

3 LSST System Performance

Steven M. Kahn, Justin R. Bankert, Srinivasan Chandrasekharan, Charles F. Claver, A. J. Connolly, Kem H. Cook, Francisco Delgado, Perry Gee, Robert R. Gibson, Kirk Gilmore, Emily A. Grace, William J. Gressler, Željko Ivezić, M. James Jee, J. Garrett Jernigan, R. Lynne Jones, Mario Jurić, Victor L. Krabbandam, K. Simon Krughoff, Ming Liang, Suzanne Lorenz, Alan Meert, Michelle Miller, David Monet, Jeffrey A. Newman, John R. Peterson, Catherine Petry, Philip A. Pinto, James L. Pizagno, Andy Rasmussen, Abhijit Saha, Samuel Schmidt, Alex Szalay, Paul Thorman, Anthony Tyson, Jake VanderPlas, David Wittman

In this chapter, we review the essential characteristics of the LSST system performance. We begin with descriptions of the tools that have been developed to evaluate that performance: the Operations Simulator (§ 3.1), the Exposure Time Calculator (§ 3.2), the Image Simulator (§ 3.3), and raytrace calculations used to evaluate stray and scatter light (§ 3.4). We then discuss the expected photometric accuracy that will be achieved (§ 3.5), and the expected accuracy of trigonometric parallax and proper motion measurements (§ 3.6). Next, we provide estimates of discrete source counts in the main LSST survey, both for stars in the Milky Way (§ 3.7.1), and for galaxies as a function of redshift (§ 3.7.2). We conclude with a discussion of the accuracy with which redshifts of galaxies can be determined from LSST photometry (§ 3.8).

3.1 Operations Simulator

Philip A. Pinto, R. Lynne Jones, Kem H. Cook, Srinivasan Chandrasekharan, Francisco Delgado, Željko Ivezić, Victor L. Krabbandam, K. Simon Krughoff, Michelle Miller, Cathy Petry, Abhijit Saha

During its ten-year survey, LSST will acquire ~ 5.6 million 15-second images, spread over ~ 2.8 million pointings. Their distribution on the sky, over time, and among its six filters has a strong impact on how useful these data are for almost any astronomical investigation. The LSST Project has developed a detailed operations simulator (LSST OpSim : <http://www.noao.edu/lsst/opsim>) in order to develop algorithms for scheduling these exposures – addressing the question “what observation should be made next?” – and to quantitatively evaluate the observing strategies discussed in § 2.1. These algorithms will become a fundamental component of the LSST design, as part of the scheduler driving the largely robotic observatory. In addition, the simulator will remain an important tool allowing LSST to adapt and evaluate its observing strategy in response to the changing scientific demands of the astronomical community.

The operations simulator incorporates detailed models of the site conditions and hardware performance, as well as the algorithms for scheduling observations. It creates realizations of the set of visits (back-to-back 15 second exposures in a given filter) that the LSST will make during a

ten-year survey, this being the primary output of the OpSim. These outputs include the position on the sky, time, and filter of each visit, and the signal-to-noise ratio achieved. These outputs can be further processed to generate estimates of the depth of the final stacked images in each filter as a function of position on the sky (Figure 2.1), histograms of the airmass distribution of visits (Figure 3.3), or other figures of merit relevant to particular science goals.

The simulation of observing conditions includes a model for seeing drawn from observing records at Cerro Tololo (Figure 2.3). This model is consistent with the auto-correlation spectrum of seeing with time over intervals from minutes to seasons as measured at the site on Cerro Pachón. Weather data, including their correlations, are taken from ten years of hourly measurements made at nearby Cerro Tololo. The 5σ PSF depth of each observation is determined using a sky background model which includes the dark sky brightness in each filter passband, the effects of seeing and atmospheric transparency, and an explicit model for scattered light from the Moon and/or twilight at each observation.

The time taken to slew the telescope from one observation to the next is given by a detailed model of the camera, telescope, and dome. It includes such effects as the acceleration/deceleration profiles employed in moving in altitude, azimuth, camera rotator, dome azimuth, and wind/stray light screen altitude, the time taken to damp vibrations excited by each slew, cable wrap, and the time taken for active optics lock and correction as a function of slew distance, filter changes, and focal plane readout. The detail of this model ensures an accurate relation between system parameters and modeled performance, making the operations simulator a valuable tool for optimizing design.

After each visit, all possible next visits are assigned a score according to a set of scientific requirements, which depend upon the current conditions and the past history of the survey. For example, if a location in the ecliptic has been observed in the r -band, the score assigned to another r -band visit to the same location will initially be quite low, but it will rise with time to peak about an hour after the first observation, and decline thereafter. This results in these observations being acquired as pairs of visits roughly an hour apart, enabling efficient association of near-Earth object (NEO) detections. To ensure uniform sky coverage, locations on the sky with fewer previous visits will be scored more highly than those observed more frequently. Observations with higher expected signal-to-noise ratio are ranked more highly, leading to a majority of visits being made near the local meridian, and a tendency for visits in redder bands to be made closer to twilight and at brighter phases of the Moon. Higher scores are given to observations in the r - and i -bands during periods of better seeing to aid in shape determination for weak lensing studies.

Once all possible next visits have been ranked for scientific priority, their scores are modified according to the cost of making the observation. Visits to locations which require more slew time are penalized, as are those which require filter changes, unwrapping cables in the camera rotator, and so on. After this modification according to cost, the highest-ranked observation is performed, and the cycle repeats. The result of a simulator run is a detailed history of which locations have been observed when, in what filter, and with what sky backgrounds, airmass, seeing, and other observing conditions. A detailed history of all telescope, dome, and camera motions is also produced for engineering studies.

Each of the two exposures in a visit requires 16 seconds to complete; while every pixel is exposed for 15 seconds, the shutters require one second to traverse the entire 63 cm of the active area in the focal plane. Two seconds are required to read out the array between exposures. After the

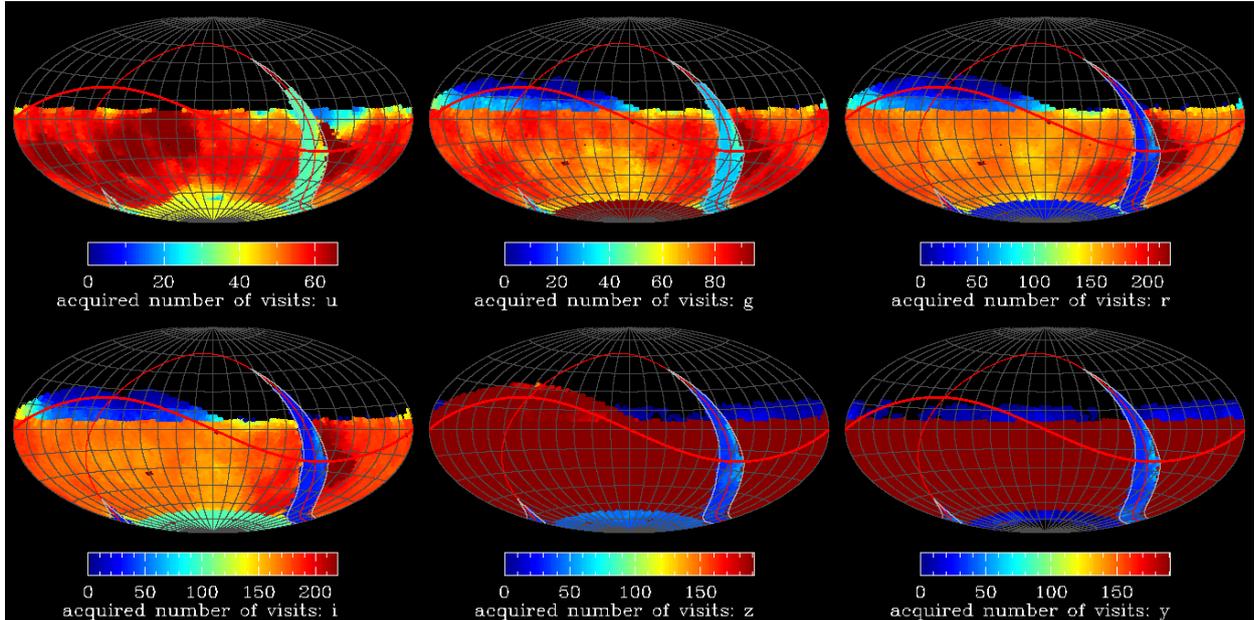


Figure 3.1: The number of visits in one realization of a simulated ten-year survey in all six LSST filters, shown in Equatorial coordinates. The project goals are to have 56, 80, 180, 180, 164, and 164 visits in the u , g , r , i , z , y filters, respectively, over $20,000 \text{ deg}^2$ of sky. One of the deep-drilling fields is apparent at $\alpha = 90^\circ$, $\delta = -32^\circ$.

second exposure, a minimum of five seconds is required to slew to an adjacent location on the sky, settle, and acquire active optics lock and correction, during which time the array is read out for the second time. Thus a complete visit to an adjacent field, with no filter change, takes a minimum of 39 seconds to perform; this amounts to spending 87% of the time exposing the detector to the sky. This, of course, does not take into account the time spent in changing filters (two minutes per change) or any of the scientific requirements on scheduling. In one specific realization of the full ten-year survey, 80% of the available time (i.e., when weather permitted) was spent exposing on the sky, which is about 92% of the naïve estimate above.

Figure 3.1 shows the number of visits across the sky in this simulation, while Figure 3.2 shows the 5σ limiting magnitude for point sources achieved in the stacked images. Figure 3.3 shows a histogram of the air-mass and seeing delivered during observations in each filter.

The current output from the OpSim assumes each visit is taken with the field centers placed onto a fixed grid on an optimally packed tessellation. This gives a variation of the effective depth across the sky, as is shown in the dashed line in Figure 3.4. To evaluate the effects of dithering on LSST performance, we simply added a small (< 0.5 times the field of view) dithering pattern to the position of each pointing, leaving other aspects of the simulation unchanged. We added a different offset in right ascension (RA) and declination (dec) for each night of the survey, following a pattern which stepped through a hexagonal grid on the field of view. This dithering makes the coverage substantially more uniform, as is shown by the solid line in Figure 3.4.

We are continuing to work on developing improved scheduling algorithms, replacing the algorithm which simply observes the field with the highest score at each step with one which looks ahead for a few hours, using a path optimization algorithm to further reduce the slew time required,

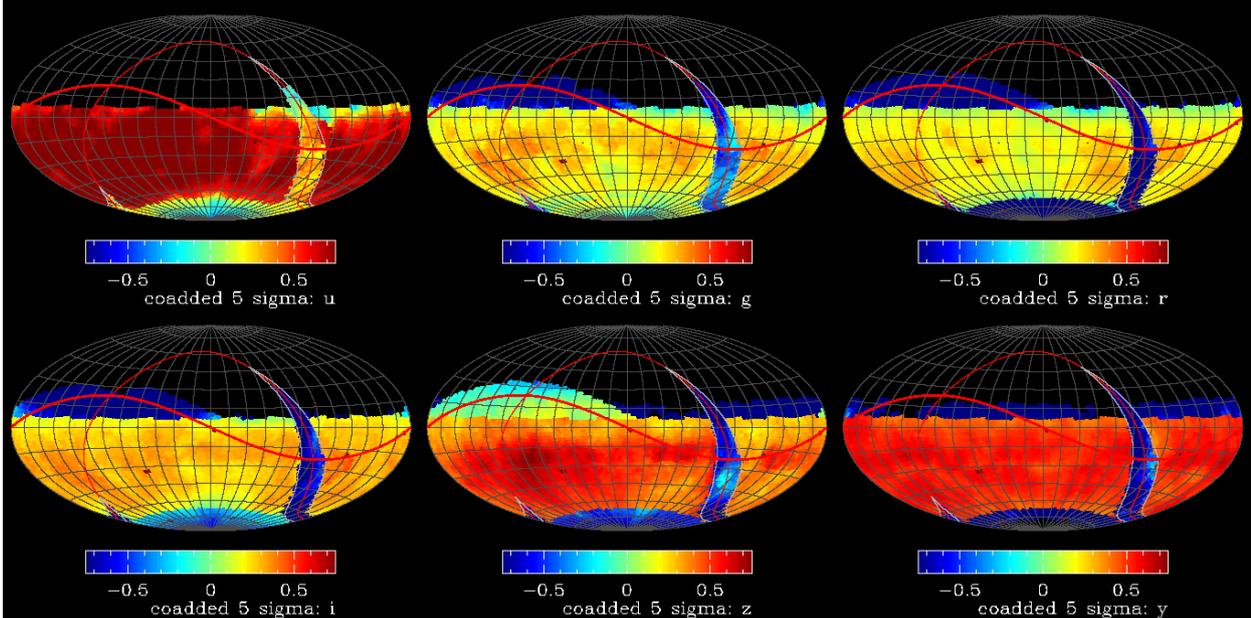


Figure 3.2: The 5σ stacked point-source depth of the simulated ten-year survey shown in Figure 3.1. The scale in each panel shows the depth of the stack relative to the fiducial values of 25.8, 27.0, 27.2, 27.0, 25.7, and 24.4 in u , g , r , i , z , y respectively.

and including more feedback from science metrics (on already acquired data) into the scheduling algorithms. We are also working with the LSST Science Collaborations to refine our current cadences to enhance the utility of the LSST data set for the widest possible applicability.

3.2 Exposure Time Calculator

Anthony Tyson, Perry Gee, Paul Thorman

In order to enable fast predictions of the signal-to-noise ratio for the detection of both point and diffuse sources, we have developed an Exposure Time Calculator (ETC; <http://lsst.org/etc>). The ETC incorporates models of the extinction, telescope and camera optics, detector efficiency, and filter response to calculate the throughput of the system in each band. It uses a sky brightness model based on data taken at CTIO, United Kingdom Infra-red Telescope (UKIRT), and SDSS.

An input source model is shifted to the correct redshift and normalized to a selected brightness or surface brightness. The resulting flux density is multiplied by the system response as a function of wavelength for a given filter band to produce a predicted photon count-rate within a specified aperture. The integral sky brightness is also calculated for the same aperture, so that the signal-to-noise ratio for detection can be calculated. The aperture is fully adjustable, and an option for PSF-weighted photometry is also provided.

