

QUASAR CLASSIFICATION USING COLOR AND VARIABILITY

CHRISTINA M. PETERS^{1,8}, GORDON T. RICHARDS¹, ADAM D. MYERS², MICHAEL A. STRAUSS³, KASPER B. SCHMIDT⁴,
 ŽELJKO IVEZIĆ⁵, NICHOLAS P. ROSS⁶, CHELSEA L. MACLEOD⁶, AND RYAN RIEGEL⁷

¹Department of Physics, Drexel University, Philadelphia, PA 19104, USA; Christina.M.Peters@drexel.edu

²Department of Physics and Astronomy, University of Wyoming, Laramie, WY 82071, USA

³Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

⁴Department of Physics, University of California, Santa Barbara, CA 93106, USA

⁵Astronomy Department, University of Washington, Seattle, WA 98195, USA

⁶Institute for Astronomy, The University of Edinburgh, Edinburgh EH9 3HJ, UK

⁷Skytree, Inc., 1731 Technology Drive, Suite 700, San Jose, CA 95110, USA

Received 2015 May 6; accepted 2015 August 12; published 2015 September 28

ABSTRACT

We conduct a pilot investigation to determine the optimal combination of color and variability information to identify quasars in current and future multi-epoch optical surveys. We use a Bayesian quasar selection algorithm to identify 35,820 type 1 quasar candidates in a 239 deg² field of the Sloan Digital Sky Survey (SDSS) Stripe 82, using a combination of optical photometry and variability. Color analysis is performed on 5-band single- and multi-epoch SDSS optical photometry to a depth of $r \sim 22.4$. From these data, variability parameters are calculated by fitting the structure function of each object in each band with a power-law model using 10 to >100 observations over timescales from ~ 1 day to ~ 8 years. Selection was based on a training sample of 13,221 spectroscopically confirmed type-1 quasars, largely from the SDSS. Using variability alone, colors alone, and combining variability and colors we achieve 91%, 93%, and 97% quasar completeness and 98%, 98%, and 97% efficiency, respectively, with particular improvement in the selection of quasars at $2.7 < z < 3.5$ where quasars and stars have similar optical colors. The 22,867 quasar candidates that are not spectroscopically confirmed reach a depth of $i \sim 22.0$; 21,876 (95.7%) are dimmer than coadded i -band magnitude of 19.9, the cutoff for spectroscopic follow-up for SDSS on Stripe 82. Brighter than 19.9, we find 5.7% more quasar candidates without confirming spectra in sky regions otherwise considered complete. The resulting quasar sample has sufficient purity (and statistically correctable incompleteness) to produce a luminosity function comparable to those determined by spectroscopic investigations. We discuss improvements that can be made to the process in preparation for performing similar photometric selection and science on data from post-SDSS sky surveys.

Key words: catalogs – quasars: general – surveys

Supporting material: FITS file

1. INTRODUCTION

The identification of large numbers of quasars/active galactic nuclei (AGNs) over a broad range of redshift and luminosity is crucial for many science projects. Work that requires object densities higher than have been provided to date by spectroscopic surveys includes cross-correlating the catalogs with the cosmic microwave background (Giannantonio et al. 2008) to constrain dark energy; using quasars to measure cosmic magnification (Scranton et al. 2005); finding binary quasars which can be used to test the merger hypothesis of quasars (Hennawi et al. 2010); finding gravitationally lensed quasars (Oguri et al. 2006); constraining quasar evolution (Myers et al. 2006); studying dust in galaxies (Ménard et al. 2010); and broader cosmological studies (Leistedt et al. 2013).

Historically, quasar candidates have been identified by virtue of their colors, variability, and (lack of) proper motion—but generally not through all of these methods combined. The standard way of identifying large numbers of candidate quasars is to make “color cuts” using optical (or infrared) photometry (e.g., Warren et al. 2000; Richards et al. 2002; Croom et al. 2004; Lacy et al. 2004; Stern et al. 2005; Maddox et al. 2012; Assef et al. 2013). This is because the majority of unobscured

quasars at $z < 2.5$ are much bluer than the majority of stars in the optical and are much redder in the infrared. However, this process is neither complete (identifying all true quasars) nor efficient (minimizing false positives). Such methods do an effective job of identifying a large number of interesting objects with relatively little effort; however, better methods are needed to scale to future surveys in a way that allows scientific analysis without the need for spectroscopic confirmation.

In addition to classification by color, time-domain data make variability a promising way for classifying objects. For examples of such work, see Koo et al. (1986), Hughes et al. (1992), Vanden Berk et al. (2004), de Vries et al. (2005), Sesar et al. (2007), Kelly et al. (2009), Kozłowski et al. (2010), MacLeod et al. (2010, 2011, 2012), Schmidt et al. (2010), Butler & Bloom (2011), and Graham et al. (2014). Specifically, quasars exhibit stochastic, aperiodic variability with variations of the order of 10% on the timescale of years (de Vries et al. 2003; Vanden Berk et al. 2004). The amplitude and timescale of this variability are sufficiently distinctive to allow one to identify an object as a candidate quasar.

Many current and future astronomical imaging surveys (SkyMapper: Keller et al. 2007; Palomar Transient Factory: Law et al. 2009; Pan-STARRS: Kaiser et al. 2010; DES: The Dark Energy Survey Collaboration 2005; LSST: Ivezić et al. 2008) are focusing on time-domain astronomy and in anticipation it is important to determine the effectiveness of

⁸ For correspondence regarding this article, please write to C. M. Peters.

classification using variability information. These surveys will observe areas of the sky many times. There is great hope that variability selection will fill in the gaps in color selection methods (or replace color selection entirely). Indeed, investigations such as Schmidt et al. (2010), MacLeod et al. (2011), and Butler & Bloom (2011) have been quite successful. However, variability-only selection suffers from its own set of problems. For example, high-redshift quasars can be lost when using a fixed observed-frame variability analysis: Ly α absorption reduces the quasar continuum in blue bands and the redder bands have larger photometric errors for fainter objects. In addition, variability increases with lower luminosity (e.g., Vanden Berk et al. 2004), but so does the host galaxy contribution—potentially complicating the selection of such objects without careful difference imaging to remove the host galaxy contribution. Thus it is important to investigate how well variability selection works by itself versus being combined with other methods (e.g., colors and astrometry).

The premise of this project is to *simultaneously* use the distinctive and quantifiable characteristics of color and variability to distinguish quasars from stars and inactive galaxies. The Sloan Digital Sky Survey (SDSS; York et al. 2000) repeatedly imaged a 2:5 equatorial section of the sky referred to as Stripe⁹ 82 (Abazajian et al. 2009; Annis et al. 2014; Jiang et al. 2014). The light curves of spectroscopically confirmed quasars and stars from Stripe 82 give us the information we need to develop and test classification of quasars.

The specific goal of this project is to use color, variability, and astrometric data in combination with modern machine learning techniques to uncover previously unidentified quasars in the SDSS Stripe 82 region and to pave the way for improved multi-faceted selection in the future. The goal is not necessarily to produce the most complete or efficient catalog possible, but to test the combined use of colors and variability data in classification. In this pilot investigation we make some simplifications to the process that will be explored in more detail in future work. Specifically, we concentrate on point sources to avoid the problem of the host galaxy washing out the variable nucleus (reducing our sensitivity to low-redshift quasars), we utilize a simple power-law model of variability as opposed to more sophisticated (but not necessarily “correct”) models such as the damped random walk, we use variability data from each band separately instead of merging them together, and we take a simplistic approach to combining photometric redshift information from different methods. Each of these simplifications for this pilot study is worthy of their own separate investigation to determine how to best deal with these issues.

A shortcoming of the traditional quasar identification process is that it usually involves selecting quasar candidates by identifying them as outliers using cuts in the *observed data space* (e.g., selecting all point sources with $u - g < 0.6$). Our classification instead makes *simultaneous* use of all of the data types available and uses modern statistical techniques (based on kernel density estimation; KDE) to make cuts in *probability space* (e.g., objects with an expected quasar probability greater than 50%). We will extend the methods developed by our group (Richards et al. 2004, 2009a, 2009b; Riegel et al. 2008) and others (e.g., Suchkov et al. 2005; Ball et al. 2006; Davoodi

et al. 2006; Bailer-Jones et al. 2008; Gao et al. 2008; D’Abrusco et al. 2009; Guy et al. 2010; Schmidt et al. 2010; Abraham et al. 2012; Bovy et al. 2012; Peng et al. 2012; Gupta et al. 2014) to create a classification algorithm for time-domain focused sky surveys. While this approach has been shown to work well in the past (e.g., Richards et al. 2004, 2009a), in future work we also intend to explore other modern statistical techniques such as described by Feigelson & Babu (2012) and the references therein.

The quasar candidates that result from application of this method are only identified *photometrically*; they lack spectroscopy which not only would confirm the type of an object, but it crucially also would determine the redshift. There are many sophisticated methods for estimating photometric redshifts (e.g., Rowan-Robinson et al. 2008; Salvato et al. 2009); we use the algorithm described in Richards et al. (2001) and Weinstein et al. (2004) which ranks among the most accurate for (luminous) quasar photometric redshift estimates. We improve this process further by using the effective prismatic effects of the Earth’s atmosphere as a low-resolution spectrograph (Kaczmarczik et al. 2009). In short, the positions of quasars, with their strong emission features, are a function of passband and redshift. This behavior of quasars allows us to uniquely incorporate *astrometric* information into our photometric redshift estimates.

Our work provides a stepping stone for quasar classification for future surveys such as the Large Synoptic Survey Telescope¹⁰ (LSST). Eventually, each region of LSST will be imaged about 200 times in each filter over the 10 years of the survey, allowing for the study of the variability of the object on scales of minutes to a decade. This focus on time-domain astronomy is an exciting new era in surveys, but the lack of spectroscopy creates a problem for confirming the type of object. As the number of spectroscopic fibers allocated to quasar identification pales in comparison to the number of photometrically detected objects that merit spectroscopic follow-up, it is only through complete and efficient object classification coupled with accurate redshift estimates that we can overcome the lack of spectroscopy in LSST and other future astronomical surveys and maximize their science output.

The layout of this paper is as follows. In Section 2 we introduce the SDSS Stripe 82 data that we will use. We then describe how the variability parameters used for classification are calculated. In Section 3 we summarize the Non-parametric Bayesian Classification Kernel Density Estimation (NBC KDE) selection algorithm and describe how it is used in this case. We test the various classification parameters and determine the optimal combination in Section 4. Then, in Section 5, we build the quasar candidate catalog using these optimal parameters, first using the full quasar training set, then using the training set divided into redshift bins to perform simultaneous classification and redshift estimation. In Section 6 we describe how the astrometric parameters are calculated, then estimate photometric and astrometric redshifts for all the candidate quasars. Next, we describe a cut to remove contamination and describe the final catalog of quasar candidates in Section 7. In Section 8 we compare to cuts in variability space and to color-based quasar selection, and calculate number counts and a luminosity function for the

⁹ sdss.org/legacy/stripe82.html

¹⁰ lsst.org

Table 1
Master Quasar Catalog

Source	Description	w/Spectra	w/o Spectra	Training Set
Table 5 from Schneider et al. (2010)	SDSS I/II	105472	0	6082
Croom et al. (2004)	2QZ	9663	0	0
Croom et al. (2009)	2SLAQ	8881	0	1576
S. M. Croom et al. (2015, in preparation)	AUS	2200	0	1706
Kochanek et al. (2012)	AGES	2844	4	0
Lilly et al. (2007) and Elvis et al. (2009)	COSMOS	259	0	0
Fan et al. (2006) and Jiang et al. (2008)	$z > 5.8$	27	0	0
Pâris et al. (2014)	SDSS-III/BOSS	168820	0	7383
Ross et al. (2012)	MMT	836	0	278
Richards et al. (2009b)	NBCKDE Photometric Catalog	174663	965542	9061
Bovy et al. (2011)	XDQSO Photometric Catalog	142567	682831	7088
Table 5 of Papovich et al. (2006)	BROADLINE objects	104	0	0
Table 5 of Glikman et al. (2006)	$z \sim 4$	10	0	0
Tables 4 and 6 of Maddox et al. (2012)	KX-selected	3608	0	986
Total		274329	1301846	13221

candidates. We discuss possible next steps in Section 9 and conclude in Section 10.

Cosmology-dependent parameters are determined using $H_o = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.3$, and $\Omega_\Lambda = 0.7$ (Hinshaw et al. 2013). Throughout this paper magnitudes will be reported on the AB system of Oke & Gunn (1983).

