude limit of deepest image (ok only)
fracflux	average fractional flux in each band (ok only)

Note. Fields in the DECaPS merged catalogs. A more complete description is available at the survey website. Most fields are 32-bit floating point numbers, except for fields pertaining to epochs and positions, which are 64-bit floating point numbers. The fields describing the number of detections of sources are 16-bit integers.

Only detections of sources considered “ok” are included in the averages. An ok detection has no problematic flags (Table 5), and has a “quality factor” greater than 0.85 (see Section 7.2.1). A detection with the NEBULOSITY flag set may be considered ok, since excluding these detections would exclude all measurements of objects in regions with substantial nebulosity.

The object catalog contains entries for slightly over two billion objects. Table 6 lists and briefly describes the fields included in the catalog.

Fluxes in the merged catalogs are given in units of 3631 Jy (equivalently, in Mgy, Finkbeiner et al. 2004). This is intended to simplify conversion to AB magnitudes; the AB magnitude is $AB = -2.5 \log_{10} \text{flux}$.

Additional statistics about the flux are available in columns labeled *lbs*, indicating that they refer to local-background-subtracted fluxes (see Section 7.2.1).

8. Conclusion

The DECam Plane Survey comprises *grizY* imaging of the southern Galactic plane, $5^\circ > l > -120^\circ$, $|b| \lesssim 4^\circ$, to depths between 24th and 21st mag. The photometric calibration is accurate to 1%–2%. Crowded field photometry is used to identify and photometer more than 20 billion detections of 2 billion sources. Raw and processed images are available online through the NOAO Archive. Individual-image and band-merged object catalogs are available through the survey website. DECaPS coadd images, model images, and residuals images are available through the DECam Legacy Survey viewer for rapid inspection.

In combination with PS1, DECaPS provides *grizY* coverage of 360° of the Milky Way midplane. DECaPS can detect main-sequence turn-off stars of old populations out to beyond 8 kpc when the dust reddening $E(B - V) < 1.5$ mag, probing a large portion of the Milky Way. Existing longer wavelength surveys such as VVV and GLIMPSE are natural counterparts to DECaPS, with less dust-sensitive bands that can see through denser regions of the interstellar medium.

We anticipate that the DECaPS photometry will serve as the foundation for new maps of stars and dust in the southern Galactic plane, complementing maps of the northern Galactic plane like that of Green et al. (2015). The combination of DECaPS photometry and APOGEE-2 spectroscopy (Abolfathi et al. 2017) will similarly enable the shape of the dust extinction curve and its variation to be measured throughout the same volume (e.g., Schlafly et al. 2017). Inaccurate dust maps place important limitations on the accuracy that can be achieved by the *Gaia* mission (Bovy et al. 2016), motivating deeper photometry and further effort in this area to realize the potential of *Gaia*’s Milky Way studies. We are excited to learn what other uses the community can find for the billions of stars cataloged by DECaPS.

E.S. acknowledges support for this work provided by NASA through Hubble Fellowship grant *HST*-HF2-51367.001-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. G.M.G. and D.P.F. acknowledge support from NSF grant AST-1312891. Travel to CTIO was partially supported by NSF grant AST-1614941.

DECaPS is based on observations at Cerro Tololo Inter-American Observatory, National Optical Astronomy Observatory (NOAO Prop. ID: 2014A-0429, 2016A-0327, and 2016B-0279; PI: Finkbeiner), which is operated by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation.

The DECaPS 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 DECaPS nebosity CNN was trained on the XStream computational resource, supported by the National Science Foundation Major Research Instrumentation program (ACI-1429830).

This project used data obtained with the Dark Energy Camera (DECam), which was constructed by the Dark Energy Survey (DES) collaboration. Funding for the DES Projects has been provided by the U.S. Department of Energy, the U.S. National Science Foundation, the Ministry of Science and Education of Spain, the Science and Technology Facilities Council of the United Kingdom, the Higher Education Funding Council for England, the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, the Kavli Institute of Cosmological Physics at the University of Chicago, the Center for Cosmology and Astro-Particle Physics at the Ohio State University, the Mitchell Institute for Fundamental Physics and Astronomy at Texas A&M University, Financiadora de Estudos e Projetos, Fundação Carlos Chagas Filho de Amparo à Pesquisa do Estado do Rio de Janeiro, Conselho Nacional de Desenvolvimento Científico e Tecnológico and the Ministério da Ciência, Tecnologia e Inovação, the Deutsche Forschungsgemeinschaft, and the Collaborating Institutions in the Dark Energy Survey. The Collaborating Institutions are Argonne National Laboratory, the University of California at Santa Cruz, the University of Cambridge, Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas-Madrid, the University of Chicago, University College London, the DES-Brazil Consortium, the University of Edinburgh, the Eidgenössische Technische Hochschule (ETH) Zürich, Fermi National Accelerator Laboratory, the University of Illinois at Urbana-Champaign, the Institut de Ciències de l'Espai (IEEC/CSIC), the Institut de Física d'Altes Energies, Lawrence Berkeley National Laboratory, the Ludwig-Maximilians Universität München and the associated Excellence Cluster Universe, the University of Michigan, the National Optical Astronomy Observatory, the University of Nottingham, the Ohio State University, the University of Pennsylvania, the University of Portsmouth, SLAC National Accelerator Laboratory, Stanford University, the University of Sussex, and Texas A&M University.