The ETC allows the source spectral energy distribution, surface brightness profile, the extinction, and the redshift to be varied, and includes a library of stellar and extragalactic source spectra. For specified seeing, Moon and cloud conditions, and for multiple exposures of a specified time

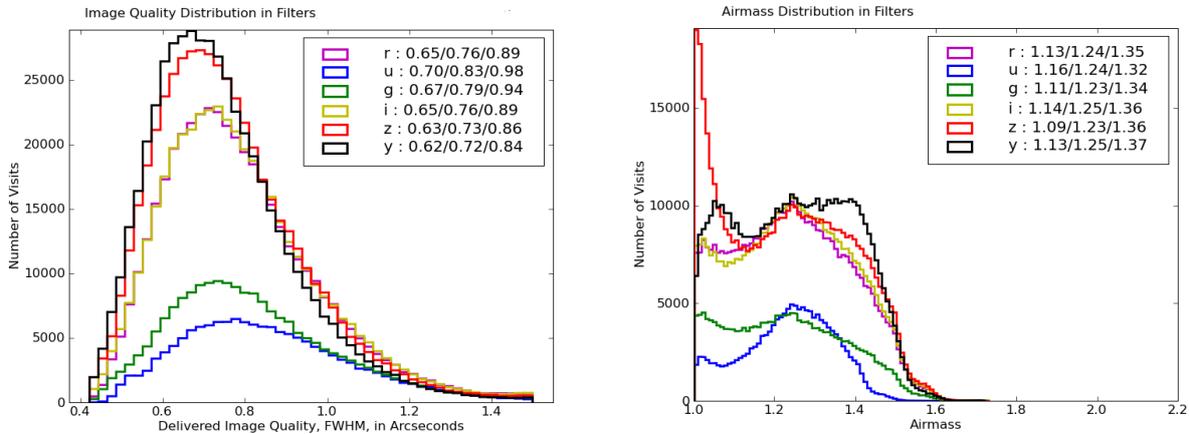


Figure 3.3: Histograms of the delivered seeing and airmass distributions for all visits in the simulated survey shown in Figure 3.1. Also shown in the legend are the 25-, 50-, and 75-percentile values in each filter.

and fraction of exposure on source, the ETC estimates both the signal-to-noise ratio for a single exposure and the exposure time required to achieve a threshold signal-to-noise ratio.

In Table 3.1 we provide the predicted signal-to-noise ratios (S/N) for some canonical source types. For each object, we quote S/N based on a single visit, and on the full ten-year survey. The calculations in the table are based on $z = 0$ template spectra of galaxies, quasars, and stars, without evolution (although the absolute magnitudes of galaxies at each redshift are rough estimates of M^*). The quoted S/N includes sky subtraction and PSF-optimized filtering for galaxies of typical angular size at the given redshift, but no provision for other systematic errors (thus values of S/N more than several hundred should be taken with a grain of salt). The sky background was estimated assuming three-day-old lunar phase and solar minimum sky levels. The seeing was assumed to be $0.7''$ in r -band with clear skies. For the high-redshift quasars, no entries are given below the Lyman limit; the flux is taken to be essentially zero at shorter wavelengths.

The ETC also allows estimates of the saturation limits of the LSST camera. In $0.7''$ seeing under photometric skies, and for a 15 sec exposure, the detectors will saturate with a star of $u, g, r, i, z, y = 14.7, 15.7, 15.8, 15.8, 15.3$ and 13.9 , respectively.

3.3 Image Simulator

John R. Peterson, J. Garrett Jernigan, Justin R. Bankert, A. J. Connolly, Robert R. Gibson, Kirk Gilmore, Emily A. Grace, M. James Jee, R. Lynne Jones, Steven M. Kahn, K. Simon Krughoff, Suzanne Lorenz, Alan Meert, James L. Pizagno, Andrew Rasmussen

The project team has developed a detailed Image Simulator (<http://lsst.astro.washington.edu>) to evaluate the system sensitivity for particular science analyses, and to test data analysis pipeline performance on representative mock data sets. The simulated images and catalogs that it produces extend to $r = 28$ (deeper than the expected ten year depth of the LSST stacked images). These have proven useful in designing and testing algorithms for image reduction, evaluating the

Table 3.1: Typical apparent magnitudes and Signal-to-Noise ratios, S/N

Object	u		g		r		i		z		y	
	mag	S/N visit/full										
Stars												
O5V, 100 kpc	18.6	190/1600	19.0	250/3500	19.6	240/3600	20.0	140/2200	20.2	65/930	20.5	17/250
A0V, 100 kpc	21.0	46/380	20.3	170/1700	20.5	140/2100	20.7	86/1300	20.9	39/560	20.6	13/180
G2V, 100 kpc	25.9	0.8/6.6	24.7	5.6/56	24.3	6.6/100	24.1	4.6/69	24.1	2.2/31	24.0	0.7/9.9
K4III, 100 kpc	23.5	6.6/55	20.6	154/1500	19.4	260/3900	19.0	270/4000	18.8	200/2800	18.3	95/1300
M3V, 1 kpc	25.6	1.1/8.9	23.0	23/230	21.8	55/830	20.6	92/1400	20.0	79/1100	19.3	42/590
L1 Dwarf, 500 pc	-	-	-	-	27.1	0.48/7.4	24.8	2.4/37	23.2	4.8/69	22.2	4.2/59
Elliptical												
$z = 0.5, M_B = -20.8$	24.2	3.4/28	22.8	31/330	21.1	91/1400	20.2	130/1900	19.8	93/1300	19.5	47/670
$z = 1, M_B = -21.3$	25.3	1.3/11	25.0	4.2/45	23.7	11/170	22.6	17/260	21.6	20/280	21.0	13/190
$z = 2, M_B = -21.9$	25.7	0.9/7.3	25.6	2.2/23	25.6	2.0/30	25.2	1.6/24	24.5	1.4/20	23.6	1.2/17
Spiral												
$z = 0.5, M_B = -20.9$	22.6	14/120	22.0	59/630	20.8	110/1700	20.2	120/1800	20.2	67/950	19.5	43/610
$z = 1, M_B = -21.2$	23.7	5.7/47	23.6	15/160	23.0	19/290	22.1	26/390	21.7	18/250	21.0	12/170
$z = 2, M_B = -21.9$	23.6	6.9/57	24.0	9.9/110	23.9	8.8/130	23.7	6.3/95	23.5	3.7/52	23.1	1.9/27
Quasar												
$z = 1, M_B = -24.4$	18.8	170/1400	18.7	430/4600	18.5	430/6500	18.5	360/5400	18.4	250/3500	18.2	140/2000
$z = 3, M_B = -26.0$	22.8	14/120	21.3	93/930	20.8	110/1700	20.7	84/1300	20.7	44/620	20.5	19/270
$z = 5, M_B = -26.0$	-	-	26.9	0.7/7.2	24.1	7.3/110	21.6	41/630	21.3	26/370	21.2	10/140
$z = 7, M_B = -26.0$	-	-	-	-	-	-	-	-	-	-	21.4	9.1/130

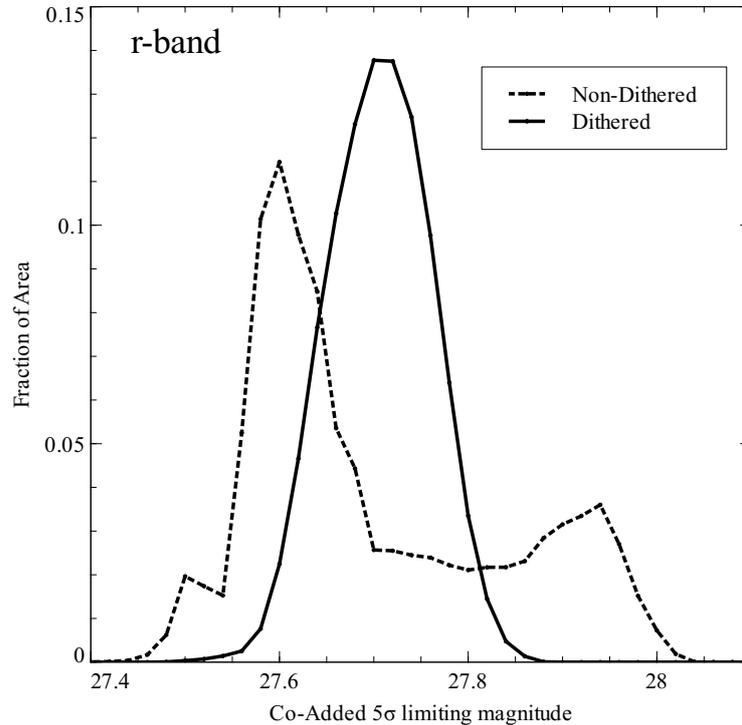


Figure 3.4: Histogram of the r band 5σ limiting magnitude of the ten-year stacked image depth. The histogram represents the stacked limiting magnitude evaluated over the full survey footprint on a grid with resolution of $0.6'$, much finer than the grid of field centers. The dashed line indicates the 5σ stacked depth in the non-dithered simulation, with two discrete peaks where fields are not overlapped (peak near 27.6 mag) and where they do overlap (peak near 27.9). The solid line indicates the 5σ stacked depth evaluated in the same simulation, with dithering added to each field’s central position. Dithering increases the median 5σ stacked depth by approximately 0.2 magnitudes.

capabilities and scalability of the analysis pipelines, testing and optimizing the scientific returns of the LSST survey, and providing realistic LSST data to the science collaborations. Figure 3.5 shows the flow of data through the LSST simulation framework.

The image simulator (Peterson et al. 2005) is a set of fast codes that begins with a catalog of objects (possibly lensed), and then traces photons through the atmosphere and the refractive and reflective optics, and into the detector where they photo-convert into electrons. The simulation can have extremely high fidelity in that all wavelength-dependent optics, detector, and atmospheric effects can be readily incorporated. The code is also sufficiently fast that a single 15-second 3.2 Gigapixel image from the LSST camera can be simulated in ~ 6 or 7 hours using a large pool of parallel machines.

The simulator constructs catalogs of objects drawn from cosmological and Galactic structure models (*base catalogs*), which are then used to generate a view of the sky above the atmosphere. These base catalogs include the spectral, photometric, astrometric, and morphological properties of the astronomical sources. The base catalogs are queried based on simulated observation sequences from the Operations Simulator (§ 3.1) creating a series of *instance catalogs*. Given the time, RA and Dec of each pointing, the appropriate airmass, sky background, and observing conditions can be determined.

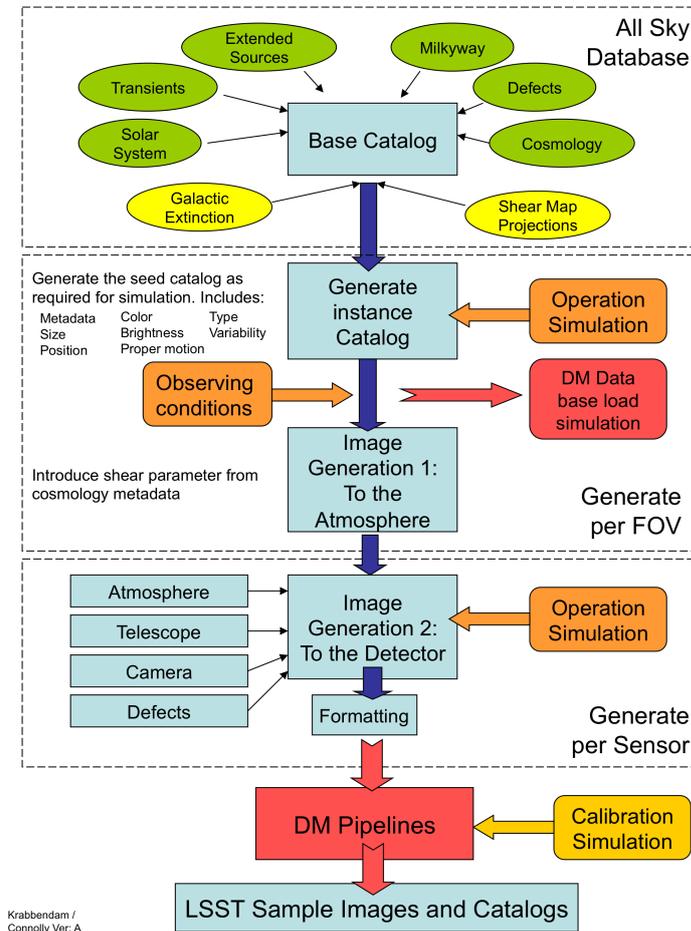


Figure 3.5: The flow of information through the image simulation. The top level describes a static view of the sky that is sampled to provide instance catalogs (based on the operations simulations, § 3.1). These catalogs are then passed into the Image Simulator resulting in a set of FITS images and catalogs.

Each object in the catalog has a sky position, magnitude at some wavelength, and spectral energy distribution (SED) file to determine the probabilities with which wavelengths are assigned to photons. Each object can either be represented by a set of parameters describing a spatial model or a hyper-resolved image to determine the probability distribution of where the photons are emitted. Additionally, objects can vary in flux during the exposure, they can move during the exposure (in the case of Solar System objects), or can be distorted due to gravitational lensing. Photons are drawn from this catalog in proportion to their magnitude and both the SED and spatial model are sampled. In this way, photons are chosen one at a time with two sky positions, a wavelength, and a time.