2. DATA

In this section, we describe the origin of the data and the parameters used for classification by our algorithm. Section 2.1 describes the imaging data and 2.2 the spectroscopic data. Sections 2.3 and 2.4 discuss derivation of the color and variability classification parameters, respectively. In principle, we could use astrometric information for classification as well; however, for this pilot study we have limited astrometric data to estimate photometric redshifts as discussed in Section 6. Machine learning algorithms need both training sets to find patterns in the data and a test set of data to verify that these patterns are useful; these data sets are described in Section 2.5.

2.1. SDSS Stripe 82

The SDSS is an optical survey that has used the 2.5 m Sloan telescope (Gunn et al. 2006) at Apache Point Observatory in New Mexico to map 14,500 deg² of the sky (Aihara et al. 2011). Photometry was performed with a drift-scan CCD camera (Gunn et al. 1998) taking nearly simultaneous 54.1 s exposures in five broad optical bands (u , g , r , i , and z) between 3000 and 10,000 Å (Fukugita et al. 1996).

The imaging data used in our analysis consists of objects solely from the SDSS Stripe 82 area, which were made available as part of SDSS Data Release 7 (DR7; Abazajian et al. 2009) and includes observations from 1999 October to 2007 November. The Stripe 82 region covers a 2°5 wide “stripe” on the celestial equator from right ascension $\sim 300^\circ$ to $\sim 60^\circ$ in the Southern Galactic Cap. Repeated observations were performed on this region throughout the SDSS I/II, with increasing frequency as part of the SDSS Supernova Survey (Frieman et al. 2008), with ~ 100 repeat imaging scans by the end of observations. The initial observations were done under optimal seeing, sky brightness, and photometric conditions. The supernova survey runs were done on useable nights, but

under less than optimal conditions. We limit our analysis to those objects detected as point sources.

The multiple observations on Stripe 82 were aligned and stacked into a coadded catalog described in Annis et al. (2014; see also Jiang et al. 2014 and Huff et al. 2014). This catalog uses 20–40 observations on the region, mostly the early runs under optimal conditions. The data were downloaded from the SDSS Stripe 82 Catalog Archive Server (CAS).¹¹ Database entries having SDSS “run” numbers of 106 and 206, representing objects with co-added photometry, were extracted along with the individual epoch photometry for each of these objects in order to generate light curves.¹² The single epoch images go to a depth of $r \sim 22.4$ (5σ), with a median seeing of 1".4. Coaddition of the imaging data reaches ~ 2 mag deeper and improves the median seeing to 1".1. The improvement in using coadded magnitudes over single epoch magnitudes for classification is demonstrated in Section 4.2; see also Ivezić et al. (2007).

2.2. Master Quasar Catalog

Definition of our quasar training set requires a subsample with spectroscopic confirmation. Our primary source of spectroscopy comes from a “Master” Quasar Catalog (MQC), described in Section 2.1 of Richards et al. (2015), containing over 1.5 million sources, for which over 250,000 have confirming spectroscopy. This data set consists of sources within the SDSS survey areas and draws objects from the sources described in Table 1.

This quasar sample represents nearly every quasar known fainter than $i \sim 16$ (including candidate photometric quasars) at the time of Data Release 10 (DR10; Ahn et al. 2014) of SDSS-III (Eisenstein et al. 2011; Dawson et al. 2013). The majority of the confirmed quasars come from the SDSS I/II quasar catalog, which is described in detail by Richards et al. (2002) and Schneider et al. (2010) and from the SDSS-III/BOSS quasar catalog, which is described in detail by Ross et al. (2012) and Pâris et al. (2014).

¹¹ <http://cas.sdss.org/stripe82/en>

¹² This process has since been made somewhat easier through the use of a unifying “thingIndex” table in Data Release 12 (Alam et al. 2015): <http://skyserver.sdss.org/dr12/en/help/browser/browser.aspx>

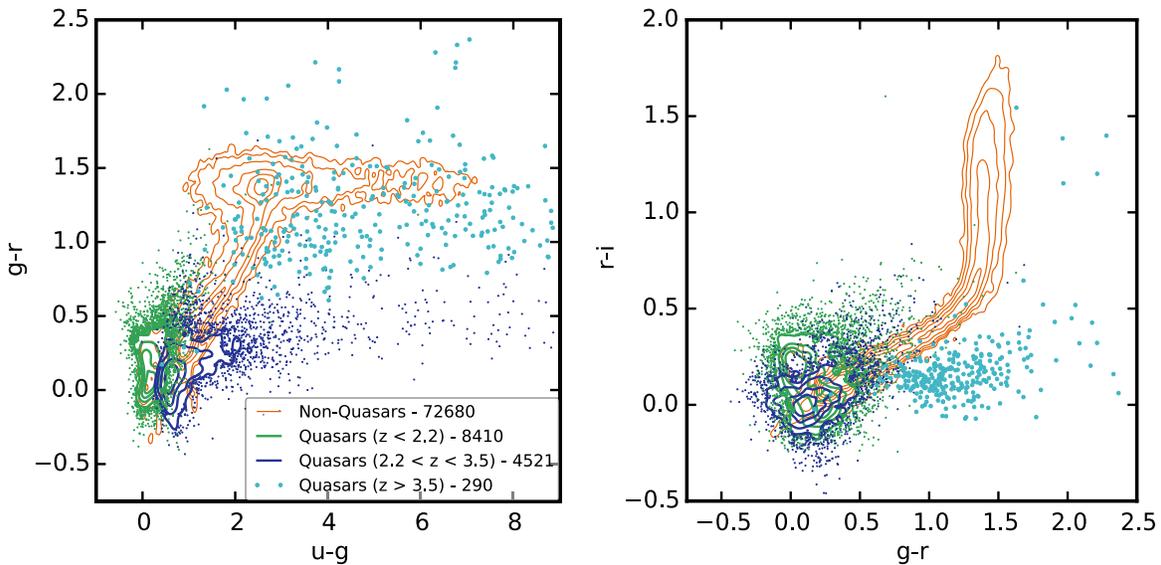


Figure 1. Quasar and non-quasar training sets in two projections of the SDSS color space using coadded photometry. Non-quasars (shown in orange contours), such as stars and compact galaxies, are considered contaminants when trying to accurately classify quasars (shown in cool colors). Notice the number of non-quasars in the region in which mid-redshift quasars ($2.2 < z < 3.5$; shown as dark blue contours and scatter points for outliers) lie. This overlap makes it difficult to accurately classify an object in this region as a quasar or non-quasar and motivates searches for alternative methods of classification, like variability. Quasars are shown as three redshift regions: low-redshift ($z < 2.2$; shown as green contours and scatter points for outliers), mid-redshift, and high-redshift ($z > 3.5$; shown as light blue dots). The extension of the non-quasar color space at $g - r \sim 1.4$ is not real, but is an artifact of including objects with large u -band photometric errors (and thus spilling into the true quasar parameter space).

The SDSS I/II quasars were primarily color selected (with some radio and X-ray selection) over a broad redshift range ($0 < z < 5$). Richards et al. (2002) describe the quasar target selection of the main quasar survey, which went to $i < 19.1$ for quasars with colors consistent with $z < 3$ and to $i < 20.2$ for quasars expected to be at higher redshifts. On Stripe 82, deeper targeting was performed (Adelman McCarthy et al. 2006) going to $i = 19.9$ and $i = 20.4$, respectively, in targeting “chunk” 22; to $i = 20.2$ (for low-redshift sources) and $i = 20.65$ (for radio sources) in targeting chunk 48; and to $i < 21$ for sources more variable (between two epochs) than 3σ (and 0.1 mag) in both g and r in targeting chunk 73. The BOSS quasars (focused on $2.2 < z < 3.5$; Ross et al. 2012) were, in addition to color selection, also targeted by variability (on Stripe 82). This variability selection is described in Palanque-Delabrouille et al. (2011) and uses an algorithm that was also based on the same parameterization of variability as used herein (see Section 2.4). Thus it is interesting to see if our method finds additional quasars beyond those already spectroscopically confirmed. Quasar candidates in our catalog that are previously known from SDSS-I/II and SDSS-III spectroscopy are indicated as such in our catalog; see the Appendix.

2.3. Classification Parameters: Colors

The optical color information used in our analysis consists of the four adjacent SDSS colors ($u - g$, $g - r$, $r - i$, and $i - z$), which were determined from the cataloged photometry using point-spread function (PSF) magnitudes, corrected for Galactic extinction (Schlegel et al. 1998). We used both single-epoch colors, from a single observation of the object, and the coadded colors, from the Annis et al. (2014) catalog.

The level of contamination from stars and galaxies varies significantly in various regions of colorspace; see Figure 1. Optical surveys for quasars often use relatively simple color cuts (drawing lines of demarcation in these color spaces) to

select objects that are likely to be quasars. In SDSS, outliers from the stellar locus in the color space were potential spectroscopic target candidates (Richards et al. 2002). The $ugri$ bands were used to identify low-redshift quasars and the $griz$ bands for high-redshift quasars. For low- and high-redshift quasars, selecting by colors is effective, but mid-redshift quasars ($2.2 < z < 3.5$) occupy the same region of color space as many stars and contamination becomes a serious problem. Note how the mid-redshift quasars, shown as dark blue contours and scatter points in Figure 1, overlap with the non-quasars, shown as orange contours. It is most efficient to choose quasars outside of this redshift region for spectroscopic follow-up, but this creates a strong selection effect in the quasar sample. For efficient selection of mid-redshift quasars, it becomes necessary to have another method to distinguish the quasars from non-quasars and this is where the variable nature of quasars becomes particularly useful.

2.4. Classification Parameters: Variability

Most quasars vary at optical wavelengths by about 10% over several years, which distinguishes them from most normal galaxies and stars (de Vries et al. 2003; Vanden Berk et al. 2004). Most variable stars vary periodically or quasi-periodically (Richards et al. 2012) and with smaller amplitude, but quasars generally show no periodic variability (Bailer-Jones 2012; Andrae et al. 2013). While the physical causes for the variability in quasars are not well understood (see Dexter & Agol 2011 for a recent investigation), the nature of the variability enables one to distinguish quasars from non-quasars.

We use the structure function (SF) to characterize variability by quantifying the amplitude of variability as a function of the time difference between paired observations. For this analysis, based on empirical experiment (balancing the number of epochs with the quality of the data), we required that the

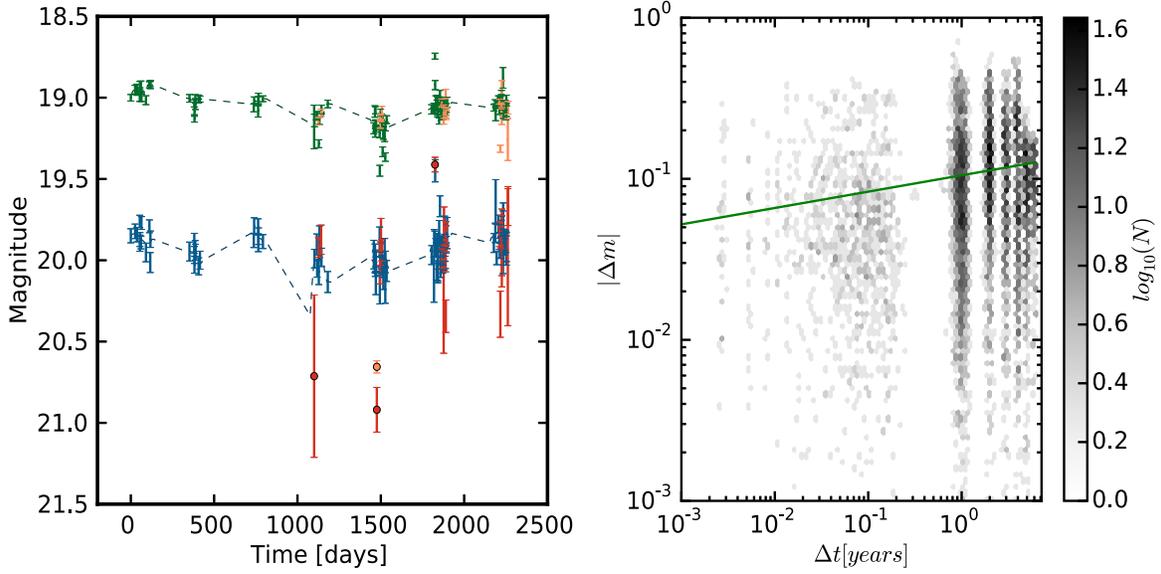


Figure 2. g and u -band light curves (left panel) and g -band structure function fit with a power-law model (right panel) of SDSS J013417.81–005036.2, a redshift 2.26 quasar from SDSS Stripe 82 (also shown in Figure 12). This quasar is shown as an example representative of the data set. Left panel: there are 126 observations in the g -band. The 106 observations that meet the PSF-width and the airmass requirements are shown as green points with error bars, while those that were removed are shown in orange. The dark green dashed line is the running median (with a window of 50 days and steps of 5 days) calculated from the g -band observations. The orange dot was removed from the light curve because it is more than 0.25 mag from the median. The u -band observations are similarly shown in blue and red. Right panel: the pairs of photometric points from the g -band light curve in the left panel are shown as a hex-bin density plot where the darkness of the hex bin indicates the number of points in that bin. The power law fit is shown as a green line. The method for calculating the structure function and the equation used to fit the structure function are detailed in Section 2.4. In the case of this object, the fitting algorithm gives $A_g = 0.105$ and $\gamma_g = 0.102$. The points removed as outliers in the left panel would only contribute $|\Delta m| > 0.25$ mag values.