The Pan-STARRS1 Surveys (PS1) and the PS1 public science archive have been made possible through contributions by the Institute for Astronomy, the University of Hawaii, the Pan-STARRS Project Office, the Max Planck Society and its participating institutes, the Max Planck Institute for Astronomy, Heidelberg and the Max Planck Institute for Extraterrestrial Physics, Garching, The Johns Hopkins University, Durham University, the University of Edinburgh, the Queen's University Belfast, the Harvard-Smithsonian Center for Astrophysics, the Las Cumbres Observatory Global Telescope Network Incorporated, the National Central University of Taiwan, the Space Telescope Science Institute, the National Aeronautics and Space Administration under Grant No. NNX08AR22G issued through the Planetary Science Division of the NASA Science Mission Directorate, the National Science Foundation Grant No. AST-1238877, the University of Maryland, Eotvos Lorand University (ELTE),

the Los Alamos National Laboratory, and the Gordon and Betty Moore Foundation.

Facility: Blanco (DECam).

Appendix Nebosity Network Structure

Our CNN takes histogram-equalized 512×512 -pixel images as input, selected from the CP `InstCal` stage images. The network consists of 14 convolutional layers, interspersed with 6 maximum pooling layers. A global average pooling layer reduces the activations of the last maximum pooling layer to 32 activations, each representing a different learned feature in the input image. These 32 features are finally fed into a two-layer dense neural network, which classifies each image as one of our four hand-classified image types: NEBULOSITY, NEBULOSITY_LIGHT, NORMAL, or SKY_ERROR.

Each convolutional and dense layer is followed by a ReLU activation layer (Hahnloser et al. 2000). We use categorical cross-entropy as our loss function. In all, our neural network has 75352 trainable parameters. In order to avoid overfitting, we use L2 weight regularization in the convolutional layers and dropout in the dense layers.





Table 7 summarizes our network architecture.

Table 7
Network Architecture

Layer	Output Shape	Details
conv2d_1	$512 \times 512, 12$	5×5 , same padded
maxpool2d_1	$256 \times 256, 12$	2×2
conv2d_2	$256 \times 256, 24$	5×5 , same padded
maxpool2d_2	$128 \times 128, 24$	2×2
conv2d_3	$128 \times 128, 24$	$\begin{pmatrix} 3 \times 3, \text{same padded} \\ 3 \times 3, \text{same padded} \\ 1 \times 1 \end{pmatrix}$
maxpool2d_3	$64 \times 64, 24$	2×2
conv2d_4	$64 \times 64, 32$	$\begin{pmatrix} 3 \times 3, \text{same padded} \\ 3 \times 3, \text{same padded} \\ 1 \times 1 \end{pmatrix}$
maxpool2d_4	$32 \times 32, 32$	2×2
conv2d_5	$32 \times 32, 32$	$\begin{pmatrix} 3 \times 3, \text{same padded} \\ 3 \times 3, \text{same padded} \\ 1 \times 1 \end{pmatrix}$
maxpool2d_5	$16 \times 16, 32$	2×2
conv2d_6	$12 \times 12, 32$	$\begin{pmatrix} 3 \times 3 \\ 3 \times 3 \\ 1 \times 1 \end{pmatrix}$
maxpool2d_6	$6 \times 6, 32$	2×2
global_avg_pool2d	32	...
dense_1	12	20% dropout
dense_2	4	10% dropout
softmax	4	...

Note. All convolutional and dense layers use ReLU activation. The final output one-hot encodes the class, and the cross-entropy loss function is used.

ORCID iDs

E. F. Schlafly  <https://orcid.org/0000-0002-3569-7421>
 D. Lang  <https://orcid.org/0000-0002-1172-0754>
 T. Daylan  <https://orcid.org/0000-0002-6939-9211>
 D. P. Finkbeiner  <https://orcid.org/0000-0003-2808-275X>
 A. Lee  <https://orcid.org/0000-0002-1441-770X>
 A. M. Meisner  <https://orcid.org/0000-0002-1125-7384>
 D. Schlegel  <https://orcid.org/0000-0002-5042-5088>
 F. Valdes  <https://orcid.org/0000-0001-5567-1301>