Galaxy positions and properties in the simulations are taken from the Millennium cosmological Simulation, with baryon physics included following [Kitzbichler & White \(2007\)](#). Galaxy SEDs use [Bruzual & Charlot \(2003\)](#) models, giving apparent magnitudes in all the LSST bands. Every galaxy is further assigned a realistic morphological profile via a disk-to-total flux ratio, position angle in the sky, inclination along the line-of-sight, bulge radius, and disk radius. More accurate galaxy

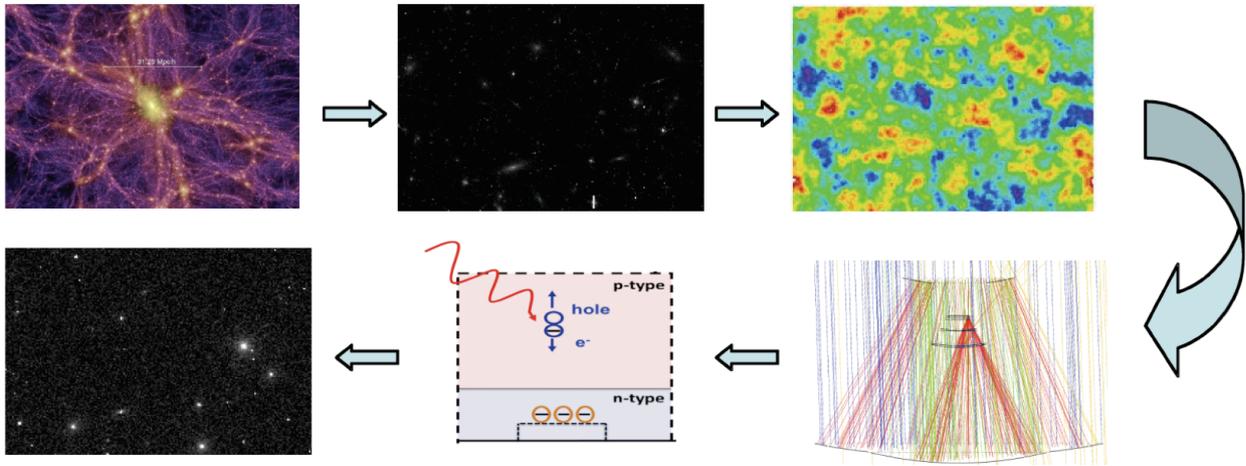


Figure 3.6: A schematic of the key steps leading to the production of a simulated image. First, a cosmological simulation is used to produce a three-dimensional dark matter map of a limited region of sky (upper left). This is then decorated with galaxies, which, along with a set of stars generated from an associated Milky Way model, are collected into a catalog of objects in the field (upper middle). This catalog is sampled to generate Monte Carlo photons in angle and color, which are propagated through a set of turbulent atmospheric screens (upper right) that move as a function of time according to input wind velocity vectors. Photons are then reflected and refracted through the mirrors and lenses of the LSST optics with an assumed set of displacements and distortions (lower right), and propagated into the detector (lower middle) where they convert to photoelectrons detected in a pixel. Background sky counts are added to produce the final simulated image of a single 15-second exposure at the lower left.

profiles, including high-frequency spatial structure such as H II regions and spiral arms, can be simulated using FITS images as input into the Image Simulator. The use of more detailed galaxy morphological profiles in the Image Simulator will allow LSST to study how galaxy morphology varies with environment and redshift.

Currently, stars are included in the Image Simulator with full SEDs, spatial velocities, and positions. The SEDs for stars are derived from Kurucz models. The model used to generate main sequence stars is based on work done by Mario Jurić and collaborators. The model includes user-specified amounts of thick-disk, thin-disk, and halo stars. Each version of a catalog contains metadata on metallicity, temperature, luminosity-type, and surface gravity, allowing the user to search for correlations between observed LSST photometry and physical information about stars using the simulated data. The catalog will be updated to include dwarf and giant stars.

After the photons are selected from the astronomical source list, they are propagated through the atmosphere and are refracted due to atmospheric turbulence. The model of the atmosphere is constructed by generating roughly half a dozen atmospheric screens as illustrated in Figure 3.6. These model screens incorporate density fluctuations following a Kolmogorov spectrum, truncated both at an outer scale (typically known to be between 20 m and 200 m) and at an inner scale (representing the viscous limit). In practice the inner scale does not affect the results. The screens are moved during the exposure according to wind velocity vectors, but, consistent with the well-established “frozen-screen approximation,” the nature of the turbulence is assumed to stay approximately fixed during the relatively short time it takes for a turbulent cell to pass over the aperture. With these screens, we start the photons at the top of the atmosphere and then alter their trajectory according to the refractions of the screen at each layer. The part of the screen

that a given photon will hit depends on the time that photon is emitted in the simulation and the wind vector of the screen.

After passing through the atmosphere, photons enter the telescope and are appropriately reflected and refracted as they hit the mirrors and lenses. On the surface of the mirrors we introduce a spectrum of perturbations that has been obtained by inverting wavefront data from existing telescopes. We also allow the orientation of each optic to be perturbed in six degrees of freedom within expected tolerances. The raytrace uses extremely fast techniques to find the intercepts on the aspheric surface and alter the trajectory by reflection or wavelength-dependent refraction. Photons can be “destroyed” as they pass through the filter in a Monte Carlo sense with a probability related to the wavelength and angle-dependent transmission function. The raytrace for the LSST configuration is illustrated in [Figure 3.6](#). The raytrace has been compared with commercial raytrace codes and is found to be accurate to a fraction of a micron. We also incorporate diffraction spikes appropriate for the design of the spider of the telescope.

In the last step, photons are ray-traced into the silicon in the detector. Both the wavelength and temperature dependent conversion probability and refraction at the interface of the silicon are included. The photons are then converted into photoelectrons which drift to the readout electrodes according to the modeled electric field profile. The misalignments and surface roughness of the silicon can also be included. The positions of the photoelectrons are pixelated and can include blooming, charge saturation, cross-talk, and charge transfer inefficiency to simulate the readout process. Finally, a simulated image is built as the photoelectrons reach the readout. The read noise and sky background are added in a post-processing step. The sky background is generated based on an SED for the full Moon and an SED for the dark sky, with an empirically derived angular function for the Rayleigh scattering of the Moon’s light. The background is vignetted according to the results of raytrace simulations.

The simulator can generate about 22,000 photons per second on a typical workstation. For bright stars that saturate, it can simulate photons much faster since tricks can be used during the simulation to figure out if a pixel will saturate. Thus, we have a fast simulator with high fidelity. [Figure 3.7](#) shows images of stars with various components of the simulator turned on or off. [Figure 3.8](#) shows a simulated image from one LSST CCD.

3.4 Stray and Scattered Light

Charles F. Claver, Steven M. Kahn, William J. Gressler, Ming Liang

Stray and scattered light is a major concern given the extremely large field of view of LSST. There are two major categories of stray and scattered light: structured and diffuse. Structured stray light comes from diffraction, ghosts from the refractive optics, and semi-focused scattering from surfaces nearby the focal plane. Diffuse scattered light is generated from the micro-surface qualities of the mirrors, dust on the optical surfaces, and scattering from non-optical surfaces in the telescope and dome.

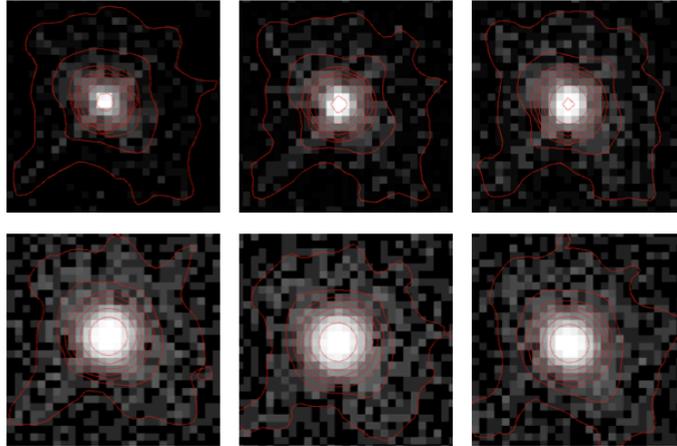


Figure 3.7: The capabilities of the simulator are demonstrated by examples of the point-spread function (PSF) for a single star 1.32° off-axis seen in the r filter, in which various components of the simulator are turned on or off. The images show a region $5.6''$ on a side, and the stretch is logarithmic. The panels are from top left to bottom right: only optical aberrations of the system design, adding mirror perturbation and misalignments, adding diffusion of charge in the detector, adding a static atmosphere, adding an atmosphere with wind, and a full atmosphere but a perfect telescope. Both atmosphere and the optics contribute to the ellipticity of the PSF. The FWHM of the PSF with telescope, atmosphere and wind is about $0.6''$, with an ellipticity of 7%.

3.4.1 Structured Stray Light

The fast optical beam and physical geometry of LSST help to minimize the impact of structured stray light at the focal plane. The relatively small cross-section ($\sim 0.7\text{m}^2$) of the support vanes holding the secondary and camera assemblies results in very low intensity diffraction spikes. The diffraction spike in the r band (see [Figure 3.9](#)) is down by six orders of magnitude from the peak at a radius of $4''$.

Structured stray light from ghosting in the refractive elements is further reduced by using state-of-the-art broad band anti-reflection coatings. The relative surface brightness of the ghost images are calculated using optical ray tracing with the lens surface treated both transmissively (first pass) and reflectively (second pass); see [Figure 3.10](#). The reflective properties of the detector are assumed to be $1 - \text{QE}(\lambda)$. This overestimates the ghost brightness at the extreme red wavelength since the QE performance is dominated by the mean free path of the photon in silicon rather than the reflection at the surface. In any case, for any reasonably bright source in the LSST's field of view, the ghost image surface brightness will be well below that of the natural night sky.

3.4.2 Diffuse Scattered Light

The first line of defense for unwanted diffuse scattered light is the dome enclosure. LSST's dome, like most modern domes, is well ventilated to equalize the inside temperature with the exterior ambient temperature, and is also well-baffled to reject external sources of light. A key feature in the LSST dome vent design is light-restricting blackened louvers that have been aerodynamically optimized to minimize restriction in air flow. Light passing through the vents must scatter from



Figure 3.8: A simulated image of a 15-second exposure of one LSST CCD ($4K \times 4K$) with $0.2''$ pixels, $0.4''$ seeing and a field of view $13.7' \times 13.7'$, representing roughly 0.5% of the LSST focal plane. The brightest stars in the image are ~ 12 magnitude. An object of brightness ~ 33 magnitude would emit ~ 1 photon in a 15 second exposure. The image is a true color composite of three images, with the g , r , and i filters mapped into B, G, and R colors respectively. Each color channel is on a logarithmic intensity scale. In its ten-year survey, LSST will produce $\sim 2 \times 10^9$ single-band images of the same size.

at least two louver surfaces before entering the dome. Using a specialized coating (Aeroglaze Z306) these dome vents will allow $< 3\%$ of the incident light through, while having $> 95\%$ air flow efficiency. The wind screen within the dome slit will provide a circular aperture to restrict unwanted light outside the LSST's field of view from entering the dome. Even with these measures some light will naturally enter the dome and illuminate objects in a way that will create unwanted light at the focal plane. A detailed analysis using non-sequential ray tracing and three-dimensional

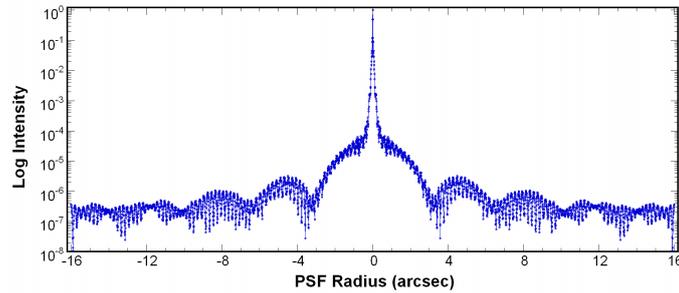


Figure 3.9: The log intensity of the r-band point-spread function along a diffraction spike. The plot spans $32''$.

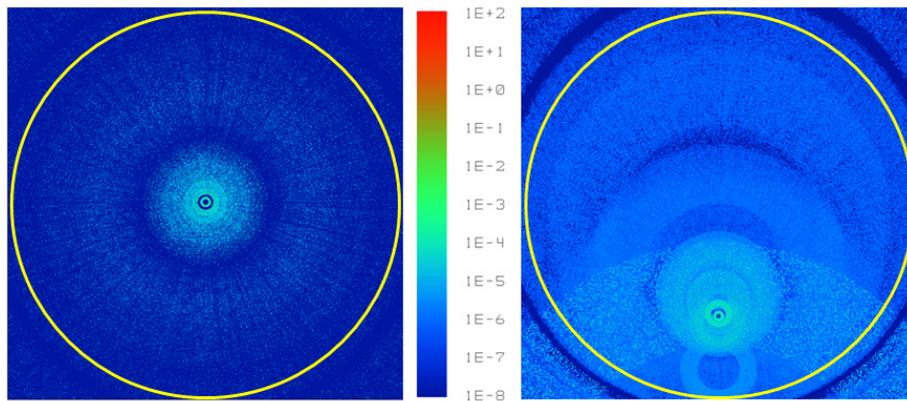


Figure 3.10: Calculated two-surface ghost images on-axis (left panel) and 1° off axis (right panel). The LSST's full field of view is represented by the yellow circle. Note that this does not yet take into account the reduction of detected surface brightness for the designed anti-reflection coating performance, and thus somewhat overestimates the effect of ghosts.

CAD models of the dome, telescope, and camera has been done to quantify the diffuse scattering contribution to the overall natural sky background. The initial analysis (Figure 3.11) computes the Point Source Transmittance (PST) for a single point source at various angles with respect to the telescope's optical axis. The PST is the integrated flux over the entire focal plane from the point source including the desired optical path and all first- and second-order scattered light. Each surface is specified with properties anticipated for the realized design, including contamination on the optical surfaces, micro-surface roughness, paint on non-optical surfaces, and so on.

The PST analysis shown in Figure 3.11 indicates that the LSST has excellent rejection of diffuse scattered light from out-of-field objects, with the PST dropping nearly three orders of magnitude beyond the imaging field of view (Ellis et al. 2009). Spreading this over the field of view of the LSST, the surface brightness contribution of a point source from diffuse scattering is at least 11 orders of magnitude below that of the direct image of the source.

3.5 The Expected Accuracy of Photometric Measurements

Željko Ivezić

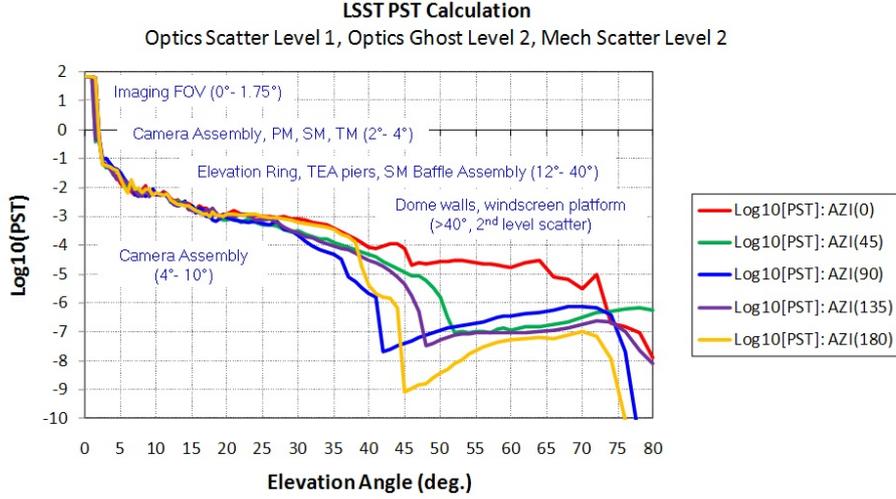


Figure 3.11: The LSST Point Source Transmittance (PST) as a function of source angle along five selected azimuth scans. The primary sources of scattering are identified along the scans.