FWHM of the PSF fit in the r -band be less than $2''$ and the airmass in the r -band be less than 1.575 for the observation to be included. These cuts remove approximately 15% of observations. After this procedure, we found that a small number of non-astrophysical outliers in the light curve still must be removed; these points are such strong outliers that we are not concerned that removing them is compromising the variability analysis. Similar to the approach in Schmidt et al. (2010), we accomplish this by calculating a running median light curve then removing all measurements with a difference between the median light curve and the observed magnitude greater than 0.25 mag (Figure 2 left panel). The SF is calculated in all of the SDSS bands where at least 10 observations remain after these cuts.

There are other methods currently being used to characterize the variability of quasars, including Slepian wavelet variance (SWV; Graham et al. 2014), AutoRegressive Moving Average, or ARMA, processes (Kasliwal et al. 2015), or damped random walk (DRW; Kelly et al. 2009; Kozłowski et al. 2010). Future work could consider using these methods instead of the SF.

In our work, the SF is defined as the root mean square magnitude difference as a function of time lag between magnitude measurements:

$$V^2(\Delta t) = \langle (m(t) - m(t + \Delta t))^2 \rangle. \quad (1)$$

In the above equation, $m(t) - m(t + \Delta t)$ is the measured magnitude difference between two observations in a given band and Δt is the time difference between the two observations in the observer’s frame. The SDSS has a high cadence of observations during the fall months each year and then gaps of ~ 9 months before the next set of observations. This irregular sampling in the light curve (Figure 2 left panel) results in a SF with gaps (Figure 2 right panel).

The SF can be modeled as a power law (Equation (3) in Schmidt et al. 2010):

$$V_{\text{PowerLaw}}(\Delta t|A, \gamma) = A \left(\frac{\Delta t}{1 \text{ year}} \right)^\gamma. \quad (2)$$

Such a parameterization is not effective at describing the underlying type of variability or the mechanism for it, but provides a sufficiently robust statistical description for the timescales (~ 1 day to ~ 8 years) covered by our data (Schmidt et al. 2010) to distinguish variable sources from non-variable sources, which is our objective. Using this model for the SF, we find that 93% of quasars are more variable than non-variable stars on average (using white dwarfs as representative) and show more growth in variability at longer timescales than 80% of non-quasar point sources.

The variability can also be modeled as a DRW (Kelly et al. 2009, Kozłowski et al. 2010, MacLeod et al. 2010), which predicts the following form of the SF:

$$V_{\text{DRW}}(\Delta t|\sigma, \tau) = \sqrt{2} \sigma (1 - e^{-\Delta t/\tau})^{\frac{1}{2}}. \quad (3)$$

To first order in Δt , the DRW behaves as:

$$V_{\text{DRW}}(\Delta t|\sigma, \tau) \sim \sqrt{2} \sigma \left(\frac{\Delta t}{\tau} \right)^{\frac{1}{2}}, \quad (4)$$

a realization of Equation (2) where $\gamma = 1/2$. In short, the DRW model is similar to the power-law model except that it truncates the growth of the magnitude differences at some characteristic timescale. For the sake of this proof of concept, the power-law model will suffice and is what we shall use hereafter. In future work we will investigate whether a more

sophisticated model, such as the DRW model, improves quasar selection; however, even that model may be too simplistic to describe quasar variability across the range of timescales probed by modern optical monitoring data (Mushotzky et al. 2011; Zu et al. 2013; Graham et al. 2014; Kasliwal et al. 2015).

To fit the power-law model to the observational data for each object we used the likelihood function (Equation (4) in Schmidt et al. 2010):

$$\mathcal{L}(A, \gamma) = \prod_{j,k} L_{j,k}, \quad (5)$$

where $L_{j,k}$ is the likelihood of observing one particular magnitude difference $\Delta m_{j,k}$ between two light curve points separated by $\Delta t_{j,k}$. To determine the maximum likelihood of a Gaussian distribution, as in the case of the noise and intrinsic photometric variability, the likelihood function is:

$$\mathcal{L} = \prod_i \frac{1}{\sqrt{2\pi} \sigma_i} \exp\left(-\frac{1}{2} \frac{(\Delta m_i)^2}{\sigma_i^2}\right). \quad (6)$$

The variance $\sigma^2 = (A(t_j - t_k)^\gamma)^2 + \sigma_{\text{phot},j}^2 + \sigma_{\text{phot},k}^2$ represents the scatter around the line that we are fitting and includes both intrinsic variability and noise. The $\sigma_{\text{phot},j}$ and $\sigma_{\text{phot},k}$ are the measured photometric errors on the measurements. Both the noise and the intrinsic photometric variability are assumed to have a Gaussian distribution.

If there is no variability or measurement noise, the SF would be equal to zero for all Δt . The likelihood function now has the form:

$$\mathcal{L} = \prod_{j>k} \frac{1}{\sqrt{2\pi} \left[(A(t_j - t_k)^\gamma)^2 + \sigma_{\text{phot},j}^2 + \sigma_{\text{phot},k}^2 \right]} \times \exp\left(-\frac{1}{2} \frac{(m_j - m_k)^2}{(A(t_j - t_k)^\gamma)^2 + \sigma_{\text{phot},j}^2 + \sigma_{\text{phot},k}^2}\right). \quad (7)$$

The product only counts those observations where $j > k$, so there is no double counting and there are $\frac{n(n-1)}{2}$ data pairs where n is the number of observations. We require the fitting to return physical values, $A > 0$ and $\gamma > 0$, so that the power-law exponent and the average variability on a 1-year timescale are positive. This is because we are fitting $|\Delta m|$ and $|\Delta t|$ and all light curves will have some level of measurement noise, causing $A > 0$. Non-variable stars generally have γ approaching 0. The expected increasing deviation from the mean for quasars with increasing $|\Delta t|$ will cause $\gamma > 0$.

We found a strong degeneracy between A and γ when maximizing the likelihood. To break this degeneracy, we applied a Gaussian prior to the likelihood on A . With a typical observing cadence of ~ 1 year, the prior is centered on the observed median $|\Delta m|$ value, \hat{A} , at $0.5 \text{ years} < |\Delta t| < 1.5$ years and the standard deviation, σ_A , for those values. We place no explicit prior on γ in the likelihood, but the requirement that $\gamma > 0$ functions as a flat prior. In addition to breaking the degeneracy, this prior encourages the minimization routine to converge on a realistic A value more quickly. The cadence of the Stripe 82 data gives sufficient data points over this time difference to support this constraint. We combine the log of the

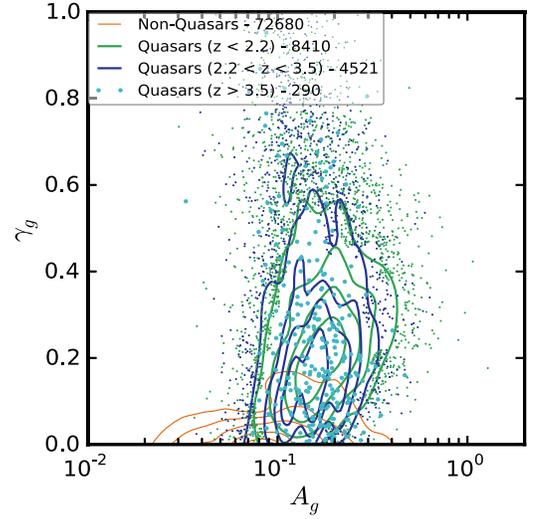


Figure 3. Quasar and non-quasar training sets in variability parameter space for the g -band observations. Note that, unlike in the color-color plots in Figure 1, there are no distinct changes in the variability parameters as a function of quasar redshift in this parameter space. This is advantageous because it allows us to separate the quasars from the non-quasars in the variability space without extreme changes in completeness at specific redshifts, as seen with color selection. Non-quasars, such as stars and normal galaxies, are shown in orange contours. Quasars are shown in cool colors as three redshift regions: low-redshift ($z < 2.2$; shown as green contours and scatter points for outliers), mid-redshift ($2.2 < z < 3.5$; shown as dark blue contours and scatter points for outliers), and high-redshift ($z > 3.5$; shown as light blue dots).

likelihood function and the prior as follows:

$$\begin{aligned} S &= -2 \frac{1}{N} \log(\mathcal{L}) + P(A) \\ &= \frac{2}{n(n-1)} \sum_{j>k} \left[\log\left((A(t_j - t_k)^\gamma)^2 + \sigma_j^2 + \sigma_k^2 \right) \right. \\ &\quad \left. + \frac{(m_j - m_k)^2}{(A(t_j - t_k)^\gamma)^2 + \sigma_j^2 + \sigma_k^2} \right] + \frac{(A - \hat{A})^2}{\sigma_A^2}, \quad (8) \end{aligned}$$

where N is the number of terms in the sum and $P(A)$ is the prior on A .

The posterior probability is maximized, by minimizing¹³ Equation (8) (the negative of the posterior probability) for each object in each of the five bands, so that for each object there are now ten variability parameters that can be used for classification: $A_u, \gamma_u, A_g, \gamma_g, A_r, \gamma_r, A_i, \gamma_i, A_z,$ and γ_z . Figure 3 shows an example for the g -band variability parameters; note that the different redshift ranges are well mixed (but are largely distinct from non-quasars) in this case. In practice, our implementation of the likelihood method is biased (10%–20% in the best-fit values) which becomes relevant when light curves are much better sampled than those discussed here. An approach such as that described in the appendices of Kozłowski et al. (2010) or Hernitschek et al. (2015) would be more robust. However, for the sake of this pilot investigation, our approach is more than sufficient, particularly because any bias in the variability parameters is the same for both selection by variability only, and by combined color and variability selection.

¹³ Using Scipy's Optimization package, Powell's method: `scipy.optimize.fmin_powell`