References

- Abolfathi, B., Aguado, D. S., Aguilar, G., et al. 2017, arXiv:1707.09322
 Benjamin, R. A., Churchwell, E., Babler, B. L., et al. 2003, *PASP*, **115**, 953
 Bohlin, R. C. 2014, *AJ*, **147**, 127
 Bovy, J., Rix, H.-W., Green, G. M., Schlafly, E. F., & Finkbeiner, D. P. 2016, *ApJ*, **818**, 130
 Brewer, B. J., Foreman-Mackey, D., & Hogg, D. W. 2013, *AJ*, **146**, 7
 Burke, D., Rykoff, E. S., Allam, S., et al. 2017, arXiv:1706.01542
 Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv:1612.05560
 Churchwell, E., Babler, B. L., Meade, M. R., et al. 2009, *PASP*, **121**, 213
 Daylan, T., Portillo, S. K. N., & Finkbeiner, D. P. 2017, *ApJ*, **839**, 4
 DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016a, arXiv:1611.00036
 DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016b, arXiv:1611.00037
 Dolphin, A. E. 2000, *PASP*, **112**, 1383
 Drew, J. E., Gonzalez-Solares, E., Greimel, R., et al. 2014, *MNRAS*, **440**, 2036
 Drew, J. E., Greimel, R., Irwin, M. J., et al. 2005, *MNRAS*, **362**, 753
 Finkbeiner, D. P., Padmanabhan, N., Schlegel, D. J., et al. 2004, *AJ*, **128**, 2577
 Flaugher, B., Diehl, H. T., Honscheid, K., et al. 2015, *AJ*, **150**, 150
 Green, G. M., Schlafly, E. F., Finkbeiner, D. P., et al. 2015, *ApJ*, **810**, 25
 Hahnloser, R. H., Sarpeshkar, R., Mahowald, M. A., Douglas, R. J., & Seung, H. S. 2000, *Natur*, **405**, 947
 Hanson, R. J., Bailer-Jones, C. A. L., Burgett, W. S., et al. 2016, *MNRAS*, **463**, 3604
 High, F. W., Stubbs, C. W., Rest, A., Stalder, B., & Challis, P. 2009, *AJ*, **138**, 110
 Juric, M. 2012, LSD: Large Survey Database Framework, Astrophysics Source Code Library, ascl:1209.003
 Lang, D., Hogg, D. W., & Schlegel, D. J. 2016, *AJ*, **151**, 36
 LeCun, Y., Bengio, Y., & Hinton, G. 2015, *Natur*, **521**, 436
 Li, T. S., DePoy, D. L., Marshall, J. L., et al. 2016, *AJ*, **151**, 157
 Lucas, P. W., Hoare, M. G., Longmore, A., et al. 2008, *MNRAS*, **391**, 136
 Magnier, E. A., Schlafly, E., Finkbeiner, D., et al. 2013, *ApJS*, **205**, 20
 Magnier, E. A., Schlafly, E. F., Finkbeiner, D. P., et al. 2016a, arXiv:1612.05242
 Magnier, E. A., Sweeney, W. E., Chambers, K. C., et al. 2016b, arXiv:1612.05244
 Minniti, D., Lucas, P. W., Emerson, J. P., et al. 2010, *NewA*, **15**, 433
 Moffat, A. F. J. 1969, *A&A*, **3**, 455
 Morganson, E., Conn, B., Rix, H.-W., et al. 2016, *ApJ*, **825**, 140
 Murphy, D. C., & May, J. 1991, *A&A*, **247**, 202
 Oke, J. B., & Gunn, J. E. 1983, *ApJ*, **266**, 713
 Padmanabhan, N., Schlegel, D. J., Finkbeiner, D. P., et al. 2008, *ApJ*, **674**, 1217
 Paige, C. C., & Saunders, M. A. 1982, *ACM Trans. Math. Software*, **8**, 43
 Plazas, A. A., Bernstein, G. M., & Sheldon, E. S. 2014, *PASP*, **126**, 750
 Portillo, S. K. N., Lee, B. C. G., Daylan, T., & Finkbeiner, D. P. 2017, arXiv:1703.01303
 Sandrelli, S., Bragaglia, A., Tosi, M., & Marconi, G. 1999, *MNRAS*, **309**, 739
 Schlafly, E. F., Finkbeiner, D. P., Jurić, M., et al. 2012, *ApJ*, **756**, 158
 Schlafly, E. F., Finkbeiner, D. P., Schlegel, D. J., et al. 2010, *ApJ*, **725**, 1175
 Schlafly, E. F., Green, G., Finkbeiner, D. P., et al. 2015, *ApJ*, **799**, 116
 Schlafly, E. F., Meisner, A. M., Stutz, A. M., et al. 2016, *ApJ*, **821**, 78
 Schlafly, E. F., Peek, J. E. G., Finkbeiner, D. P., & Green, G. M. 2017, *ApJ*, **838**, 36
 Scolnic, D., Casertano, S., Riess, A., et al. 2015, *ApJ*, **815**, 117
 Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, **131**, 1163
 Slater, C. T., Bell, E. F., Schlafly, E. F., et al. 2014, *ApJ*, **791**, 9
 Stetson, P. B. 1987, *PASP*, **99**, 191
 Tonry, J. L., Stubbs, C. W., Lykke, K. R., et al. 2012, *ApJ*, **750**, 99
 Valdes, F., Gruendl, R., & DES Project 2014, in ASP Conf. Ser. 485, Astronomical Data Analysis Software and Systems XXIII, ed. N. Manset & P. Forshay (San Francisco, CA: ASP), 379
 Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, *AJ*, **140**, 1868