The expected photometric error for a point source in magnitudes (roughly the inverse of signal-to-noise ratio) for a single visit (consisting of two back-to-back 15-second exposures) can be written as

$$\sigma_1^2 = \sigma_{sys}^2 + \sigma_{rand}^2, \quad (3.1)$$

where σ_{rand} is the random photometric error and σ_{sys} is the systematic photometric error (which includes errors due to, for example, imperfect modeling of the point spread function, but does not include errors in the absolute photometric zeropoint). For more details, see Section 3.3.1 in the LSST overview paper, [Ivezić et al. \(2008b\)](#). For N stacked observations, we assume $\sigma_{rand}(N) = \sigma_{rand}/\sqrt{N}$. This theoretically expected behavior has been demonstrated for repeated SDSS scans for N as large as 50 ([Ivezić et al. 2007](#); [Sesar et al. 2007](#); [Bramich et al. 2008](#)). The LSST calibration system and procedures are designed to maintain $\sigma_{sys} < 0.005$ mag and this is the value we adopt for a single LSST visit. Some effects that contribute to σ_{sys} will be uncorrelated between observations (e.g., errors due to imperfect modeling of the point spread function) and their impact will decrease with the number of stacked observations similarly to random photometric errors. For the final irreducible errors in LSST stacked photometry, we adopt $\sigma_{sys}=0.003$ mag (which will be probably dominated by errors in the transfer of the photometric zeropoints across the sky). LSST’s photometry will be limited by sky noise, and the random photometric error as a function of magnitude (per visit) can be described by

$$\sigma_{rand}^2 = (0.04 - \gamma)x + \gamma x^2 \text{ (mag}^2\text{)}, \quad (3.2)$$

with $x = 10^{0.4(m-m_5)}$. Here m_5 is the 5σ depth (for point sources) in a given band, and γ depends on the sky brightness, readout noise, and other factors. Using the LSST exposure time calculator (§ 3.2), we have obtained the values of γ listed in [Table 3.2](#). The 5σ depth for point sources is determined from

$$m_5 = C_m + 0.50(m_{sky} - 21) + 2.5 \log_{10} \frac{0.7}{\theta} + 1.25 \log_{10} \frac{t_{vis}}{30} - k_m(X - 1) \quad (3.3)$$

Table 3.2: The Parameters from Equation 3.2 and Equation 3.3

	u	g	r	i	z	y
m_{sky}^a	21.8	22.0	21.3	20.0	19.1	17.5
θ^b	0.77	0.73	0.70	0.67	0.65	0.63
γ^c	0.037	0.038	0.039	0.039	0.040	0.040
C_m^d	23.60	24.57	24.57	24.47	24.19	23.74
k_m^e	0.48	0.21	0.10	0.07	0.06	0.06
m_5^f	23.9	25.0	24.7	24.0	23.3	22.1
Δm_5^g	0.21	0.16	0.14	0.13	0.13	0.13

^a The expected median zenith sky brightness at Cerro Pachón, assuming mean solar cycle and three-day old Moon (mag/arcsec²).

^b The expected delivered median zenith seeing (arcsec). For larger airmass, X , seeing is proportional to $X^{0.6}$.

^c The band-dependent parameter from Equation 3.2.

^d The band-dependent parameter from Equation 3.3.

^e Adopted atmospheric extinction.

^f The typical 5σ depth for point sources at zenith, assuming exposure time of 2×15 sec, and observing conditions as listed. For larger airmass the 5σ depth is brighter; see the bottom row.

^g The loss of depth at the median airmass of $X = 1.2$ due to seeing degradation and increased atmospheric extinction.

where m_{sky} is the sky brightness (mag/arcsec²), θ is the seeing (FWHM, in arcsec), t_{vis} is the exposure time (seconds), k is the atmospheric extinction coefficient, and X is airmass. The constants, C_m , depend on the overall throughput of the instrument and are determined using the LSST exposure time calculator. The assumptions built into the calculator were tested using SDSS observations and by comparing the predicted depths to the published performance of the Subaru telescope (Kashikawa et al. 2003). The adopted values for C_m and k are listed in Table 3.2, as well as the expected m_5 in nominal observing conditions. See also Table 3.3 for the expected photometric accuracy at higher S/N.

3.6 Accuracy of Trigonometric Parallax and Proper Motion Measurements

Željko Ivezić, David Monet

Given the observing sequence for each sky position in the main survey provided by the LSST Operations Simulator (§ 3.1), we generate a time sequence of mock astrometric measurements. The assumed astrometric accuracy is a function of S/N . Random astrometric errors per visit are modeled as $\theta/(S/N)$, with $\theta = 700$ mas and S/N is determined using expected LSST 5σ depths for point sources. The estimated proper motion and parallax accuracy at the bright end ($r < 20$) is driven by systematic errors due to the atmosphere. Systematic errors of 10 mas are added in

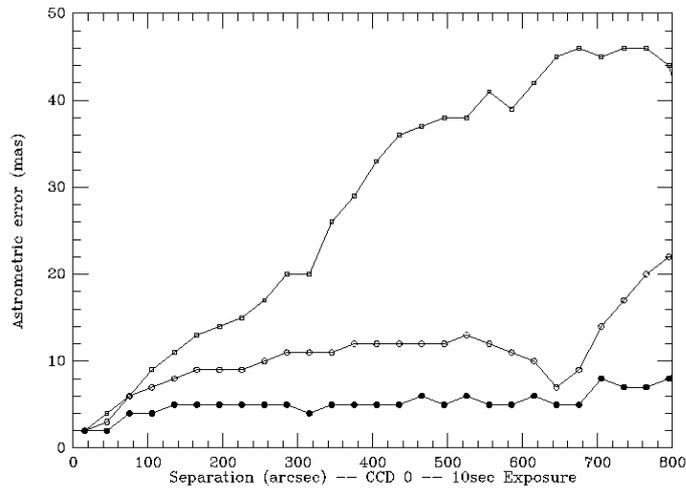


Figure 3.12: Differential astrometric error as a function of angular separation derived from a sequence of 10-second Subaru Suprime-Cam observations. The upper curve is computed from transformation using only offsets between frames. The middle curve includes linear transformation coefficients and the bottom curve includes cubic transformation coefficients. The improvement in astrometric accuracy suggests that low-order polynomials are a reasonable model for the geometric impact of atmospheric turbulence over spatial scales of several arcminutes. From Saha & Monet (2005), with permission.

quadrature, and are assumed to be uncorrelated between different observations of a given object. Systematic and random errors become similar at about $r = 22$, and there are about 100 stars per LSST sensor (0.05 deg^2) to this depth (and fainter than the LSST saturation limit at $r \sim 16$) even at the Galactic poles.

Precursor data from the Subaru telescope (Figure 3.12) indicate that systematic errors of 10 mas on spatial scales of several arc-minutes are realistic. Even a drift-scanning survey such as SDSS delivers uncorrelated systematic errors (dominated by seeing effects) at the level of 20-30 mas rms per coordinate (measured from repeated scans, Pier et al. 2003), and the expected image quality for LSST will be twice as good as for SDSS. Furthermore, there are close to 1000 galaxies per sensor with $r < 22$, which will provide exquisite control of systematic astrometric errors as a function of magnitude, color, and other parameters, and thus enable absolute proper motion measurements.

The astrometric transformations for a given CCD and exposure, and proper motion and parallax for all the stars from a given CCD, are simultaneously solved for using an iterative algorithm (§ 2.7). The astrometric transformations from pixel to sky coordinates are modeled using low-order polynomials and standard techniques developed at the U.S. Naval Observatory (Monet et al. 2003). The expected proper motion and parallax errors for a ten-year long baseline survey, as a function of apparent magnitude, are summarized in Table 3.3. Roughly speaking, trigonometric parallax errors can be obtained by multiplying the astrometric errors by 0.039, and proper motion errors (per coordinate) can be obtained by multiplying the single-visit astrometric errors by 0.014 yr^{-1} .

Blue stars (e.g., F and G stars) fainter than $r \sim 23$ will have about 50% larger proper motion and

Table 3.3: The expected proper motion, parallax, and accuracy for a ten-year long baseline survey.

r mag	σ_{xy}^a mas	σ_{π}^b mas	σ_{μ}^c mas/yr	σ_1^d mag	σ_C^e mag
21	11	0.6	0.2	0.01	0.005
22	15	0.8	0.3	0.02	0.005
23	31	1.3	0.5	0.04	0.006
24	74	2.9	1.0	0.10	0.009

^a Typical astrometric accuracy (rms per coordinate per visit).

^b Parallax accuracy for 10-year long survey.

^c Proper motion accuracy for 10-year long survey.

^d Photometric error for a single visit (two 15-second exposures).

^e Photometric error for stacked observations (see Table 1).

parallax errors than given in the table due to decreased S/N in z and y . The impact on red stars is smaller due to the relatively small number of observations in the u and g bands, but extremely red objects, such as L and T dwarfs, will definitely have larger errors, depending on details of their spectral energy distribution. After the first three years of the survey, the proper motion errors are about five times as large, and parallax errors will be about twice as large as the values given in Table 3.3; the errors scale as $t^{-3/2}$ and $t^{-1/2}$ respectively.

For comparison with Table 3.3, the SDSS-POSS proper motion measurements have an accuracy of ~ 5 mas/yr per coordinate at $r = 20$ (Munn et al. 2004). Gaia is expected to deliver parallax errors of 0.3 mas and proper motion errors of 0.2 mas/yr at its faint end at $r \sim 20$. Hence, LSST will smoothly extend Gaia’s error vs. magnitude curve four magnitudes fainter, as discussed in detail in § 6.12.

3.7 Expected Source Counts and Luminosity and Redshift Distributions

Željko Ivezić, A. J. Connolly, Mario Jurić, Jeffrey A. Newman, Anthony Tyson, Jake VanderPlas, David Wittman

The final stacked image of LSST will include about ten billion galaxies and ten billion stars, mostly on the main sequence. The data sources and assumptions used to derive these estimates are described here. Of course, LSST will also detect very large samples of many other types of objects such as asteroids, white dwarfs, and quasars (roughly ten million in each category). We defer discussion of those more specific topics to the relevant science chapters that follow.

3.7.1 Stellar Counts

In order to accurately predict stellar source counts for the LSST survey, both a Galactic structure model and a detailed estimate of the stellar luminosity function are required. SDSS data can be used to guide these choices. [Figure 3.13](#) shows the stellar counts, as a function of distance and color, for stars observed with SDSS towards the North Galactic Pole. Stars are selected to have colors consistent with main sequence stars following criteria from [Ivezić et al. \(2008a, hereafter I08\)](#). This color selection is sufficiently complete to represent true stellar counts, and sufficiently efficient that contamination by giants, white dwarfs, quasars, and other non-main sequence objects is negligible. Distances are computed using the photometric parallax relation and its dependence on metallicity derived by I08. The displayed density variation in the horizontal direction represents the luminosity function, and the variation in the vertical direction reflects the spatial volume density profiles of disk and halo stars. Both effects need to be taken into account in order to produce reliable counts for the LSST survey.

To extrapolate stellar counts from the SDSS faint limit at $r = 22.5$ to the faint limit of the stacked LSST map ($r = 27.5$), we use the Milky Way model by [Jurić et al. \(2008, hereafter, J08\)](#). This model reproduces the SDSS count data to within 10% (except in regions with significant substructure) as shown in [Figure 3.13](#), as well as the count variation as a function of position on the sky. Using photometric data for 50 million stars from SDSS Data Release 4, sampled over a distance range from 100 pc to 15 kpc, J08 showed that the stellar number density distribution, $\rho(R, Z, \phi)$ can be well described (apart from local overdensities; the J08 best-fit was refined using residual minimization algorithms) as a sum of two cylindrically symmetric components,

$$\rho(R, Z, \phi) = \rho_D(R, Z) + \rho_H(R, Z). \quad (3.4)$$

The disk component can be modeled as a sum of two exponential disks

$$\rho_D(R, Z) = \rho_D(R_\odot) \times \left[e^{-|Z+Z_\odot|/H_1 - (R-R_\odot)/L_1} + \epsilon_D e^{-|Z+Z_\odot|/H_2 - (R-R_\odot)/L_2} \right], \quad (3.5)$$

and the halo component requires an oblate power-law model

$$\rho_H(R, Z) = \rho_D(R_\odot) \epsilon_H \left(\frac{R_\odot^2}{R^2 + (Z/q_H)^2} \right)^{n_H/2}. \quad (3.6)$$

The best-fit parameters are discussed in detail by J08. For LSST simulations, we have adopted parameters listed in the second column of their Table 10.