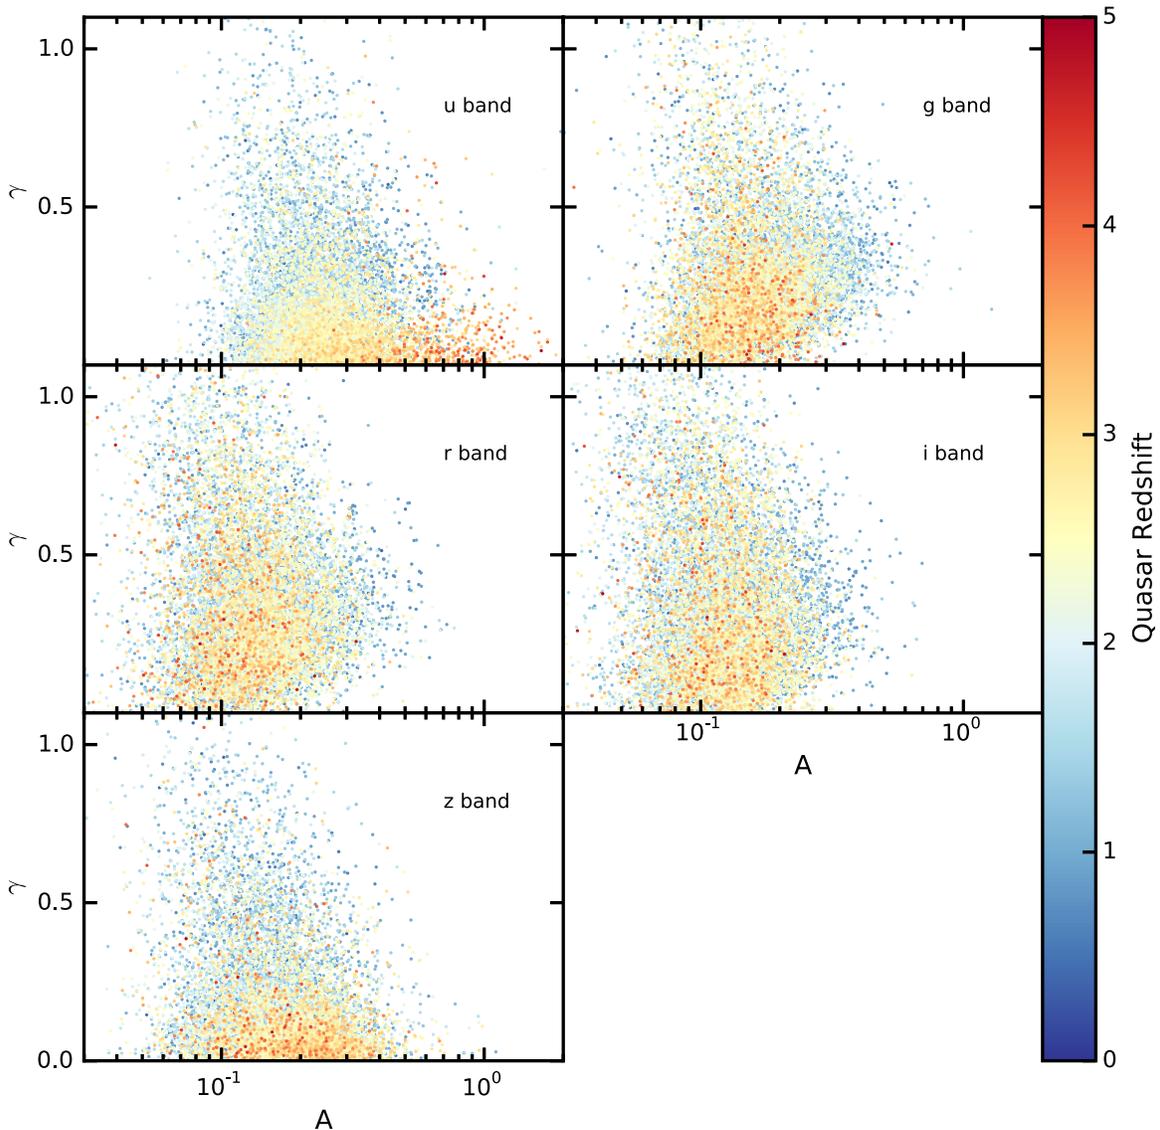


Figure 4. All spectroscopically confirmed quasars shown in A vs. γ space in each of the SDSS bands, colored by redshift. Shown to demonstrate the difficulty involved in combining the observations in all five bands to obtain one light curve and one structure function in order to describe the overall variability without previously knowing the object’s redshift. Note how the distribution of points shifts with band and with redshift. In particular, A and γ values agree well in the g , r , and i bands, but the large photometric errors in u and z bands artificially increase the apparent amplitude of the variability.

We currently fit the SF to the multi-epoch data for all bands separately to compare their performance in the NBC KDE selection algorithm (see Section 3). However, there are several ideas on how best to combine the observations in all five bands to obtain one light curve and one SF to describe the overall variability. These methods are complicated by differences in how quasars vary in the different bands. For example, different bands represent different distances in the accretion disk resulting in a time lag between the bands and different characteristic timescales.

As shown in Figure 4, there are different amplitudes of variability in different bands. Additionally, $\text{Ly}\alpha$ absorption obscures the true variability of quasars at high redshift. This is quite apparent in the u -band (top left panel) where the measured variability parameters for high- z quasars are caused by the high photometric errors of the u -band dropouts. It is also recognized that quasars become more luminous as they become bluer (Schmidt et al. 2010, 2012) and that bluer quasars in

general are more variable (Vanden Berk et al. 2004; MacLeod et al. 2010). Both of these effects must be taken into account when combining observations to describe the overall variability. A further complication for LSST will be the non-simultaneity of the observations in different bands. Thus, proper treatment of the combined variability data is complex and beyond the scope of this paper. For our purposes, describing the variability in each band is sufficient, and we therefore proceed with fitting the SF for each of the bands separately.

2.5. Test Set and Training Sets

Now that we have described the data inputs to our algorithm we can formally define the *test* and *training* sets. The test set begins with all stellar morphology (`objc_type == 6`) objects on the SDSS Stripe 82 with observations in DR7. Restricting our sample to point sources allows us to concentrate on the improvements gained by combining colors and variability

without having to worry about the differences in color and variability at redshifts and luminosities where the host galaxy contributes significantly to these properties. This set of observations was then limited by the following criteria: $-40^\circ < R. A. < 55^\circ$, $g - i < 6.0$, $g < 23.5$, $i < 22$, $\sigma_g < 0.5$, and $\sigma_i < 0.33$. These cuts are intended to reduce scatter due to high stellar density near the Galactic plane, high dust obscuration, and non-astrophysical colors. Observations with flags indicating poor photometry, such as those discussed in Section 3.2 of Richards et al. (2002) were also excluded. There are 1,163,174 objects with 49,274,136 observations that meet these cuts.

Only objects where we had sufficient observations to calculate variability parameters in all five bands and astrometric parameters in u and g were included in the test and training sets. Additionally, we require coadded colors $-1.0 < u - g < 9.0$, $-0.75 < g - r < 2.5$, $-0.5 < r - i < 3.0$, and $-1.5 < i - z < 1.75$, to constrain the parameter space for the NBC KDE to limit the necessary computational time for objects with unusually deviant colors. After these cuts, 916,587 objects remain. These objects compose the *cleaned data set*. The test set consists of the 903,366 sources from the cleaned data set that have not been spectroscopically identified as quasars.

The quasar training set is formed from the 13,221 spectroscopically confirmed quasars in the MQC that have matches in the cleaned data set. To keep computational time reasonable, we select a subsample of 72,680 non-matches for the non-quasar training set. As with our previous work (e.g., Richards et al. 2009b), we note that the vast majority of these non-quasar training set objects are not actually spectroscopically confirmed to be non-quasars and thus there will be some level of contamination as is discussed further in Section 3. We do not explicitly include or exclude spectroscopically confirmed stars or galaxies in the non-quasar training set as most of these were selected as quasars (and found to be contaminants) and are thus biased in their color-space distribution. In practice, when we run the classification on the test set we include the training set objects so that our catalog of candidate objects includes the known quasars, making it easier to determine our completeness of these sources.

3. NBC KDE ALGORITHM

Using training sets described in Section 2.5, classification of the test set objects (based on parameters described in Sections 2.3 and 2.4) was performed using NBC based on applying KDE to select quasars; see Richards et al. (2004), Gray et al. (2005), and Riegel et al. (2008). The algorithm takes training sets of objects divided into quasars and non-quasars. It creates an N-dimensional probability space for each of the classes, where N is the number of parameters that describe each type of object and the parameter space is normalized to give equal weight to each parameter (Gray et al. 2005). A probability density function (PDF) is constructed for each class of objects using KDE, by representing each individual object within a class by an N-dimensional Gaussian distribution and summing together the result for each object. Using the NBC KDE selection algorithm it is possible to combine all the classification parameters ($u - g$, $g - r$, $r - i$, $i - z$, A_u , γ_u , A_g , γ_g , A_r , γ_r , A_i , γ_i , A_z , and γ_z) and perform the classification simultaneously considering all the characteristics to determine if the object is a quasar or a non-quasar.

From this PDF, the probability of an unclassified object being a quasar or non-quasar can be calculated, but first we need an understanding of the real-world ratio of quasars to non-quasars. When a new point is placed in the PDF, the probability of it being a quasar or a non-quasar is weighted by its prior probability. This prior is an expectation of how many of the unknown objects are non-quasars. This weighting is an application of Bayes' Theorem:

$$P(M|D, I) = \frac{p(D|M, I)P(M|I)}{p(D|I)}. \quad (9)$$

In Equation (9), Bayes' Theorem (Bayes 1763; Ivezić et al. 2014), D stands for data, M for model, and I for prior information. This relates the posterior for the model based on the likelihood given the data and a prior. The pair of multi-dimensional weighted PDFs measures the probability of an unknown object being a quasar or a non-quasar, while taking into account the expected ratio of quasars to non-quasars, and classifies it accordingly. Throughout this work we use a prior of 0.95, meaning that we expect 95% of the objects to be non-quasars. The lower limit for the prior is determined by the fraction of known quasars in the test set. In Richards et al. (2009b) the ratio of quasar candidates to the test set was 2.6%. We use a slightly lower prior to capture some of the quasars that Richards et al. (2009b) did not. We assumed the prior to be independent of position on the sky and magnitude. Changing the prior by 1% does not change the number of quasar candidates by 1% of the test set, but changes the number by roughly 1% of the quasar candidates (Richards et al. 2015, submitted).

The algorithm requires a bandwidth for each of the training sets. The bandwidth controls the width of the kernel (a Gaussian distribution in our case) used to build the KDE. It is important to choose an optimal bandwidth when calculating the KDE or the distribution will be too smooth (under-fit) or will be too structured (over-fit)—in the same way as choosing an incorrect bin size for a histogram. The optimal bandwidth was found by performing leave-one-out cross-validation (leaving one object out and using the remainder of the training set to classify) over a range of bandwidths. We also refer to this as a *self test*.

This process was repeated to find the optimal bandwidth based on the product of completeness and efficiency. *Completeness* is defined as the number of known quasars correctly classified as quasars divided by the number of known quasars. It is also referred to as sensitivity. *Efficiency* is defined as the number of known quasars correctly classified as quasars divided by the number of objects (known quasars and non-quasars) classified as quasars. It is also referred to as purity. Different metrics could be chosen depending on the desired science and whether completeness is needed over efficiency, but we use the product of completeness and efficiency as a middle ground for this proof of concept. That is, an efficiency of 85% and a completeness of 70% is considered a better selection than efficiency of 99% and a completeness of 55%.

After an initial self-classification of the training set is done, all those objects in the non-quasar training set that were classified as quasars in the self test are removed. This process is expected to remove the majority of quasars that may have contaminated the non-quasar training set due to lack of prior spectroscopic confirmation. This new "cleaned" non-quasar training set is used for the final classification. This cleaning

process is a single iteration process and is performed separately for each of the classifications that we attempt below.

Having established the quasar prior probability, the quasar training set, a “cleaned” non-quasar training set, and the bandwidths for each of the training sets, we can proceed to classification of the unknown sources (i.e., the test set). Application of the NBC KDE algorithm results in each object receiving a binary quasar versus non-quasar classification, bifurcated at $P(M|D, I) = 0.5$. In the future, it may make more sense to simply output a probability for each object to facilitate combining this information with other data, but for the sake of this pilot study, we have chosen to make a hard cut (but in probability space rather than color space).

We explore which set of parameters (color, variability, or both) produces the best results in Section 4, then we will apply the algorithm to the test set to obtain a set of quasar candidates in Section 5.

4. TESTING CLASSIFICATION PARAMETERS

Our goal is to establish whether combining color and variability information in quasar selection is superior to using just colors or variability alone. To accomplish this goal, the NBC KDE algorithm was used in a series of *self tests*, which consists of performing leave-one-out cross-validation on the training sets (rather than on a test set). The object being classified is not included in the training set and the process is repeated for each object in the training sets. The classifications returned by the algorithm are compared to the known classifications of the objects to estimate the completeness and efficiency of selection using those particular input parameters.

Section 4.1 uses the NBC KDE algorithm with the above quasar and non-quasar training sets to perform a self test using colors alone. This process serves as our basis of comparison: do other parameters enable more robust quasar selection than colors alone? In Section 4.2, we attempt variability-only classification along with combined color and variability classification. We then compare the results of these self tests. This process reveals which variability (and color) parameters yield the most robust classification.

4.1. Classification Using Color

Our first self test was performed using only the *single-epoch* SDSS adjacent colors ($u - g$, $g - r$, $r - i$, $i - z$) as inputs to the algorithm. In practice, we chose a random epoch (meeting our requirements for good photometric and astrometric data) for each object. Using single epoch data is the most fair comparison for the majority of the objects in the SDSS footprint and we can use this as a control to compare how our method improves selection by adding variability. We could have chosen the “best” epoch for optimal classification by single-epoch colors alone; however, as we are testing the improvement from adding variability to the color classification, any epoch with quality data will serve.

The results of the classification are shown in Table 2, row 1, which indicates that these parameters are successful at not classifying non-quasars as quasars, at the expense of missing more than 37% of known quasars. Indicative of the well-known problem of separating high-redshift quasars from the locus of moderate-to-cool temperature stars (e.g., Richards et al. 2002), most of these missing quasars are at high redshift as can be seen from Figure 5. On the other hand, low-redshift quasars,

which can be selected robustly by traditional color cuts, are also easily identified using the NBC KDE algorithm as shown in Richards et al. (2004).

The completeness of our single-epoch selection is distinctly different from Richards et al. (2006): it is seemingly too high at low- z (given our restriction to point sources) and too low at high- z . For low- z this merely reflects the completeness of point sources. At high- z it is important to realize that in Richards et al. (2006) the purpose was to perform as complete a selection as possible, with efficiency as low as 50%, using hard color cuts. We will discuss how complete our selection is for all quasars, including extended sources, in Section 8.