This Galaxy model gives star counts accurate only to about a factor of two, due to our incomplete knowledge of the three-dimensional dust distribution in the Galactic plane, and the uncertain location of the edge of the stellar halo. As illustrated in [Figure 3.13](#), if this edge is at 100 kpc or closer to the Galactic center, it will be detected as a sudden drop in counts of blue faint stars beyond some color-dependent flux limit. For example, blue turn-off stars with $M_r < 5$ should display a sharp decrease in their differential counts for $r > 25$, if there is a well-defined end to the distribution of halo stars at 100 kpc. We obtain approximate estimates by extrapolating counts for $r < 21$ from USNO-B all-sky catalog to fainter magnitudes using models described above. There

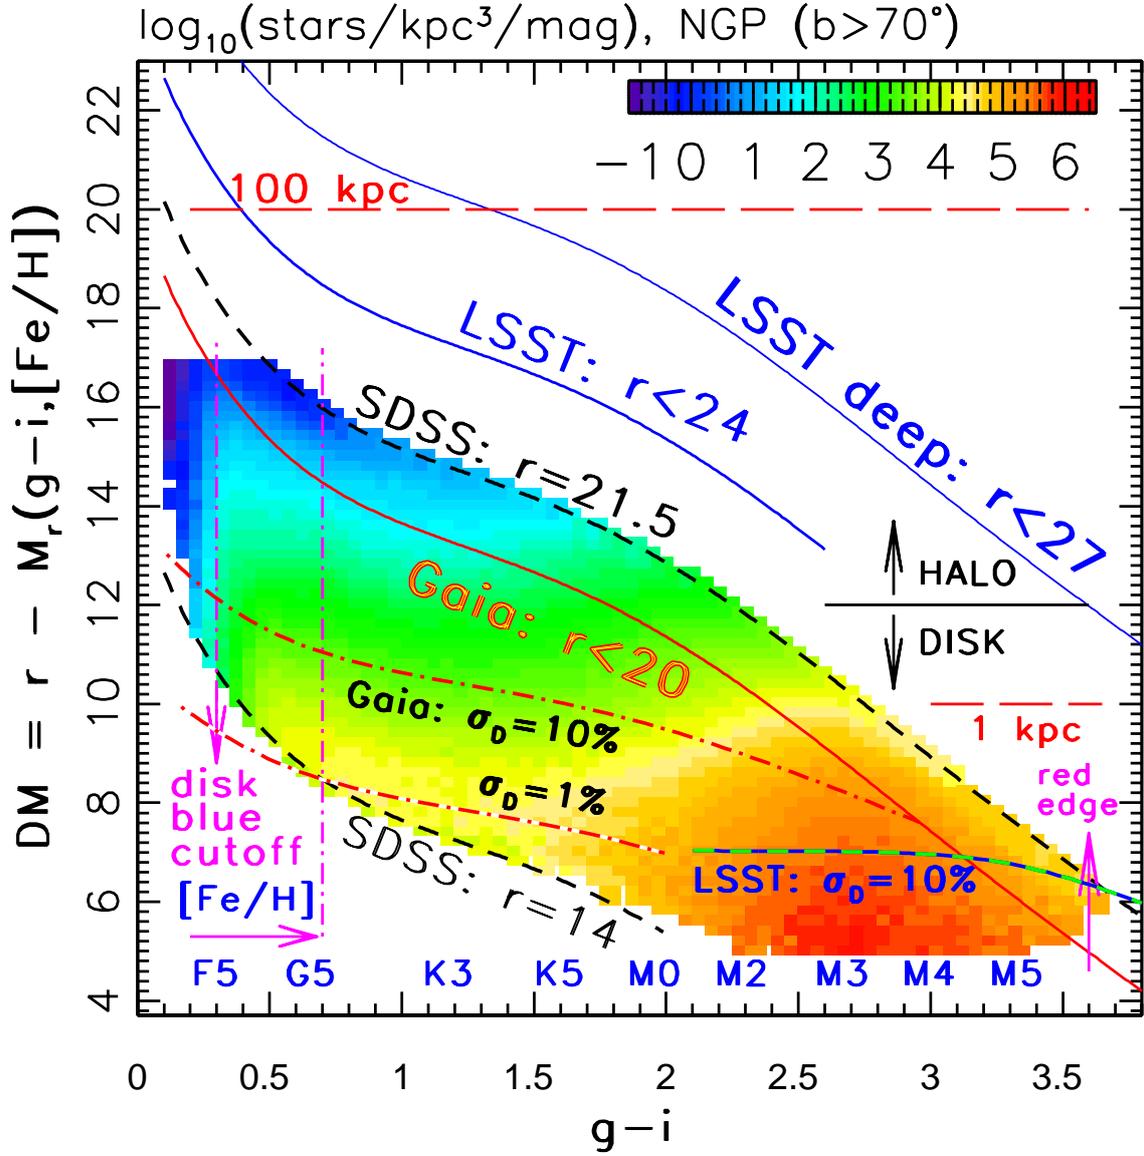


Figure 3.13: The volume number density (stars/kpc³/mag, log scale according to legend) of ~ 2.8 million SDSS stars with $14 < r < 21.5$ and $b > 70^\circ$, as a function of their distance modulus (distances range from 100 pc to 25 kpc) and their $g-i$ color. The sample is dominated by color-selected main sequence stars. The absolute magnitudes are determined using the photometric parallax relation from I08. The metallicity correction is applied using photometric metallicity for stars with $g-i < 0.7$, and by assuming $[Fe/H] = -0.6$ for redder stars. The relationship between the MK spectral type and $g-i$ color from Covey et al. (2007) is indicated above the $g-i$ axis; $g-i = 0.7$ roughly corresponds to G5. The two vertical arrows mark the turn-off color for disk stars and the red edge of the M dwarf color distribution. The $[Fe/H]$ label shows the color range ($g-i < 0.7$) where the photometric metallicity estimator from I08 performs best. The two diagonal dashed lines, marked $r = 14$ and $r = 21.5$, show the apparent magnitude limits for SDSS data. At a distance of $\sim 2-3$ kpc ($DM = 12$), halo stars begin to dominate the counts. The diagonal solid lines mark the apparent magnitude limits for Gaia ($r < 20$), LSST’s single epoch data ($r < 24$, 10σ), and LSST’s stacked data ($r < 27$, 10σ). The dashed line in the lower right corner marks the distance limits for obtaining 10% accurate trigonometric distances using LSST data. The two dot-dashed lines mark analogous limits for obtaining 1% and 10% accurate trigonometric distances using Gaia’s data (§ 6.12).

Table 3.4: Stellar counts based on USNO-B and model-based extrapolations

	$N(r < 24.5)^a$	ratio(24.5/21)	ratio(27.8/24.5)	ratio(27.8/21)
Galactic center	172	6.4	3.8	24
anticenter	120	4.5	2.4	11
South Galactic Pole	4	2.6	2.0	5

^a The number of stars with $r < 24.5$ in thousands per deg². The entries are computed using counts based on the USNO-B catalog and extrapolated from its $r = 21$ limit using model-based count ratios, listed in the second column. LSST will detect ~ 4 billion stars with $r < 24.5$ and 10 billion stars with $r < 27.8$.

are 10^9 stars with $r < 21$ in the USNO-B catalog, and this count is probably accurate to better than 20%, which is a smaller uncertainty than extrapolations described below.

The ratio of stellar counts to $r < 24.5$ and $r < 27.8$ (LSST’s single visit and stacked depths) to those with $r < 21$ varies significantly across the sky due to Galactic structure effects and the interstellar dust distribution. For the dust distribution, we assume an exponential dependence in radial and vertical directions with a scale height of 100 pc and a scale length of 5 kpc. We assume a dust opacity of 1 mag/kpc (in the r band) which produces extinction of 0.1 mag towards the North Galactic pole, 20 mag towards the Galactic center, and 5 mag towards the anticenter, in agreement with “common wisdom.” Using the stellar counts model described above, and this dust model, we evaluate the counts’ ratios as a function of location on the sky and integrate over the sky to be covered by LSST’s main survey. In the regions observed by SDSS, the predicted counts agree to better than 20% (the models were tuned to SDSS observations, but note that the normalization comes from USNO-B). The counts’ ratios for several special directions are listed in [Table 3.7.1](#). The predicted total number of stars is 4 billion for $r < 24.5$ with an uncertainty of $\sim 50\%$, and 10 billion for $r < 27.8$, with an uncertainty of at most a factor of 2.

3.7.2 Galaxy Counts

Model-independent, empirical estimates of galaxy counts with LSST can be gleaned from a number of deep multicolor photometric surveys that have been performed over the last decade. These are sufficient to predict the counts for the LSST galaxies (e.g., [Ilbert et al. 2006b](#)) with an uncertainty of about 20% (most of this uncertainty comes from photometric systematics and large-scale structure). Based on the CFHTLS Deep survey ([Hoekstra et al. 2006](#); [Gwyn 2008](#)), the cumulative galaxy counts for $20.5 < i < 25.5$ are well described by

$$N_{gal} = 46 \times 10^{0.31(i-25)} \text{ galaxies arcmin}^{-2}. \tag{3.7}$$

The so-called “gold” sample of LSST galaxies with a high S/N defined by $i < 25.3$ (corresponding to $S/N > 20$ for point sources assuming median observing conditions), will include four billion galaxies (or 55 arcmin^{-2}) over $20,000 \text{ deg}^2$ of sky (see [Figure 3.14](#)). The effective surface density of galaxies useful for weak lensing analysis in the “gold” sample will be about 40 arcmin^{-2} with an uncertainty of 20%. The total number of galaxies detected above the faint limit of the stacked

map ($r < 27.5$, corresponding to $i \sim 26$ given the typical colors of galaxies) will be close to 10 billion over 20,000 deg².

The redshift and rest-frame color distributions of these sources are much less well understood due to the lack of any complete spectroscopic sample to the depth of the LSST. To estimate the redshift distributions for the LSST we, therefore, use both simple extrapolation of observations and more sophisticated simulations that have been designed to match available observational data sets. In [Figure 3.15](#), we show a prediction for the redshift distribution of galaxies of the form

$$p(z) = \frac{1}{2z_0} \left(\frac{z}{z_0} \right)^2 \exp(-z/z_0). \quad (3.8)$$

This functional form fits DEEP2 data well (after completeness corrections) for $i < 23$. We estimate z_0 for the $i < 25$ sample by extrapolating the tight linear relationship between z_0 and limiting i magnitude in the DEEP2 survey, $z_0 = 0.0417i - 0.744$ (measured for $21.5 < i < 23$). The mean redshift of a sample is $3z_0$ and the median redshift is $2.67z_0$; for the $i < 25$ sample, the mean redshift is 0.9 and the median is 0.8.

This prediction is compared to a model based on an empirical evolving luminosity function, and to the simulations of [Kitzbichler & White \(2007, hereafter KW07\)](#). Based on the Millennium simulations ([Springel et al. 2005](#)), the baryonic physics in KW07 models includes gas cooling, heating from star formation, supernovae, and radio mode feedback. Comparisons with existing imaging surveys has shown that the model for the dust used in KW07 provides a good match to the color-luminosity relation seen in deep surveys to $z \sim 1.4$ (although the simulations predict more than the K -band number counts).

3.8 Photometric Redshifts

A. J. Connolly, Jeffrey A. Newman, Samuel Schmidt, Alex Szalay, Anthony Tyson

The estimation of galaxy redshifts from broad band photometry, i.e., photometric redshifts ([Baum 1962; Koo 1985; Loh & Spillar 1986; Connolly et al. 1995](#)), has become a widely used tool in observational cosmology ([Collister & Lahav 2004; Wadadekar 2005; Carliles et al. 2008; Gwyn & Hartwick 1996; Lanzetta et al. 1996; Sawicki et al. 1997; Budavári et al. 2000; Ilbert et al. 2006a](#)). These probabilistic redshift (and galaxy property) estimates are derived from characteristic changes in the observed colors of galaxies due to the redshifting of features in galaxy spectral energy distributions through a series of broad band filters. At optical and ultraviolet wavelengths, the Lyman and Balmer breaks (at 1000Å and 4000Å respectively) are the primary source of the redshift information. To first order, the accuracy to which we can determine the position of these breaks from the observed colors determines the scatter within the photometric redshift relation, and our ability to correctly distinguish between the breaks determines the amount of systematic error (or catastrophic outliers) in the relation.

The LSST reference filter system, covering the u, g, r, i, z , and y passbands, provides leverage for redshift estimation from $z = 0$ to $z > 6$ (although, as we will describe later, the redshift interval, $1.4 < z < 2.5$, will be less well constrained as the Balmer break has transitioned out of the y band and the Lyman break has yet to enter the u band). We describe here the expected photometric

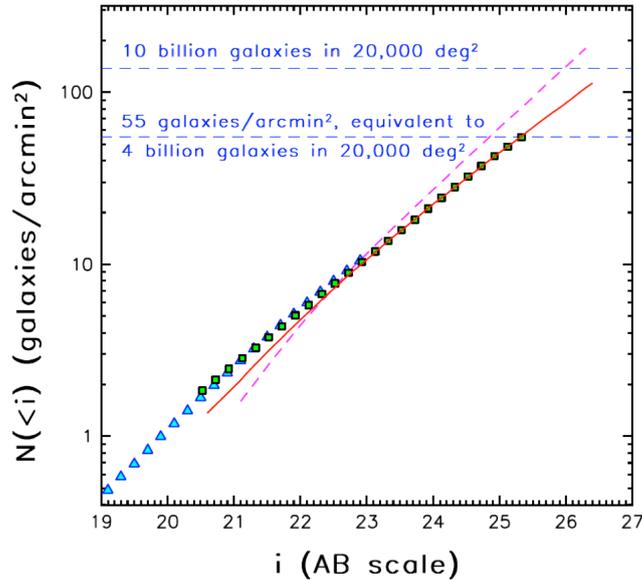


Figure 3.14: Cumulative galaxy counts in the SDSS i band. The triangles show SDSS counts from the so-called “stripe 82” region (Abazajian et al. 2009) and the squares show counts from the CFHTLS Deep survey (Gwyn 2008). The dashed diagonal line is based on the Millennium Simulation (Springel et al. 2005) and the solid line is a simulation based on a model with evolving luminosity function from the DEEP2 and VVDS surveys (measured at redshifts up to unity) and non-evolving SEDs. The two dashed horizontal lines are added to guide the eye. LSST will detect 4 billion galaxies with $i < 25.3$, which corresponds to an S/N of at least 20 for point sources in median observing conditions. The full LSST sample may include as many as 10 billion galaxies.

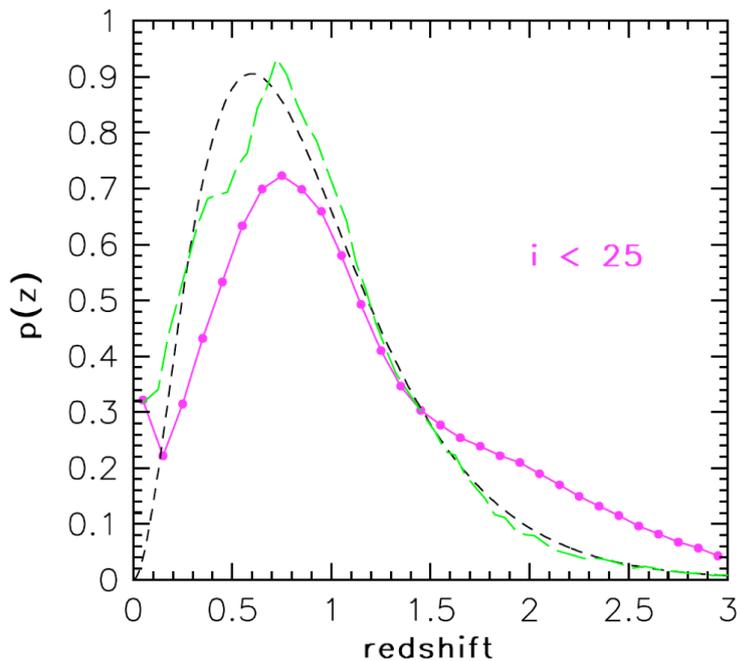


Figure 3.15: The redshift probability distributions for faint galaxies. The dashed curve shows a best fit to the measured and debiased DEEP2 redshift distribution for galaxies with $i < 23$, and extrapolated to $i < 25$ (see text). The other two curves show model predictions for galaxies with $i < 25$ (magenta: the Millennium Simulation; green: an evolving luminosity function; see Figure 3.14).

redshift performance for LSST based on empirical studies and simulations (including scatter, bias, and the fraction of sources that are catastrophic outliers) and describe ongoing work to characterize and minimize these uncertainties.

3.8.1 Photometric Redshifts for the LSST

Photometric redshifts for LSST will be applied and calibrated over the redshift range $0 < z < 4$ for galaxies to $r \sim 27.5$. For the majority of science cases, such as weak lensing and BAO, a subset of galaxies with $i < 25.3$ will be used. For this high S/N gold standard subset (§ 3.7.2) over the redshift interval, $0 < z < 3$, the photometric redshift requirements are:

- The root-mean-square scatter in photometric redshifts, $\sigma_z/(1+z)$, must be smaller than 0.05, with a goal of 0.02.
- The fraction of 3σ outliers at all redshifts must be below 10%.
- The bias in $e_z = (z_{photo} - z_{spec})/(1+z_{spec})$ must be below 0.003 (§ 14.5.1, or 0.01 for combined analyses of weak lensing and baryon acoustic oscillations); the uncertainty in $\sigma_z/(1+z)$ must also be known to similar accuracy.

Figure 3.16 and Figure 3.17 show the expected performance for the LSST gold sample on completion of the survey. These results are derived from simulated photometry that reproduces the distribution of galaxy colors, luminosities, and colors as a function of redshift as observed by the COSMOS (Lilly et al. 2009), DEEP2 (Newman et al. 2010, in preparation), and VVDS (Garilli et al. 2008) surveys. The simulations include the effects of evolution in the stellar populations, redshift, and type dependent luminosity functions, type dependent reddening, and of course photometric errors. The photometric redshifts are determined using a likelihood technique as outlined below.

Figure 3.17 shows the residuals, fraction of outliers, dispersion, and bias associated with the photometric redshifts as a function of i band magnitude and redshift. For this case, magnitude and surface brightness priors have been applied to the data and all sources with broad or multiply peaked redshift probability functions have been excluded (see §3.8.3). For the brighter sample, ($i < 24$), the photometric redshifts meet or exceed our performance goals for all except the highest redshift bin. For the gold sample, the photometric redshifts meet the science requirements on dispersion and bias at all redshifts. At redshifts $z > 2$, the fraction of outliers is a factor of two larger than the goal for LSST. These outliers reduce the size of the samples with usable photometric redshifts by approximately 10%. Other cuts and priors will reduce the outlier fraction further. This demonstrates that highly accurate photometric redshifts should be attainable with LSST photometry, assuming perfect knowledge of SED templates (or equivalently the span of galaxy properties). For selected subsets of objects (e.g., bright red sequence galaxies), we may be able to do much better attaining $\sigma_z/(1+z)$ of 0.01 or less.

3.8.2 Dependence on Filter Specification and Signal-to-Noise Ratio

The accuracy of LSST photometric redshift depends on both the characteristics of the filter system and our ability to photometrically calibrate the data. For the LSST reference filters the scatter in the photometric redshifts in simulated data scales approximately linearly with S/N, with a

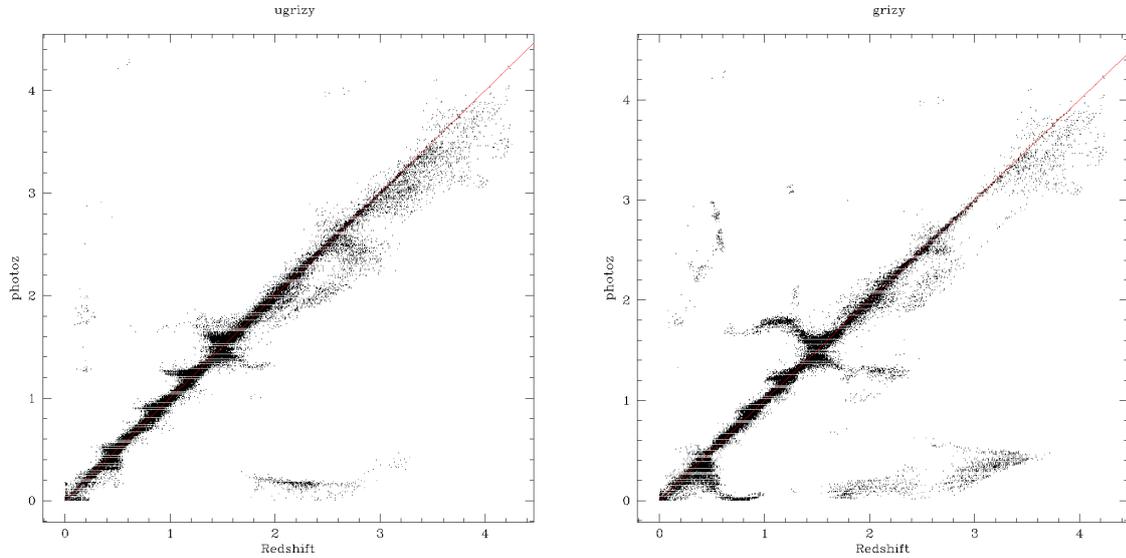


Figure 3.16: Impact of using the u filter to improve measurement and resolve degeneracies in photometrically determined redshifts. On the left is the correlation between the photometric redshifts and spectroscopic redshifts with the full complement of LSST filters. The right panel shows the photometric redshift relation for data excluding the u filter. The addition of u data reduces the scatter substantially for $z < 0.5$ and removes degeneracies over the full redshift range.

floor of $\sigma_z \sim 0.02$ (when including all galaxy types). This is consistent with the σ_z obtained for photometric redshifts obtained for $r < 17.77$ galaxies in SDSS (e.g., Ball et al. 2008; Freeman et al. 2009), several magnitudes brighter than the depths of the photometry (and hence in some ways analogous to the gold sample).

However, it is significantly better than has been achieved to date for photometric redshift algorithms for fainter samples and to higher redshifts, likely due to the fact that such samples must handle a broader range of galaxy types over a broader redshift range, and are increasingly dominated by strongly star forming galaxies (which possess only weak 4000\AA breaks) as they extend fainter and to higher redshifts. With CFHT Legacy Survey deep *ugriz* imaging, Ilbert et al. (2006a) achieve $\sigma_z \sim 0.03$ for $i < 21.5$, degrading to $0.04 - 0.05$ for $22.5 < i < 24$; while with deep 16-band photometry, and restricting to a subset of galaxies with $z < 1.2$ and $K < 21.6$ (AB), Mobasher et al. (2007) attain $\sigma_z \sim 0.03$ for a sample with $20 < i < 24$. Unfortunately, these numbers are difficult to compare due to the larger number of bands and the K band limit (which will favor massive, lower star formation rate galaxies at higher redshifts) used by Mobasher et al. (2007). The fundamental limitation which puts a floor on σ in these empirical tests is unclear (likely depending on poorly known template spectra, errors in photometric measurements due to blended galaxies, and variations in the emission line properties of galaxies with redshift and type). The number of catastrophic failures also depends on S/N, but the exact scaling remains unclear (Mobasher et al. 2007); in Ilbert et al. (2006a), the catastrophic failure rate is $< 1\%$ for $i < 21.5$, $\sim 2\%$ for $21.5 < i < 22.5$, $\sim 4\%$ for $22.5 < i < 23.5$, and $\sim 9\%$ for $23.5 < i < 24$. Regardless, based upon the SDSS experience, we can expect that with greater zero point uniformity, better bandpass characterization, and improved calibration LSST should yield significantly better photometric redshift results than previous optical broad-band surveys.

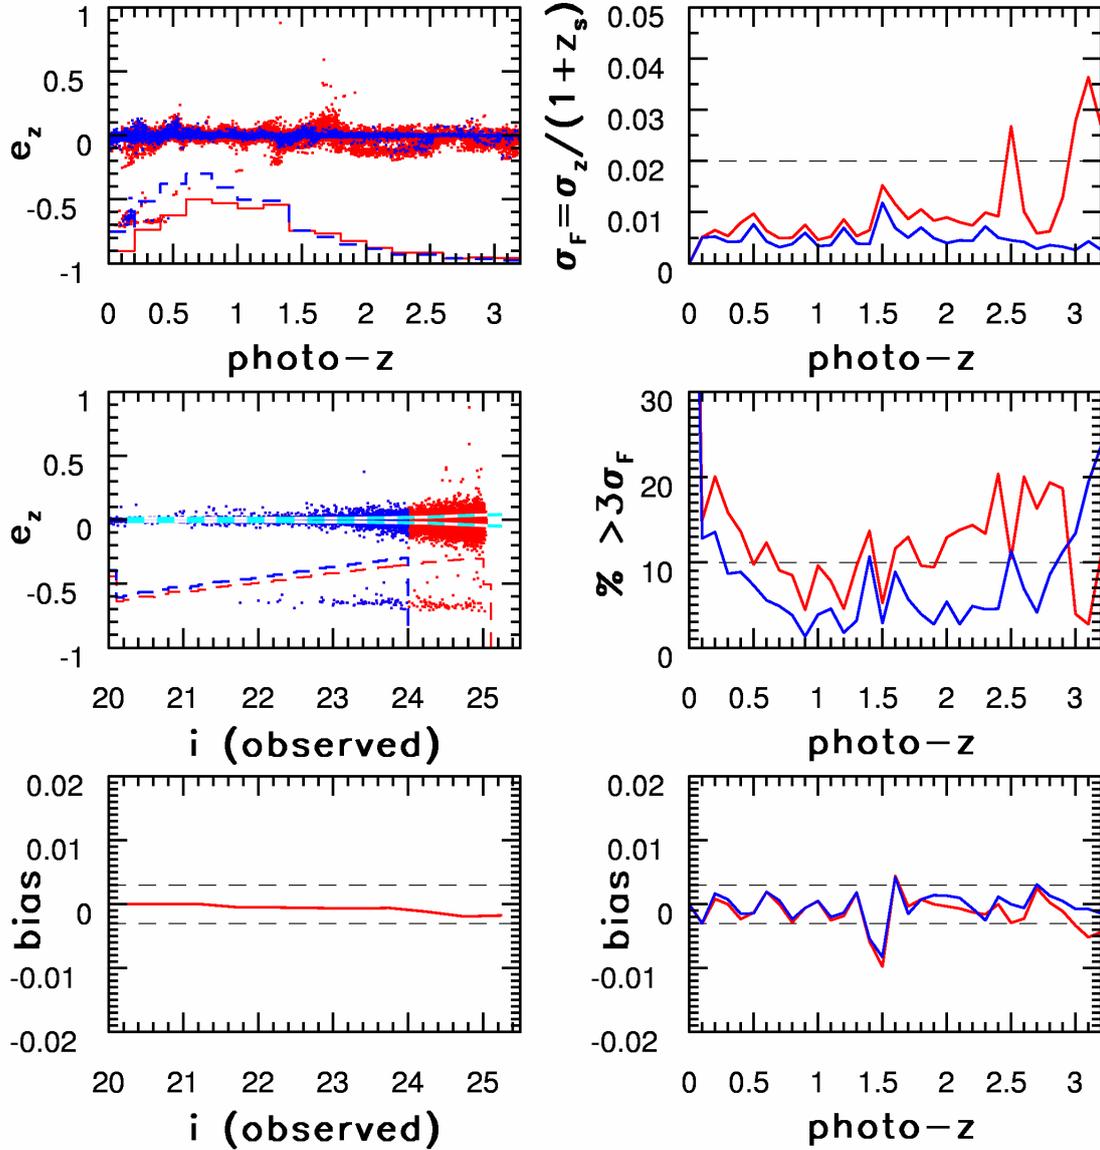


Figure 3.17: Illustration of the photometric redshift performance as a function of apparent magnitude and redshift, for a simulation based on the LSST filter set (*ugrizy*). Red points and curves correspond to the gold sample defined by $i < 25$, and blue points and curves to a subsample with $i < 24$. The photometric redshift error is defined as $e_z = (z_{photo} - z_{spec}) / (1 + z_{spec})$. *Top left*: e_z vs. photometric redshift. The two histograms show redshift distributions of the simulated galaxies. *Top right*: the root-mean-square scatter (rms, determined from the interquartile range) of e_z as a function of photometric redshift. The horizontal dashed line shows the science driven design goal. *Middle left*: e_z vs. observed i band magnitude. Two histograms show the logarithmic differential counts (arbitrary normalization) of simulated galaxies. The two horizontal cyan lines show the 3σ envelope around the median e_z (where σ is the rms from the top right panel). *Middle right*: the fraction of 3σ outliers as a function of redshift. The horizontal dashed line shows the design goal. *Bottom left*: the median value of e_z (bias) as a function of apparent magnitude. The two dashed lines show the design goal for limits on bias. *Bottom right*: the median value of e_z (i.e., the bias in estimated redshifts) as a function of redshift. The two dashed lines show the design goal for this quantity.