In the SDSS Stripe 82 region, where we will conduct our experiments on variability selection of quasars, we are able to combine multiple epochs of imaging data to produce more accurate color measurements of the quasars (as discussed in Section 2.1). Thus, we perform a second self test using *coadded* colors for each object. Table 2, row 2 demonstrates that the use of coadded colors yields a small improvement in the efficiency of the sample, but a large improvement in the completeness—now being 93% complete. Figure 5 shows that most of this improvement comes from the recovery of high-redshift quasars; smaller photometric errors make it easier to distinguish the high-redshift quasar distribution from stars. However, there is still a dip at $z \sim 2.8$ where even the coadded colors do not enable better than 75% completeness.

4.2. Choosing Optimal Classification Parameters

Variability alone can be the basis for a robust quasar classification (e.g., Schmidt et al. 2010; Butler & Bloom 2011; MacLeod et al. 2011), so we next perform a self test by applying KDE to the pair of variability parameters for each band (as defined in Section 2.4) and then on combinations of variability parameters from the multiple bands. The results are shown in Table 2 and Figure 5. It is interesting to compare the performance of the bands because each represents different distances from the center of the accretion disk, different characteristic timescales, and different (redshift-dependent) peak amplitudes.

Particularly important is that variability selection has a higher completeness in the range $2.6 < z < 3.0$ than do colors. There are no significant trends with redshift in the $A-\gamma$ space in the g , r , and i bands, so the quasars can be separated out from the non-quasars in the variability space without completeness issues at specific redshifts (unlike the dramatic drops seen for color-only selection). The completeness drops off gradually with higher redshift, which is a result of changes in observed magnitude, signal-to-noise ratio, and timescale of variability in the observer’s frame. Combining g and r , r and i , and g , r , and i , we find similar trends as using just the variability parameters from a single band, with marginally higher completeness (and efficiency) at all redshifts.

Selection by u - and z -band variability performs much worse than both coadded and single epoch colors. The u band is strongly influenced by Ly α forest absorption of the (variable) quasar continuum at high redshift, thus suppressing the signal-to-noise ratio. This results in discordant variability parameters for quasars that are quite apparent in Figure 4. The lower performance of the z -band is likely due to the lower signal-to-noise ratio of the photometry and thus the larger scatter of the variability parameters as seen in Figure 4. These discrepant

Table 2
NBC KDE Results—Self Test Non-quasar and Quasar Fraction

Self Test	Non-quasars As Non-quasars			Quasars As Quasars		
	Correct	Total	Fraction	Correct	Total	Fraction
Single epoch colors	68611	69566	0.986	8232	13221	0.623
Coadded colors	69474	69738	0.996	12353	13221	0.934
<i>u</i> variability	70970	71936	0.987	5550	13221	0.420
<i>g</i> variability	69489	70040	0.992	11138	13221	0.842
<i>r</i> variability	69998	70476	0.993	11137	13221	0.842
<i>i</i> variability	69935	70397	0.993	10782	13221	0.816
<i>z</i> variability	70665	71372	0.990	5403	13221	0.409
<i>g</i> and <i>r</i> variability	69777	70054	0.996	12060	13221	0.912
<i>r</i> and <i>i</i> variability	69714	70050	0.995	11933	13221	0.903
<i>g</i> , <i>r</i> , and <i>i</i> variability	69728	70034	0.996	12150	13221	0.919
Coadded colors; <i>u</i> variability	69644	70077	0.994	12311	13221	0.931
Coadded colors; <i>g</i> variability	69822	70114	0.996	12739	13221	0.964
Coadded colors; <i>r</i> variability	69912	70229	0.996	12741	13221	0.964
Coadded colors; <i>i</i> variability	69880	70157	0.996	12634	13221	0.956
Coadded colors; <i>z</i> variability	69682	69990	0.996	12359	13221	0.935
Coadded colors; <i>g</i> and <i>r</i> variability	69663	70081	0.994	12816	13221	0.969
Coadded colors; <i>r</i> and <i>i</i> variability	69658	70096	0.994	12800	13221	0.968
Coadded colors; <i>g</i> , <i>r</i> , and <i>i</i> variability	69948	70108	0.998	12626	13221	0.955

Note. Fraction of non-quasars correctly classified as non-quasars and quasars correctly classified as quasars from the leave-one-out cross-validation of the training sets. The non-quasar total is different in the different rows because the non-quasar training set is “cleaned” before it is used for the final classification, as described in Section 3. The bandwidths are chosen to optimize the product of completeness and efficiency.

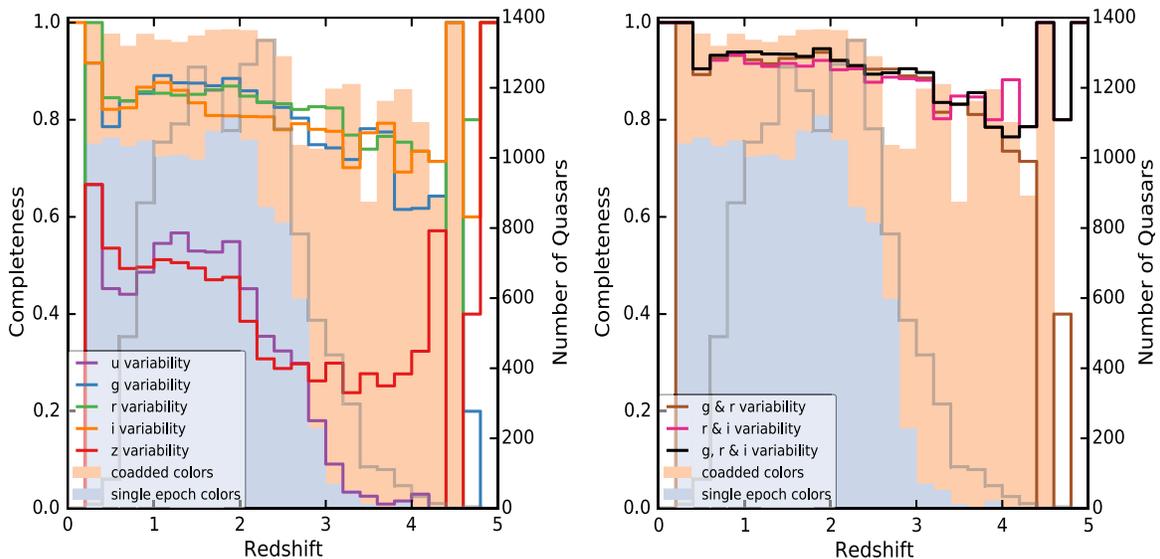


Figure 5. Fraction of quasars correctly classified as quasars (completeness). These panels demonstrate that we are able to separate the quasars from the non-quasars in the variability space without extreme changes in completeness at specific redshifts. In both panels the gray line shows the number of quasars in each bin (right axis) and light blue (single epoch) and peach (coadded epochs) histograms show the completeness of color-only selection (left axis, Section 4.1). Note the catastrophic loss of high- z quasars from single-epoch colors and the incompleteness at $z \sim 2.8$ even for coadded colors. We also show classification from variability only: single bands (left panel) and combinations of multiple bands (right panel). The g , r , and i bands are shown as blue, green, and orange lines, respectively. There are no dramatic drops in the g -, r -, or i -bands variability at distinct redshifts, just a gradual decline with increasing redshift, which is related to observed magnitude, signal to noise ratio, and timescale of variability in the observer’s frame. The overall completeness using variability alone is not as high as coadded colors alone at low redshifts, but is more successful than single-epoch colors alone at high redshifts.

values increase the probability of high-redshift quasars being classified as stars.

While variability selection produces more consistent results with redshift than color selection, we find that, at many redshifts, color selection is still superior. We thus consider coadded colors with combinations of variability parameters from single and multiple bands. The results are shown in Table 2 and Figure 6. Adding variability parameters from just

one band significantly improves the selection, especially the high signal-to-noise ratio bands g , r , and i . The addition of the u - and z -band variability to colors still fails at $z \sim 2.8$ because the variability signal is not strong enough (as demonstrated in Figures 4 and 5) to overcome color selection bias.

We graphically summarize the results of the self tests in Figure 7. Quasar completeness as a function of redshift is shown in the left panel, quasar completeness as a function of

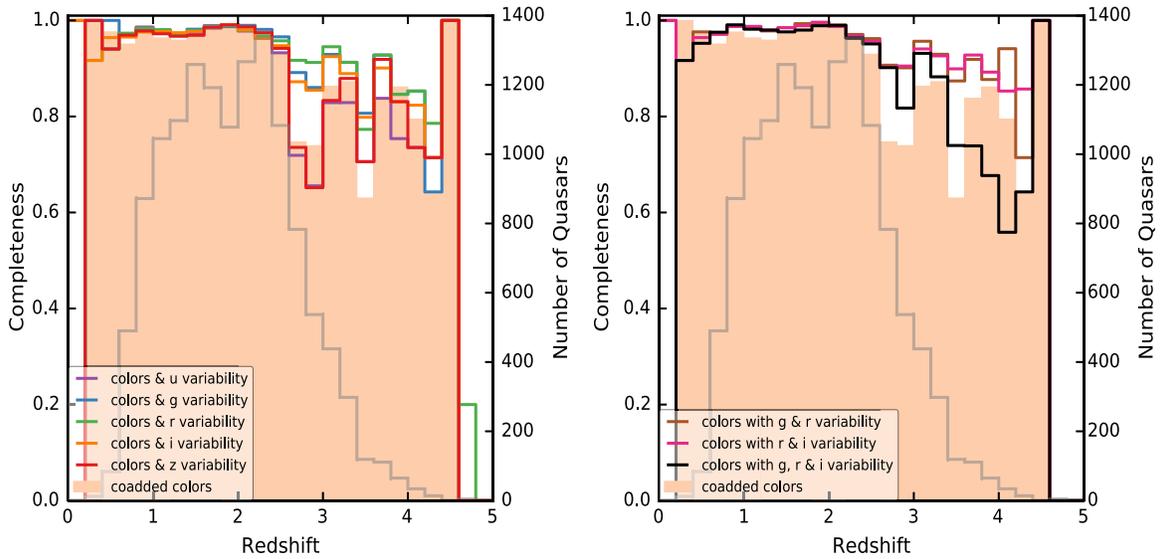


Figure 6. Fraction of quasars correctly classified as quasars using coadded colors *and* variability, as a function of redshift. Notice the improved completeness near redshifts 2.7 and 3.5, where the quasars and non-quasars overlap in color space, with the addition of variability parameters. Shown are single bands of variability combined with coadded colors (left panel) and combinations of multiple bands of variability combined with coadded colors (right panel). In both panels the gray line shows the number of quasars in each bin (right axis).

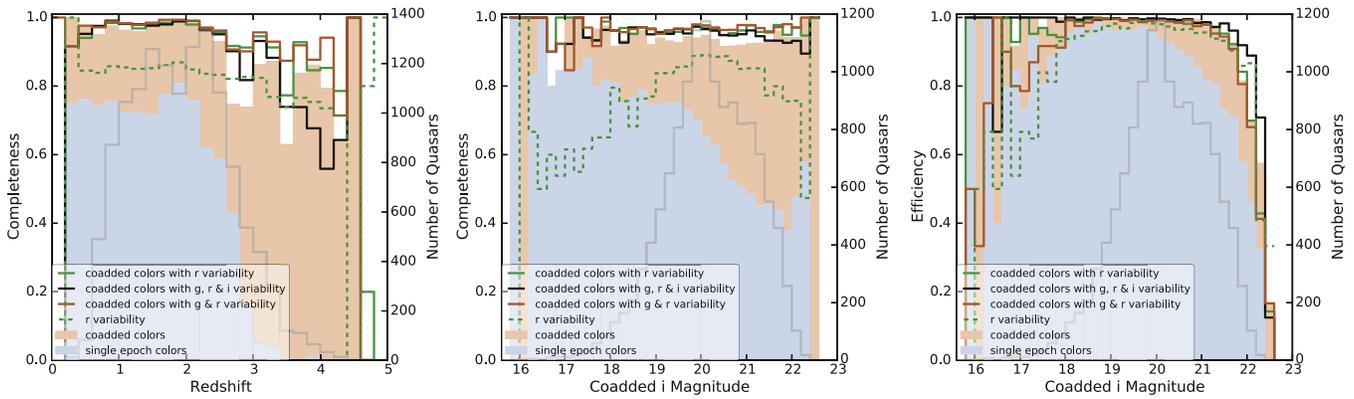


Figure 7. Comparison of self tests using different combinations of color and variability. These panels demonstrate that the combination of color and variability gives the best results for completeness and efficiency as a function of redshift and magnitude with more details in the text. Shown are the completeness (known quasars classified as quasars divided by known quasars) as a function of redshift (left panel), completeness as a function of coadded *i*-band magnitude (*center panel*), and efficiency (known quasars classified as quasars divided all objects classified as quasars) as a function of coadded *i*-band magnitude (right panel). The gray line shows the number of quasars in each bin (right axis).

i magnitude in the center panel, and quasar efficiency as a function of *i* magnitude in the right panel. For colors alone, both coadded and single epoch, there are regions of color space where the quasar training set and non-quasar training set overlap, resulting in redshift regions with poor completeness. Variability alone, as demonstrated by the *r*-band selection, does not have these redshift trends, but has a lower efficiency than coadded colors at all other redshifts. The addition of coadded colors to the *r*-band variability information helps to improve upon the colors alone at all redshifts, but in particular in the dips at $z \sim 2.7$ and $z \sim 3.5$. Using coadded colors together with variability in multiple bands improves the classification even further (e.g., compare the solid green lines to the dotted green lines). The left panel of Figure 7 shows that adding the *i*-band variability makes things worse (possibly because the *i*-band has a lower signal-to-noise ratio than *g* or *r* given that quasars generally have blue spectral energy distributions; SEDs), but note that there are relatively few high-redshift objects and the middle

panel shows that the loss of completeness is coming from very faint objects. Moreover, the right panel shows that adding the *i*-band variability improves the efficiency. Table 3 shows that while adding the *i*-band variability reduces the completeness by 1%, it compensates by increasing the efficiency by 2%.