In [Figure 3.16](#) we show the impact of the u band filter for redshift estimation assuming a survey to the nominal depth of the LSST and including magnitude and surface brightness priors. At low redshift the redshifting of the Balmer break through the u filter enables the estimation of the photometric redshifts for $z < 0.5$ (where the break moves into the g and r bands). At higher redshift, the transition of the Lyman break into the u band filter increases the accuracy of the photometric redshifts for $z > 2.5$. The result of this is two-fold; the scatter in the redshift estimation is decreased at low redshift, improving studies of the properties of galaxies in the local Universe and reducing the number of catastrophic outliers (mistaking the Lyman break for the Balmer break results in a degeneracy between $z = 0.2$ and $z = 3$ galaxies) by a factor of two. Removal of the u band results in a deterioration of the photometric redshifts for $z < 0.6$ to such an extent that they fail to meet the required performance metrics as described above.

At redshifts $1.3 < z < 1.6$, the photometric redshift constraints are most dependent upon the y filter. For $z > 1.6$ the Balmer break transitions out of the y band and hence the photometric redshifts are only poorly constrained until the Lyman break enters the u band at $z > 2.5$.

Addition of near-infrared passbands from, for example, a space-based imager yielding $S/N=10$ photometry in both the J and H bands at an AB magnitude of 25 results in a reduction in σ_z , the fraction of outliers, and the bias by approximately a factor of two for $z > 1.5$. At redshifts $z < 1.5$ there is no significant improvement in photometric redshift performance from near-infrared data, in contrast to the u band data which impacts photometric redshifts below $z = 1$ (even when the J and H bands are included already).

3.8.3 Priors in Redshift Estimation

In order to mitigate catastrophic failures in photometric redshifts, Bayesian approaches for redshift estimation have been developed ([Benítez 2000](#)). In this case we search for the two-dimensional posterior distribution $P(z, T | \mathbf{C}, \mathbf{O})$, where z is the redshift of the galaxy, T is the “template” or galaxy type, \mathbf{C} is the vector of fluxes from the data, and \mathbf{O} is a vector of galaxy observables independent of the fluxes, such as size, brightness, morphology, or environment. If we make the assumption that \mathbf{O} and \mathbf{C} are independent then,

$$P(z, T | \mathbf{C}, \mathbf{O}) = \frac{P(\mathbf{C} | z, T) P(z, T | \mathbf{O})}{P(\mathbf{C})}. \quad (3.9)$$

The posterior distribution $P(z, T | \mathbf{C}, \mathbf{O})$ is given in terms of the likelihood function $P(\mathbf{C} | z, T)$ and the prior distribution $P(z, T | \mathbf{O})$; the prior encompasses all knowledge about galaxy morphology, evolution, environment, brightness, or other quantities.

The most common prior used in photometric redshifts has been magnitude ([Benítez 2000](#)); e.g. a prediction for the overall redshift distribution of galaxies given an apparent magnitude and type. Other priors that have been considered include morphological type and surface brightness ([Stabenau et al. 2008](#)). Surface brightness for a given galaxy scales as $(1 + z)^4$, suggesting it should be a powerful constraint, but it evolves strongly with redshift, depends on spectra type, and depends on accurate measurements of the size of galaxies close to the seeing size, making it less useful. However, [Stabenau et al. \(2008\)](#) show that if a ground-based survey can precisely measure the angular area of galaxies, achieve a seeing of $1''$ or less, and attain a surface brightness

sensitivity below $26.5 \text{ mag/arcsec}^2$, then it should be able to do almost as well as one from space in constraining photometric redshifts via surface brightness (resulting in a decrease in the photometric redshift bias of up to a factor of six).

In general even with the use of priors, photometric redshifts are typically taken to be simply the redshift corresponding to the maximum likelihood point of the redshift probability function. Template-based photometric redshift algorithms, however, can provide a full probability distribution over the entire redshift and spectral type range. Using the full redshift probability distribution function for each galaxy can significantly reduce the number of catastrophic outliers by, for example, excluding galaxies with broad or multiply peaked probability distributions.

By identifying and pre-filtering problematic regions in photometric redshift space, we can exclude the galaxies most likely to produce outliers while retaining the galaxies that have redshifts that are well-constrained. For most statistical studies, it is far more important to eliminate outliers than to maximize the total number of galaxies in the sample. Application of a simple photometric redshift space filter (for example, excluding galaxies classified as blue galaxies at $1.5 < z < 1.8$, which are particularly susceptible to catastrophic failure) gives an outlier fraction for an $i < 25.3$ sample a factor of two smaller than those we’ve described in this section. Other priors can further reduce outliers.

Both the specific photometric redshift technique used, the appropriate selection methods, priors, and their weights will be science case specific. For high redshift galaxies or for searches for unusual objects, heavily weighting priors based on galaxy properties may suppress those sources. For science cases requiring galaxies of specific types (e.g., baryon acoustic oscillation measurements, § 13.3) or for galaxies with particular observed attributes (e.g., resolved galaxies for weak lensing studies), methods for optimizing priors must be defined.

3.8.4 Photometric Redshift Calibration: Training Sets

Calibration of photometric zero points, SEDs, and priors will be critical for developing photometric redshifts for the LSST. If the range of spectral types is only coarsely sampled, the uncertainty in predicted redshift will increase, as the exact SED for an individual galaxy may not be present in the training set. For example, if we use only 50% of the model galaxy templates used to generate spectra when computing photometric redshifts with the methods used for Figures 3.16 and 3.17, the scatter (σ_z) increases by 40% and the bias by 50% overusing all of the templates. This outcome highlights the need for significant numbers of spectroscopic galaxies to train our template SEDs, and also illustrates the need for training sets to span the properties of galaxies in the samples to which we apply photometric redshifts.

It remains unclear both how small a subset of the complete data is sufficient to determine the overall redshift structure, and how we might select that subset. If the objects we seek reside only in certain regions of color space or have some specific properties, then simple sampling strategies can be used to pick an appropriate subset for spectroscopy (e.g., Bernstein & Huterer 2009). We often cannot, however, isolate a problematic population a priori. We could rely on the “standard” technique of either applying a sharp selection threshold in a particular attribute (e.g., galaxy size or magnitude) or picking a suitable random fraction of the underlying sample and then studying

this population in detail. Neither of these naïve techniques is optimal in any statistical sense; their only appeal is their apparent simplicity.

For certain well-defined parameter estimation problems there are classical stratification techniques (Neyman 1938) if we want the optimum variance estimator over a sample consisting of discrete “classes,” each with its own variance. These stratified sampling strategies lie between the limiting cases of totally randomly sampling from the full ensemble or randomly sampling each category. In astronomy, however, it is quite rare that a single estimator will suffice. More likely we seek a distribution of a derived quantity; that is, we seek the distribution of an intrinsic quantity but have only the observed quantity available (consider measurements of the luminosity function: we seek to determine the probability distribution of the true physical brightnesses of a population when only apparent magnitude can be measured).

Sampling strategies using Local Linear Embedding (LLE, Roweis & Saul 2000) can preserve the distribution of spectral types of galaxies in local spectroscopic surveys with $\sim 10,000$ galaxies (compared to an initial sample of 170,000 spectra). This is done by considering how much new information is added as we increase the number of galaxies within a sample (VanderPlas & Connolly 2009). This reduction in sample size required to encapsulate the full range of galaxy types is also consistent with the sample sizes used for Principal Component Analyses of SDSS spectra (Yip et al. 2004).

Based on this fact, one approach would be to generate a series of selected fields distributed across the sky with galaxies to $r \sim 26$ calibrated via selected deep spectroscopy. We would need sufficient numbers of galaxies per redshift bin to beat down the statistical errors to at least the level of the systematic errors. If we take the previously stated dark energy systematic targets as a goal ($\delta_z = 0.003(1+z)$, $\Delta\sigma_z = 0.004(1+z)$), then we need ~ 6000 galaxies per bin. In fact, given that we need to characterize the full distribution function, as it is non-Gaussian, it is more likely that we would need $\sim 100,000$ galaxies total if the sample were split up into ten redshift bins. The number needed can, however, be reduced by almost a factor of two by sampling the redshift distribution in an optimized manner (Ma & Bernstein 2008). This number is comparable to that needed for calibration of the templates and zero points. For the gold sample, $i < 25$, obtaining redshifts for 50,000 galaxies over several calibration fields is not an unreasonable goal by 2015; there are existing samples of comparable size already down to somewhat brighter magnitude limits. For instance, the DEEP2 Galaxy Redshift Survey has obtained spectra of $> 50,000$ galaxies down to $R_{AB} = 24.1$ (Newman et al. 2010, in preparation), while VVDS (Garilli et al. 2008) and zCOSMOS (Lilly et al. 2009) have both obtained spectra of $\gtrsim 20,000$ galaxies down to $i = 22.5$, and smaller samples extending to $i = 24$.

3.8.5 Photometric Redshift Calibration: Cross-correlation

An alternative method that can get around any incompleteness issues in determining redshift distributions is to employ cross-correlation information (Newman 2008). Past experience suggests we may not be successful in obtaining redshifts for all of the galaxies selected for spectroscopy; recent relatively deep ($i < 22.5$ or $R < 24.1$) surveys have obtained high-confidence ($> 95\%$ certainty) redshifts for from 42% (VVDS; Garilli et al. 2008) to 61% (zCOSMOS; Lilly et al. 2009)

to 75% (DEEP2; Newman et al. 2010, in preparation) of targeted galaxies, and extremely high-confidence ($> 99.5\%$) redshifts for 21% (VVDS) – 61% (DEEP2). Surveys of fainter galaxies have even higher rates of failure (Abraham et al. 2004). Redshift success rate in these surveys is a strong function of both galaxy properties and redshift; i.e., the objects missed are not a fair sample.

Deep infrared spectroscopy from space has problems of its own. The field of view of JWST is quite small, resulting in large cosmic variance and small sample size, and Joint Dark Energy Mission (JDEM) or Euclid spectroscopy will be limited to emission-line objects. Even with a spectroscopic completeness as high as that of SDSS ($\sim 99\%$; Strauss et al. 2002), the missed objects are not a random subsample, enough to bias redshift distributions beyond the tolerances of dark energy experiments (Banerji et al. 2008).

Even if spectroscopic follow-up systematically misses some populations, however, any well-designed spectroscopic campaign will have a large set of faint galaxies with well-determined redshifts. These can then be used to determine the actual redshift distribution for any set of galaxies selected photometrically, such as objects in some photometric redshift bin, via angular cross-correlation methods.

Because galaxies cluster together over only relatively small distances, any observed clustering between a photometric sample and galaxies at some fixed redshift, z_s , can only arise from galaxies in the photometric sample that have redshifts near z_s (Figure 3.18). Therefore, by measuring the angular cross-correlation function (the excess number of objects in one class near an object of another class on the sky, as a function of separation) between a photometric sample and a spectroscopic sample as a function of the known spectroscopic z , we can recover information about the redshift distribution of the photometric sample (hereafter denoted by $n_p(z)$; Newman 2008). If we only measure this cross-correlation, the redshift distribution would be degenerate with the strength of the intrinsic three-dimensional clustering between the two samples; however, the two-point autocorrelation functions of the photometric and spectroscopic samples provide sufficient information to break that degeneracy. Other cross-correlation techniques for testing photometric redshifts have been developed (Zhan & Knox 2006; Schneider et al. 2006), but they do not break the clustering-redshift distribution degeneracy.

In the limit where sample cosmic variance is negligible (e.g., because many statistically independent fields on the sky have been observed spectroscopically), and spectroscopic surveys cover $\gtrsim 10 \text{ deg}^2$ on the sky, Monte Carlo simulations (Newman 2008) find that the errors in determining either $\langle z \rangle$ or σ_z for a Gaussian $n_p(z)$ for a single photometric redshift bin are nearly identical, and are fit within 1% by:

$$\sigma = 9.1 \times 10^{-4} \left(\frac{\sigma_z}{0.1} \right)^{1.5} \left(\frac{\Sigma_p}{10} \right)^{-1/2} \left(\frac{dN_s/dz}{25,000} \right)^{-1/2} \left(\frac{4 h^{-1} \text{ Mpc}}{r_{0,sp}} \right)^\gamma \left(\frac{10 h^{-1} \text{ Mpc}}{r_{max}} \right)^{2-\gamma}, \quad (3.10)$$

where σ_z is the Gaussian sigma of the true redshift distribution, Σ_p is the surface density of objects in the given photometric redshift bin in galaxies arcmin^{-2} , dN_s/dz is the number of objects with spectroscopic redshifts per unit z , $r_{0,sp}$ is the true scale length of the two-point cross-correlation function between spectroscopic and photometric galaxies (the method provides a measurement of this quantity as a free byproduct); and r_{max} is the maximum radius over which cross-correlations are measured (larger radii will reduce the impact of nonlinearities, at the cost of slightly lower

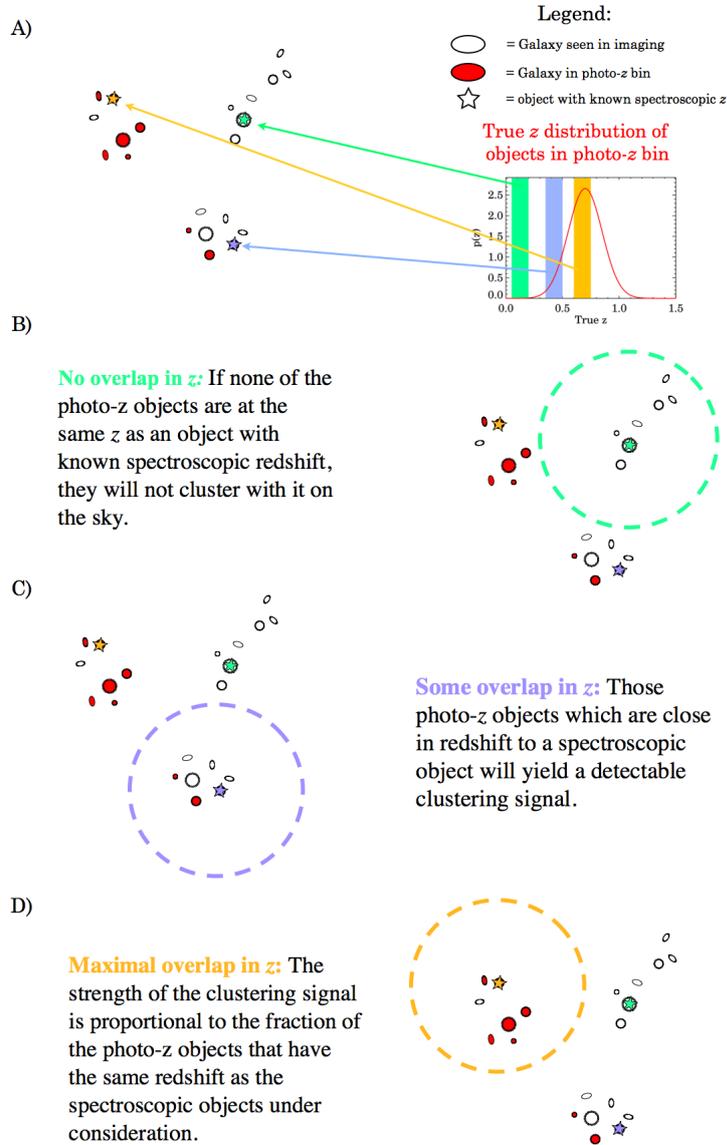


Figure 3.18: Cartoon depiction of cross-correlation photometric redshift calibration (Newman 2008). Panel A) shows the basic situation: we have imaging for many galaxies (circles/ellipses), some of which fall in a photometric redshift bin of interest (red). Galaxies that are near each other in three dimensions cluster together on the sky. We also know the spectroscopic redshifts of a smaller sample of objects (stars). The true redshift distribution for the objects in the photometric redshift bin is here assumed to be a Gaussian with mean 0.7 (plot); the stars are color-coded according to the redshift range the galaxy in question was determined to lie in with the color-coded ranges shown on the plot. B) For spectroscopic redshift objects that do not overlap in z with the photometric redshift objects, there will be no excess of neighbors that lie in the photometric redshift sample. C) If there is some overlap with the true redshift range of the photometric redshift sample, there will be some excess of neighbors around the spectroscopic object that lie in the photometric redshift bin. D) The strength of this clustering signal will be stronger the greater the fraction of the photometric redshift sample lies at the same z as the spectroscopic object in question. Because of this, we can reconstruct the true redshift distribution of the photometric redshift sample by measuring its clustering with objects of known redshift as a function of the spectroscopic z .