These self tests of the quasar and non-quasar training sets validate our hypothesis that the most successful option is a combination of coadded colors and variability. No combination of colors and variability was highest in both completeness and efficiency; however, the combination of coadded colors and both *g* and *r* variability parameters give the most robust selection with a combined product of completeness and efficiency of 93.88% (see Table 3) and was consistent in completeness across all redshift values (see Figure 6). As such, for our analysis of the test set in the next section, we have adopted coadded colors with both *g* and *r* variability parameters as our basis set.

Table 3
NBC KDE Results: Self Test Completeness and Efficiency

Self Test	Variability Only		Single Epoch Colors w/Variability		Coadded Colors w/Variability	
	Completeness	Efficiency	Completeness	Efficiency	Completeness	Efficiency
Color only	0.6226	0.8960	0.9343	0.9791
u variability	0.4198	0.8517	0.6934	0.9289	0.9312	0.9660
g variability	0.8424	0.9529	0.8372	0.9149	0.9635	0.9776
r variability	0.8424	0.9588	0.8583	0.9165	0.9637	0.9757
i variability	0.8155	0.9589	0.8126	0.9235	0.9556	0.9785
z variability	0.4087	0.8843	0.7158	0.9214	0.9348	0.9757
g and r variability	0.9122	0.9775	0.8115	0.9758	0.9694	0.9684
r and i variability	0.9026	0.9726	0.8076	0.9734	0.9682	0.9669
g , r , and i variability	0.9190	0.9754	0.8573	0.9761	0.9550	0.9875

Note. Completeness (known quasars classified as quasars divided by known quasars) and efficiency (known quasars classified as quasars divided all objects classified as quasars) for each of the self tests described in Section 4.2. This indicates that the most successful option is a combination of coadded colors and variability, but no particular variability bands stood out when in combination with colors.

5. BUILDING A QUASAR CANDIDATE CATALOG

Now that the most efficient set of parameters are chosen, in Section 5.1 the algorithm is applied to the test set using the full quasar training set. Finally, in Section 5.2 we test a process where the algorithm is used to perform simultaneous classification and redshift estimation. Specifically, the test set is classified using a series of quasar training sets that only contains quasars from limited redshift ranges.

5.1. Classifying the Test Set

In the previous section we identified coadded colors combined with both g and r variability as producing the best classification for the training set objects. We now apply the selection to the test set. The NBC KDE algorithm was used to perform an 8D classification ($u - g, g - r, r - i, i - z, A_g, \gamma_g, A_r,$ and γ_r), using the same bandwidths used during the self tests and an identical prior. The objects identified as quasar candidates, with $P(Q|d) > 0.5$, are listed in the catalog (available online) which is described in more detail in Section 7.

The results of the classification are shown in Figure 8. We will discuss the new candidate quasars, their characteristics, and contaminants in Sections 7 and 8. In general, the candidate quasars (green contours) closely mirror the distribution of the known quasars (orange contours) and extend slightly beyond in the parameter space. The incorrectly classified quasars lie in the area where quasars and non-quasars overlap in color and variability space. When comparing to the quasar distribution as a function of redshift shown in Figure 1, the candidate quasars extend beyond the known quasars into mid-redshift and high-redshift regions of color space. The candidate quasars have a higher density in the areas overlapping the non-quasars (gray contours), than the known quasars. This could be caused by the variability parameters selecting quasars that were missed by color selection because they are hidden in the stellar locus, or stellar contaminants in our selection. There are also some new candidates in the bluest corner of $g - r$ versus $r - i$ color space which are likely white dwarf contaminants that we will attempt to purge in Section 7.

5.2. Classification using Redshift Bins

Quasar colors depend on redshift as shown in Figure 1. As such, it is possible to identify quasars while simultaneously

estimating their redshifts (e.g., Suchkov et al. 2005; Bovy et al. 2012). We test the extension of our method in a similar manner simply by limiting the quasar training set to a narrow redshift region. By doing so, we are able to select quasars with colors similar to other quasars of that redshift, thereby simultaneously providing a rough estimate of the redshift.

To accomplish this, the full quasar training set (see Section 2.5) was divided into 18 separate training sets by redshift: non-overlapping redshift bins from 0.4 to 4.0 with a bin width of 0.2. The quasars outside each redshift bin were added to the *non-quasar* training set. A handful of quasars that were significant outliers (5σ) from the modal color in each bin were removed from the quasar training set. These outliers could be caused by errors in the photometry and/or heavy dust reddening. Including them caused us to find objects with those colors that are not really quasars or are quasars at a different redshift.

As above, a self test was performed on the training sets for each redshift bin to find the optimal bandwidths. Specifically, the redshift-bin training sets were used to classify the full quasar training set (13,221 quasars spanning the full redshift range). The results of these self tests are shown in Table 4 and Figures 9 and 10. These show that the completeness of quasar classification (both identifying known quasars as quasars *and* also as being in the correct redshift bin) is generally better than 75%. The contamination (here quasars from the wrong redshift bin being selected) is typically less than 10%.

Of the 13,221 training set quasars, 12,535 were classified in at least one bin (94.8% overall completeness). These objects are shown as a density plot in Figure 10 in $\Delta z = 0.2$ photometric redshift bins. The regions of misclassification at spectroscopic redshifts ~ 0.75 and ~ 2.1 stem from degeneracies in color-redshift space.

With the self test completed, we finally classify the test set described in Section 2.5, the same that was classified in Section 5.1. For each of the non-overlapping redshift bins from 0.4 to 4.0, each object in the test set is returned as either a quasar candidate or a non-quasar candidate. If it is found to be a quasar candidate, we calculate the quasar probability (in addition to the initial binary classification). Many objects were found to be quasar candidates in several bins and the classification probability in each bin was calculated. Results of the classification are given in Table 5; Figure 11 shows the results of the classification in color and variability parameter

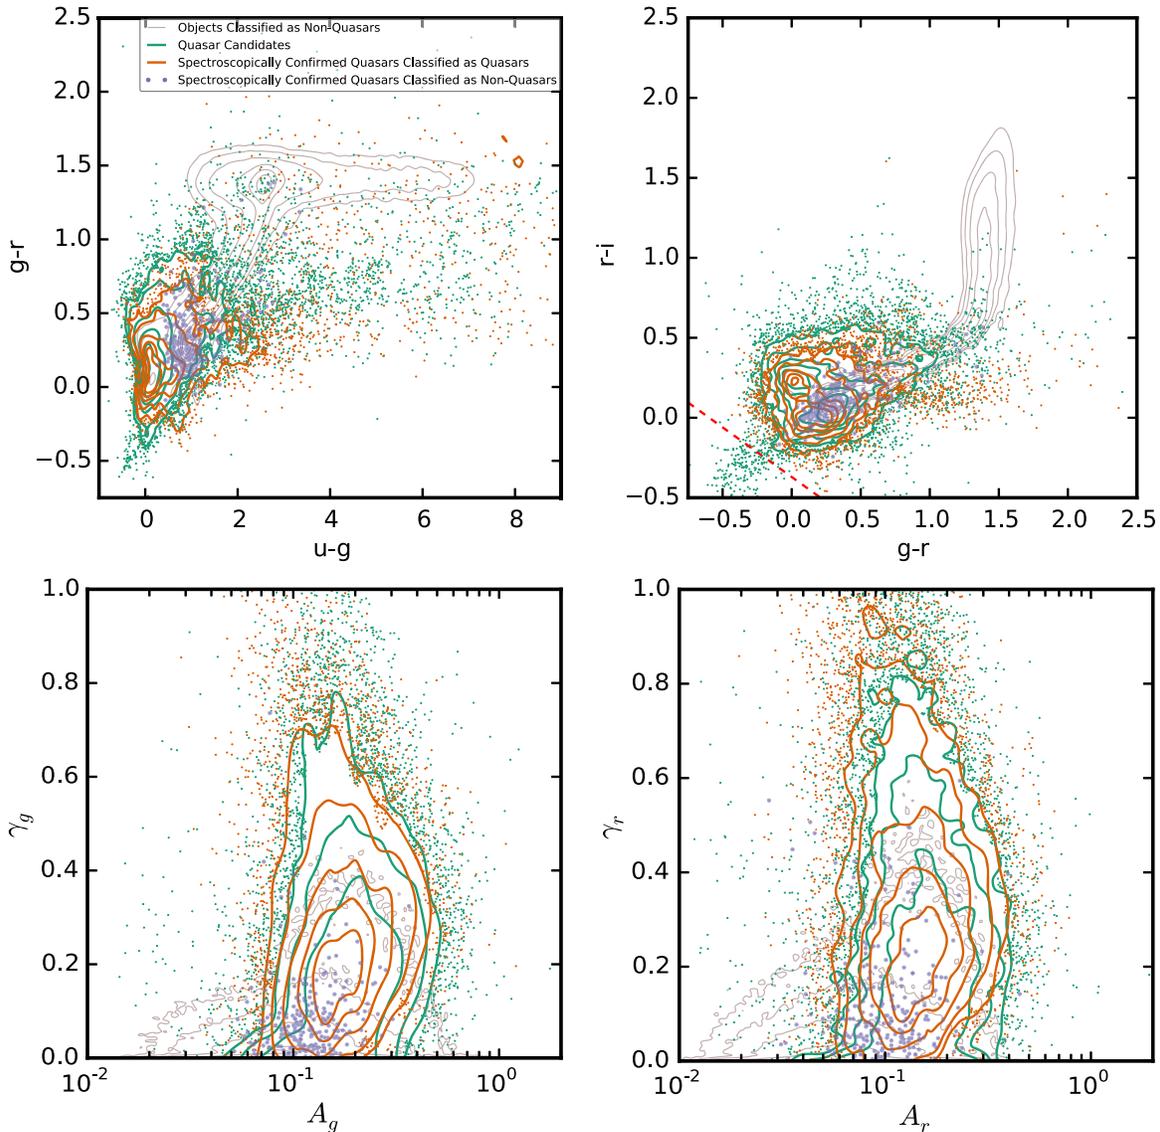


Figure 8. Color and variability parameter space plots showing the results of test set classification using a single quasar training set covering the full quasar redshift range (Section 5.1). These panels demonstrate that the incorrectly classified quasars lie in the area where quasars and non-quasars overlap in color and variability space and that the candidate quasars closely mirror the distribution of the known quasars and extend slightly beyond in the parameter space (including a region known to be inhabited by white dwarfs in the blue corner of the upper right panel). *Colors left panel:* $u - g$ color vs. $g - r$, *colors right panel:* $g - r$ vs. $r - i$, *variability left panel:* A_g vs. γ_g , and *variability right panel:* A_r vs. γ_r . Objects in the test set classified as non-quasars are shown as gray contours. Levels for contours in Figures 8 and 11: gray: colors—95%, 90%, 80%, 60%, 40%, 20%, variability—98%, 95%, 90%, 80%; green: colors—90%, 80%, 60%, 40%, 20%, variability—90%, 80%, 60%; orange: 90%, 80%, 60%, 40%, 20%, quasar candidates that are not spectroscopically identified are shown as green contours and scatter points for outliers, spectroscopically identified quasars classified as quasars are shown as orange contours and scatter points for outliers, and spectroscopically identified quasars incorrectly classified as non-quasars are shown as purple dots. The red dashed line in the upper right panel is the white dwarf cut described in Equation (12).

space, as in Figure 8. We discuss the difference in this selection and the selection in Section 5.1 in Section 7. An analysis of the quasar candidates is performed in Section 8.

6. REDSHIFT ESTIMATION

In this section we will improve on the accurate, but not precise, redshift estimation of Section 5.2 and compute photometric redshifts for the quasar candidates. First, we will describe the astrometric information (Section 6.1) and near-infrared colors (Section 6.2), that will be used in addition to optical colors (Section 2.3). We combine these inputs to calculate photometric redshifts using the method described in Weinstein et al. (2004). We compare the robustness of our different redshift estimates in Section 6.3.

6.1. Astrometry

In addition to colors, our analysis will make use of *astrometric* measurements of quasars (Kaczmarczik et al. 2009). Light rays from extraterrestrial sources are bent according to Snell’s law as they enter the Earth’s atmosphere from the vacuum of space. A celestial source observed from the Earth will appear higher in the sky than it actually is, unless it is at the zenith. The amount of this deflection depends on the index of refraction in the air and the photon’s angle of incidence. Since the index of refraction of air is a function of wavelength, shorter wavelength photons are bent more than longer wavelength photons. This effect is known as differential chromatic refraction (DCR).

Table 4
NBC KDE Results: Test Set Classification of Spectroscopically Confirmed Quasars

Redshift Bin	Number Inside Redshift Bin			Number Outside Redshift Bin		
	Correct	Total	Fraction	Correct	Total	Fraction
$0.4 < z \leq 0.6$	67	84	0.798	12788	13137	0.973
$0.6 < z \leq 0.8$	368	494	0.745	11855	12727	0.932
$0.8 < z \leq 1.0$	662	870	0.761	11704	12351	0.948
$1.0 < z \leq 1.2$	891	1043	0.854	11368	12178	0.934
$1.2 < z \leq 1.4$	949	1097	0.865	11307	12124	0.933
$1.4 < z \leq 1.6$	1100	1262	0.872	11147	11959	0.932
$1.6 < z \leq 1.8$	1085	1191	0.911	10766	12030	0.895
$1.8 < z \leq 2.0$	851	1078	0.790	11343	12143	0.934
$2.0 < z \leq 2.2$	1036	1278	0.811	11150	11943	0.934
$2.2 < z \leq 2.4$	1151	1322	0.871	10349	11899	0.870
$2.4 < z \leq 2.6$	996	1084	0.919	10572	12137	0.871
$2.6 < z \leq 2.8$	535	782	0.684	11866	12439	0.954
$2.8 < z \leq 3.0$	469	540	0.869	12093	12681	0.954
$3.0 < z \leq 3.2$	340	435	0.782	12377	12786	0.968
$3.2 < z \leq 3.4$	223	298	0.748	12587	12923	0.974
$3.4 < z \leq 3.6$	103	119	0.866	12933	13102	0.987
$3.6 < z \leq 3.8$	107	111	0.964	12966	13110	0.989
$3.8 < z \leq 4.0$	61	65	0.939	13026	13156	0.990

Note. Fraction of quasars inside the redshift bin correctly classified as inside the redshift bin and quasars outside the redshift bin correctly classified as outside the redshift bin from the leave-one-out cross-validation of the training sets, using the training sets divided into redshift bins.

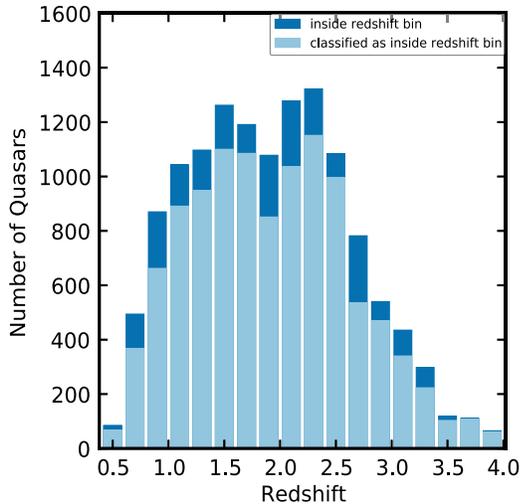


Figure 9. Classification of a test set of quasars with known spectroscopic redshifts, using the training sets divided into redshift bins. Dark blue indicates all quasars in that bin, light blue indicates quasars classified with the correct redshift. The ratio of the two is the completeness of quasars inside the redshift bin.

The automated corrections for the DCR effect to the SDSS astrometry are computed as a function of a broad-band flux ratio. The DCR for any given object depends on the effective wavelength of the bandpass (the convolution of the object’s SED and the filter transmission curve) of the object within a given bandpass, which in turn depends upon the filter’s transmission properties and on the distribution of the source’s flux within the bandpass. A pure power-law (without emission lines) changes the effective wavelength in a correctable way, but the DCR corrections become anomalous when there are emission lines. For example, adding an emission line on the blue side of the filter makes the effective wavelength bluer, while adding an emission line on the red side makes the

effective wavelength redder. For emission line objects (like quasars), the effective wavelength can be very different from the assumed power law, changing by as much as 150 \AA in the u -band (Kaczmarczik et al. 2009). The difference between the expected and observed astrometric displacements due to DCR enables the distinction of quasars and non-quasars in addition to providing an additional source of information about the redshift of the object. We examine the differential DCR offset (along the parallactic angle; Filippenko 1982) in the u -band ($auPar$) and in the g -band ($agPar$); the effect is too small to measure in r , i , and z given the astrometric errors of our data and the smaller DCR at longer wavelengths.

Kaczmarczik et al. (2009) reduced the statistical error in the astrometric offsets of individual objects by normalizing the DCR offsets at multiple epochs (each with different airmass) to some fiducial airmass. Here we take a different approach that we find to be more robust. To first order, differential refraction is linear in $\tan(Z)$, where Z is the zenith angle, with zero intercept (no DCR at airmass of one at the zenith). Thus, a plot of multiple epochs of noisy quasar DCR measurements should cluster around a line with a fixed slope (for a given bandpass and object redshift) with zero intercept.

In a manner similar to our SF fitting above, we use minimization of a log likelihood function to calculate the astrometric parameters in the u and g band. We fit the data with a straight line that runs through the origin and parameterize the DCR simply by the slope of the line. The light curve is cleaned of outliers in the same way as was done for the variability parameter calculation. We require at least 10 good observations in each band and at least one observation with airmass in the r band greater than 1.5, which is $\tan(Z) \sim 1.1$ —contrary to the variability analysis above since here higher airmass means a larger DCR signal despite greater extinction. We weight each observation by the r -band airmass since higher airmass observations are more rare and should have greater discriminatory power. Further work could be performed in the future to determine if this weighting scheme is indeed optimal.

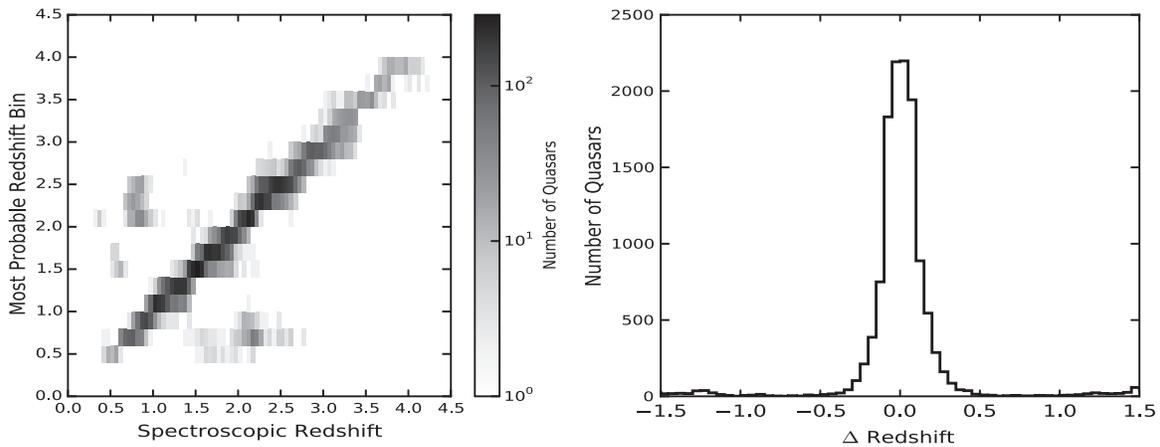


Figure 10. Comparison of spectroscopic redshift to the bin into which known quasars were classified with the highest probability. Left panel: spectroscopic redshift vs. the most probable redshift bin. Right panel: histogram of Δz (the most probable redshift bin minus the spectroscopic redshift). Only 5.6% of the quasars have $|\Delta z| > 0.5$.

Table 5
NBC KDE Results: Test Set Classification with Redshift Bins

Redshift Bin	QSO Candidates		Known QSOs Returned				
	All	qso_prob > 0.8	Known QSOs	Returned	Fraction	qso_prob > 0.8	Fraction
0.4 < z ≤ 0.6	2925	380	84	67	0.798	46	0.548
0.6 < z ≤ 0.8	3433	801	494	367	0.743	293	0.593
0.8 < z ≤ 1.0	3590	767	870	671	0.771	332	0.382
1.0 < z ≤ 1.2	4775	1920	1043	883	0.847	567	0.544
1.2 < z ≤ 1.4	6238	2981	1097	945	0.861	656	0.598
1.4 < z ≤ 1.6	5543	2237	1262	1097	0.869	754	0.598
1.6 < z ≤ 1.8	7838	3516	1191	1083	0.909	740	0.621
1.8 < z ≤ 2.0	5931	2585	1078	840	0.779	574	0.533
2.0 < z ≤ 2.2	5195	1948	1278	1034	0.809	582	0.455
2.2 < z ≤ 2.4	4162	2354	1322	1146	0.867	895	0.677
2.4 < z ≤ 2.6	4540	2477	1084	993	0.916	832	0.768
2.6 < z ≤ 2.8	3023	1028	782	524	0.670	327	0.418
2.8 < z ≤ 3.0	2246	1295	540	465	0.861	410	0.759
3.0 < z ≤ 3.2	1390	753	435	334	0.768	260	0.598
3.2 < z ≤ 3.4	1228	644	298	223	0.748	181	0.607
3.4 < z ≤ 3.6	1122	671	119	102	0.857	99	0.832
3.6 < z ≤ 3.8	596	399	111	106	0.955	106	0.955
3.8 < z ≤ 4.0	514	348	65	60	0.923	58	0.892
Total	32108	20962	13153	10940	0.831	7712	0.586

Note. Classification of the full test set of objects, using the training sets divided into redshift bins. Total will not be a sum of the above rows because many objects were classified in multiple bins.

Figure 12 shows an example of this process for a single quasar with the u -band data in blue and the g -band data in green. These astrometric data can be used to constrain photometric redshifts for quasars in surveys where there are many observations and/or observations at high airmass that can provide constraints on the DCR slope. See Figure 7 of Kaczmarszik et al. (2009). We will use the astrometric parameters $auPar$ and $agPar$ in Section 6.3 when calculating the photometric redshifts of the quasar candidates.

In the left panel of Figure 13 we plot all of the empirical DCR slopes for the quasar training set. The right panel of Figure 13 shows that non-quasars and quasars have somewhat different signals in this parameter space. We have only included point sources in this analysis, but the process should work for normal star-forming galaxies too, as the 4000 Å break

can produce significant astrometric shifts relative to the SED model assumed in the astrometric solution. In this pilot investigation, we have not used the DCR effect for classification; however, the information provided by DCR would add yet another piece of information that could be used to refine the classification probabilities of the objects in the test set. For example, objects with large negative values of $auPar$ are (empirically) more likely to be non-quasars than quasars.