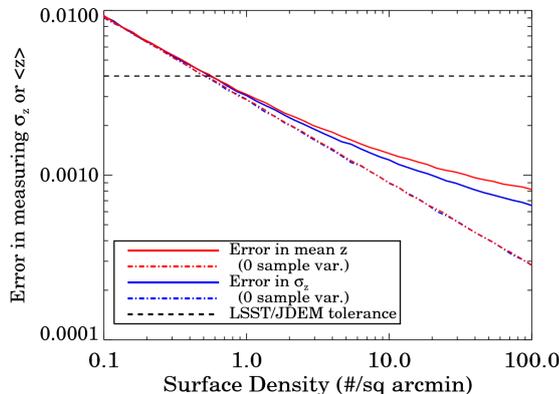


Figure 3.19: Results from Monte Carlo simulations of uncertainties in cross-correlation measurements of redshift distributions. Plotted are the rms errors in the recovery of the mean and sigma of a photometric sample distributed as a Gaussian in z with $\sigma_z = 0.1$, as a function of the surface density of that sample (representing objects in a single photometric redshift bin) on the sky. We assume a fiducial spectroscopic survey of 25,000 galaxies per unit redshift. Current and planned spectroscopic samples are sufficient to reach the required LSST photometric redshift calibration tolerances at $z < 2.5$, but larger sets of redshifts than currently available at $z > 2.5$ may be required.

S/N). Typical values of r_0 and γ for both local and $z \sim 1$ galaxy samples are $3\text{--}5 h^{-1}$ Mpc and $1.7\text{--}1.8$, respectively (Zehavi et al. 2005; Coil et al. 2006).

Errors are roughly 50% worse in typical scenarios if sample variance is significant (i.e., a small number of fields, covering relatively area, are sampled); see Figure 3.19 for an example of these scalings. Detecting non-Gaussianities such as tails in the photometric redshift distributions is straightforward in this method. The number of spectroscopic galaxies required to meet LSST photometric redshift bias and error characterization requirements is similar to the number in current and funded redshift samples for $z < 2.5$. More details on cross-correlation photometric redshift calibration and on potential systematics are given in Newman (2008).

We have tested these Monte Carlo simulations by applying cross-correlation techniques to mock catalogs produced by incorporating semi-analytic galaxy evolution prescriptions into the Millennium Run simulation (Croton et al. 2006; Kitzbichler & White 2007). Although these simulations do not perfectly match reality, they do present the same sorts of obstacles (e.g., clustering evolution) as we will encounter with LSST samples. As seen in Figure 3.20, cross-correlation techniques can accurately reconstruct the true redshift distribution of a sample of faint galaxies using only spectroscopy of a subset of bright ($R < 24.1$) objects over 4 deg^2 of sky. The dominant uncertainty in the Millennium Run reconstructions is due to the variance in the integral constraint (Bernstein 1994), which was not included in the error model of Newman (2008). This variance can be suppressed, however, by use of an optimized correlation estimator (e.g., Padmanabhan et al. 2007), and is negligible if spectroscopic surveys cover $\gtrsim 10 \text{ deg}^2$ fields.

Cross-correlation methods can accurately determine photometric error distributions for faint galaxies even if we obtain spectra of only the brightest objects at a given redshift (there are many $z = 2$ galaxies with $R < 24$, for instance, or $z = 3$ galaxies with $R < 25$). This is in contrast to methods which calibrate photometric redshifts via spectroscopic samples, due to the differences in SEDs between bright and faint galaxies and the substantial impact of confusion/blending effects on samples

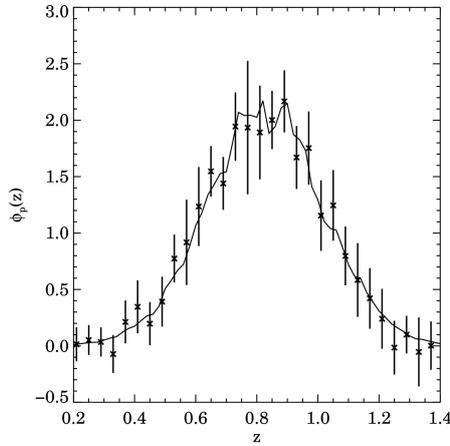


Figure 3.20: A demonstration of the recovery of redshift distributions with cross-correlation techniques. Results for a single redshift bin are shown. The solid line is the true redshift distribution of a subset of those galaxies with $M_B - 5\log h < -17$ in 24 0.5×2 degree light-cone mock catalogs constructed from the Millennium Simulation semi-analytic models of Croton et al. (2006), with the probability of being included in the set given by a Gaussian in redshift, centered at $z = 0.8$ and with dispersion $\sigma_z = 0.2$. Deviations from a Gaussian curve are due to sample (or “cosmic”) variance in the redshift distribution. Points and error bars show the median and rms variation in the cross-correlation reconstruction of this true distribution using a spectroscopic sample consisting of 60% of all galaxies down to $R_{AB} = 24.1$ in only 4 of the 24 fields. The true distribution may be reconstructed to the accuracy required by LSST using spectroscopic samples of realistic size.

of faint, high-redshift galaxies (Newman et al. 2010, in preparation). To apply cross-correlation techniques, we do not need an excessively deep sample, nor must it be uniformly complete, only well-defined. The proposed BigBOSS survey (Schlegel et al. 2009) or a proposed wide-field spectrograph on Subaru would be ideal for these purposes, providing samples of millions of galaxies, Lyman α absorbers, and QSOs with spectroscopic redshifts to $z = 2.5$, each with different clustering characteristics facilitating cross-checks. Even if BigBOSS only overlaps with LSST around the Celestial Equator, it should provide large enough numbers of redshifts to meet LSST calibration goals.

It would, of course, be preferable to obtain statistically complete spectroscopy down to a limit approaching the LSST photometric depth rather than relying on these less direct techniques. Even if spectroscopy does not prove to be sufficiently complete to test calibrations, the closer we can come to that goal, the better the photometric redshift algorithms we will be able to develop. It would be extremely difficult to tune those algorithms using cross-correlation techniques alone, without also using a set of objects with well-known redshifts and SEDs. Furthermore, as seen in Equation 3.10, the better our photometric redshifts are (i.e., the smaller σ_z is), the more precisely we can calibrate them. Making sure we are achieving the tight calibration requirements for LSST dark energy studies will require cross-checks. Cross-correlation techniques will allow us to do this by repeating the analysis with very different spectroscopic samples; if all is working properly, the redshift distributions from each spectroscopic sample should agree. As an example, one could use one set of spectroscopy going faint in the deep LSST calibration fields, and another, shallower or sparser set, covering the main wide field LSST survey. Alternatively, one could use spectroscopic subsamples with very different clustering properties (e.g., star forming galaxies versus luminous

red galaxies) to do the test. The recovered redshift distribution for a photometric redshift bin must be consistent when applying any variety of type, redshift, and magnitude cuts to the spectroscopic sample if the reconstruction is accurate.

References

- Abazajian, K. N. et al., 2009, *ApJS*, 182, 543
 Abraham, R. G. et al., 2004, *AJ*, 127, 2455
 Ball, N. M., Brunner, R. J., Myers, A. D., Strand, N. E., Alberts, S. L., & Tcheng, D., 2008, *ApJ*, 683, 12
 Banerji, M., Abdalla, F. B., Lahav, O., & Lin, H., 2008, *MNRAS*, 386, 1219
 Baum, W. A., 1962, in IAU Symposium, Vol. 15, Problems of Extra-Galactic Research, G. C. McVittie, ed., p. 390
 Benítez, N., 2000, *ApJ*, 536, 571
 Bernstein, G., & Huterer, D., 2009, ArXiv e-prints, 0902.2782
 Bernstein, G. M., 1994, *ApJ*, 424, 569
 Bramich, D. M. et al., 2008, *MNRAS*, 386, 887
 Bruzual, G., & Charlot, S., 2003, *MNRAS*, 344, 1000
 Budavári, T., Szalay, A. S., Connolly, A. J., Csabai, I., & Dickinson, M., 2000, *AJ*, 120, 1588
 Carliles, S., Budavári, T., Heinis, S., Priebe, C., & Szalay, A., 2008, in Astronomical Society of the Pacific Conference Series, Vol. 394, Astronomical Data Analysis Software and Systems XVII, R. W. Argyle, P. S. Bunclark, & J. R. Lewis, eds., p. 521
 Coil, A. L., Newman, J. A., Cooper, M. C., Davis, M., Faber, S. M., Koo, D. C., & Willmer, C. N. A., 2006, *ApJ*, 644, 671
 Collister, A. A., & Lahav, O., 2004, *PASP*, 116, 345
 Connolly, A. J., Csabai, I., Szalay, A. S., Koo, D. C., Kron, R. G., & Munn, J. A., 1995, *AJ*, 110, 2655
 Covey, K. R. et al., 2007, *AJ*, 134, 2398
 Croton, D. J. et al., 2006, *MNRAS*, 365, 11
 Ellis, S. et al., 2009, Proceedings of the August 2009 Meeting of the SPIE, San Diego. in press
 Freeman, P. E., Newman, J. A., Lee, A. B., Richards, J. W., & Schafer, C. M., 2009, *MNRAS*, 1053
 Garilli, B. et al., 2008, *A&A*, 486, 683
 Gwyn, S. D. J., 2008, *PASP*, 120, 212
 Gwyn, S. D. J., & Hartwick, F. D. A., 1996, *ApJL*, 468, L77
 Hoekstra, H. et al., 2006, *ApJ*, 647, 116
 Ilbert, O. et al., 2006a, *A&A*, 457, 841
 —, 2006b, *A&A*, 453, 809
 Ivezić, Ž. et al., 2008a, *ApJ*, 684, 287
 —, 2007, *AJ*, 134, 973
 —, 2008b, ArXiv e-prints, 0805.2366
 Jurić, M. et al., 2008, *ApJ*, 673, 864
 Kashikawa, N. et al., 2003, *AJ*, 125, 53
 Kitzbichler, M. G., & White, S. D. M., 2007, *MNRAS*, 376, 2
 Koo, D. C., 1985, *AJ*, 90, 418
 Lanzetta, K. M., Yahil, A., & Fernández-Soto, A., 1996, *Nature*, 381, 759
 Lilly, S. J. et al., 2009, VizieR Online Data Catalog, 217, 20070
 Loh, E. D., & Spillar, E. J., 1986, *ApJ*, 303, 154
 Ma, Z., & Bernstein, G., 2008, *ApJ*, 682, 39
 Mobasher, B. et al., 2007, *ApJS*, 172, 117
 Monet, D. G. et al., 2003, *AJ*, 125, 984
 Munn, J. A. et al., 2004, *AJ*, 127, 3034
 Newman, J. A., 2008, *ApJ*, 684, 88
 Neyman, J., 1938, Lectures and Conferences on Mathematical Statistics
 Padmanabhan, N., White, M., & Eisenstein, D. J., 2007, *MNRAS*, 376, 1702
 Peterson, J. R. et al., 2005, Bulletin of the American Astronomical Society, Vol. 37, LSST Atmosphere and Instrument Simulation. p. 1205

Chapter 3: References

- Pier, J. R., Saha, A., & Kinman, T. D., 2003, *Information Bulletin on Variable Stars*, 5459, 1
- Roweis, S. T., & Saul, L. K., 2000, *Science*, 290, 2323
- Saha, A., & Monet, D., 2005, *Bulletin of the American Astronomical Society*, Vol. 37, LSST Astrometric Science. p. 1203
- Sawicki, M. J., Lin, H., & Yee, H. K. C., 1997, *AJ*, 113, 1
- Schlegel, D. J. et al., 2009, ArXiv e-prints, 0904.0468
- Schneider, M., Knox, L., Zhan, H., & Connolly, A., 2006, *ApJ*, 651, 14
- Sesar, B. et al., 2007, *AJ*, 134, 2236
- Springel, V. et al., 2005, *Nature*, 435, 629
- Stabenau, H. F., Connolly, A., & Jain, B., 2008, *MNRAS*, 387, 1215
- Strauss, M. A. et al., 2002, *AJ*, 124, 1810
- VanderPlas, J., & Connolly, A., 2009, *AJ*, submitted
- Wadadekar, Y., 2005, *PASP*, 117, 79
- Yip, C. W. et al., 2004, *AJ*, 128, 585
- Zehavi, I. et al., 2005, *ApJ*, 630, 1
- Zhan, H., & Knox, L., 2006, *ApJ*, 644, 663