6.2. VISTA Hemisphere Survey

While we select objects only using optical imaging data, we can make use of near-IR (NIR) photometry to improve our photometric redshift estimates. The VISTA Hemisphere Survey (VHS) is a near-infrared survey with coverage in the southern

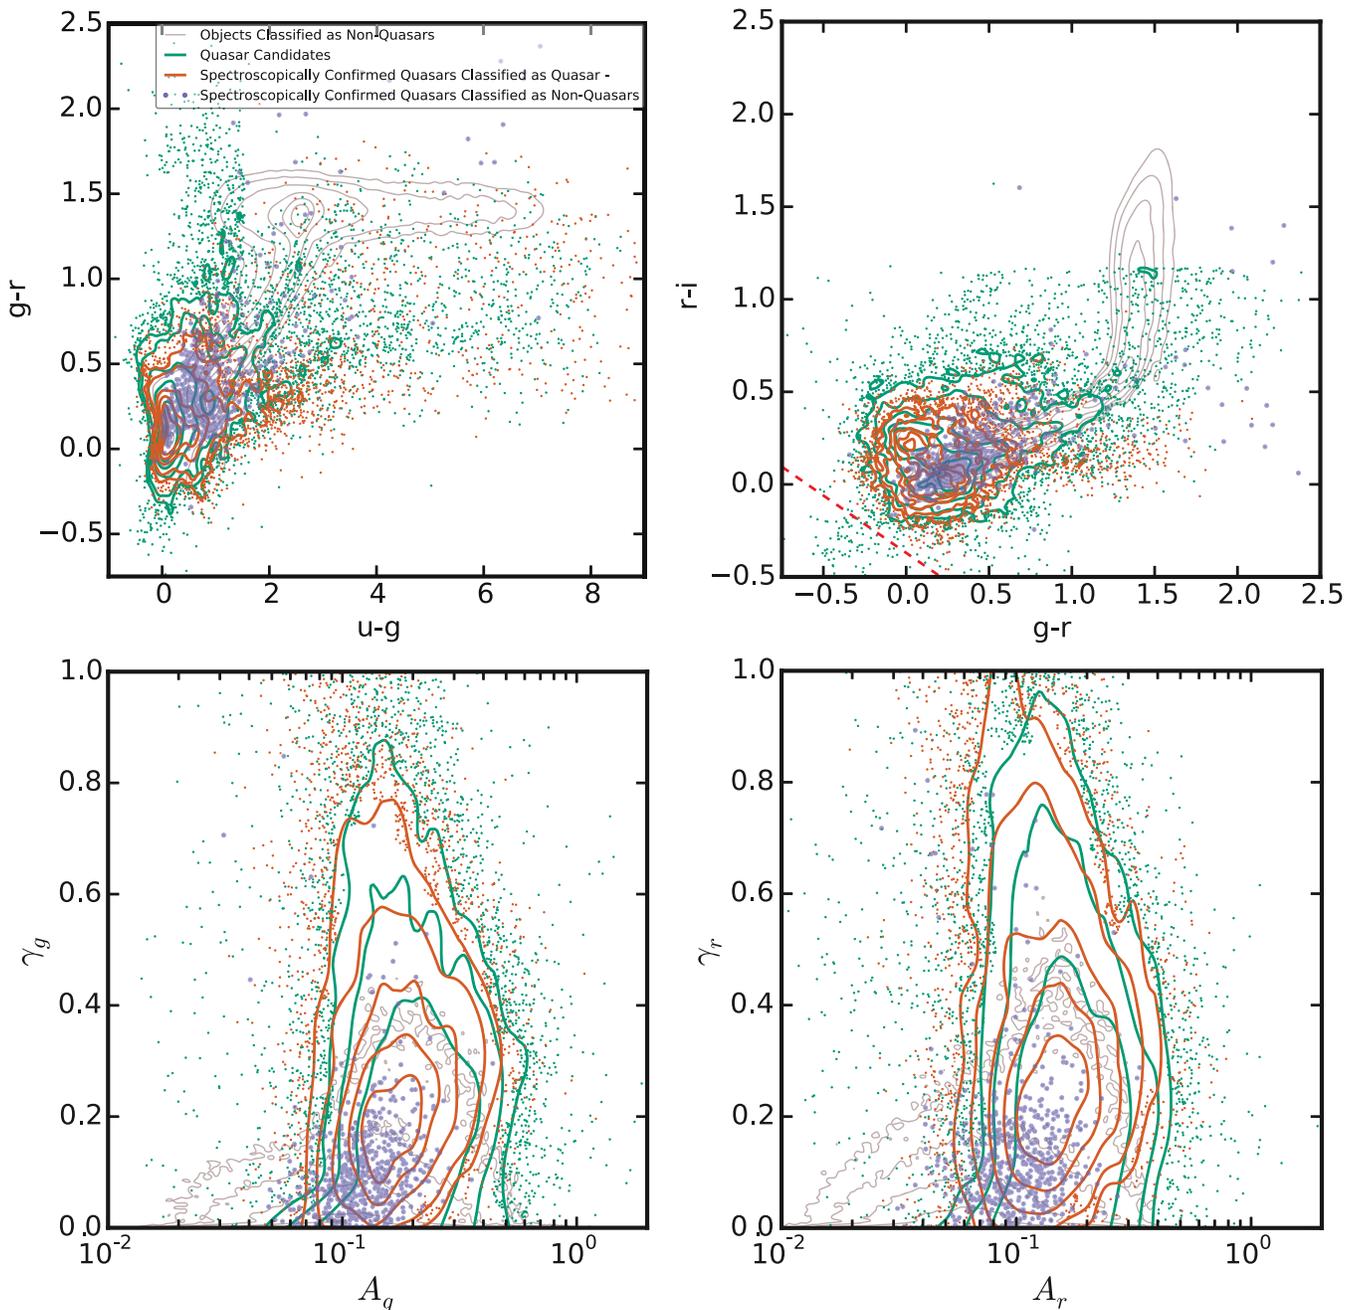


Figure 11. Same as in Figure 8 (color and variability space plots showing the results of test set classification), but using redshift bins (described in Section 5.2). In the bottom panels, we find that the selection in variability parameter space shows no noticeable difference to Figure 8, which is not surprising as A_g vs. γ_g and A_r vs. γ_r have no strong redshift trends. However, there are slight differences in color space (top panels). We discuss these further in Section 7.

hemisphere, including the full Stripe 82 footprint. The second VHS public data release (VHSDR2) was made available on the VISTA Science Archive (VSA)¹⁴ in 2014 April. These data include three bands J , H , and K_s , with (Vega) magnitude limits of $J = 20.2$, $H = 19.3$, and $K_s = 18.2$ (McMahon et al. 2013). Using the Rayleigh criteria, the surveys were matched at $1''0$ (Parejko et al. 2008): 48% of the quasar candidates had matches in all three bands. It would be beneficial to calculate photo- z estimates for the remaining non-detections to put constraints on the quasar SED, but that is beyond the scope of this work.

¹⁴ <http://horus.roe.ac.uk/vsa/index.html>

6.3. Photometric/Astrometric Redshifts

Empirical photometric redshifts (Richards et al. 2001) were calculated for all of the objects that were found to be potential quasars in Sections 5.1 or 5.2. The algorithm is described in detail in Weinstein et al. (2004) and essentially involves least-squares fitting (without error weighting) between the candidate quasar colors and the mean (sigma clipped) colors of quasars as a function of redshift. The covariance matrix used in the process was calculated using the quasars with known spectroscopically determined redshifts. The quasars are binned by redshift in bins of width 0.02. The mean color-vector and the color covariance matrix is found for the quasars in each redshift

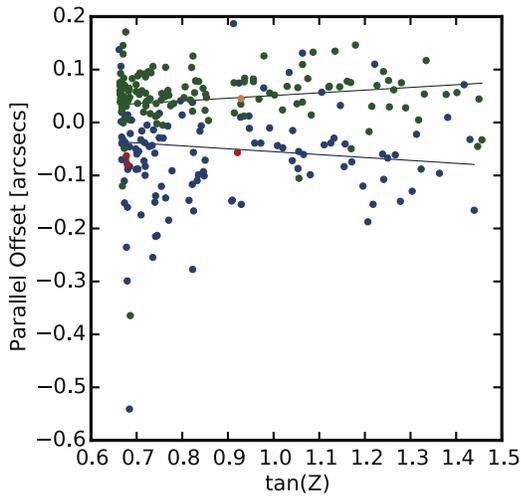


Figure 12. Measured astrometric offset along the parallactic angle as a function of $\tan(Z)$. Shown is SDSS J013417.81–005036.2, a redshift 2.26 quasar from SDSS Stripe 82, the same object shown in Figure 2. This quasar is shown as an example representative of the data set. Each point refers to a different observation of this object, at a different airmass. The astrometric accuracy is ~ 0.03 arcsecs for $g < 20.0$, but up to 0.1 arcsecs for $g \sim 22.0$ (Pier et al. 2003). u -band observations are shown in blue, with those points that were outliers removed from the light curve in Figure 2 are shown in red. g -band observations are shown in green, with outliers removed from the light curve shown in orange. The fits, shown as solid blue and green lines, have a y -axis intercept of zero. For this quasar, the slope of the line (offset along the parallactic angle) in the u -band ($auPar$) is -0.055 and g -band ($agPar$) is 0.105. The astrometric redshift is found to be 2.57.

bin; see Figure 4 of Richards et al. (2015, submitted). For each of the quasar candidates, we calculate how “far” its colors are from these calculated mean colors and convert this information into a probability distribution as a function of redshift bin, as shown in Equation (5) of Weinstein et al. (2004). The peak of the probability distribution is reported as the photometric redshift and the confidence is calculated by integrating under the curve down to a threshold. A few examples of photometric redshift PDFs are shown in Figure 14.

First, the photometric redshift was calculated using SDSS adjacent colors ($u - g$, $g - r$, $r - i$, $i - z$). The mean colors were calculated using all MQC objects with known spectroscopic redshifts (i.e., not just the Stripe 82 quasars) using coadded photometry when available. We did this to improve the constraints on the photometry for high-redshift quasars. Those objects without coadded photometry have larger photometric errors, but the increase in the number of objects overcomes the noise. The color-based photo- z PDF of 4 representative objects is shown in green in Figure 14. The 13,419 quasars on Stripe 82 with spectroscopic redshifts are shown in Figure 15 (top left panel). Of these objects, 5843 (43.5%) have a calculated photometric redshift within 0.1 of the spectroscopic redshift and 10,201 (76.0%) are within 0.3, as seen in Figure 16. The quasars around redshift 0.8 and 2.2 have particularly poor photometric redshifts because of a color-redshift degeneracy. This is described in detail in Section 4.2.3 of Weinstein et al. (2004).

Next, a redshift based on the astrometric data (the *astrometric* redshift) was calculated using the parameters described in Section 6.1. The mean vector and the covariance matrix were calculated using $auPar$ and $agPar$, using the same method as for the SDSS adjacent colors. The astrometric redshift PDF is shown in orange in Figure 14. The 13,028 quasars on Stripe 82 with spectroscopic redshifts and for which we were able to calculate astrometric redshifts are shown in Figure 15 (top right panel). This process gives poorer redshift estimates than the SDSS photometric redshifts, but the purpose is to break degeneracies in the photometric redshifts by combining photometric and astrometric information. That is, the astrometric redshift serves as an informative prior.

Next, the astrometric redshift PDFs and the photometric redshift PDFs are combined using weighted averages in a similar manner as Carrasco Kind & Brunner (2014) (Section 3.1.2 and Equation (7)) to make *astro-photometric* redshifts. Specifically, we have combined the PDFs by adding rather than multiplying in order to enable a relative weighting of the two PDFs. In future work, we will consider a multiplicative joining of the data with smoothing to provide relative weighting. The colors curve is given five times the

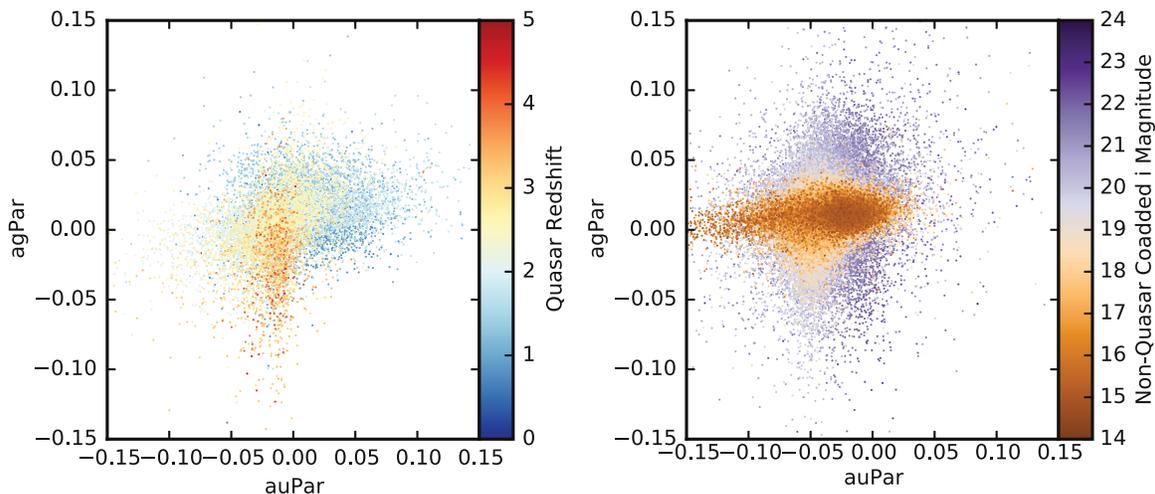


Figure 13. Slope of the line (offset along the parallactic angle) with respect to redshift in the u -band ($auPar$) and g -band ($agPar$) as a function of redshift for the quasar sample (left panel) and as a function of magnitude for non-quasars (right panel). Left panel: while the changes in these astrometric parameters are not as strong as the changes in color with redshift, they provide another source of redshift information. Right panel: the differences between the distributions in the left panel and right panel can aid in the separation of quasars from non-quasars. See Section 9. For example, objects with large negative values of $auPar$ are more likely to be non-quasars than quasars.

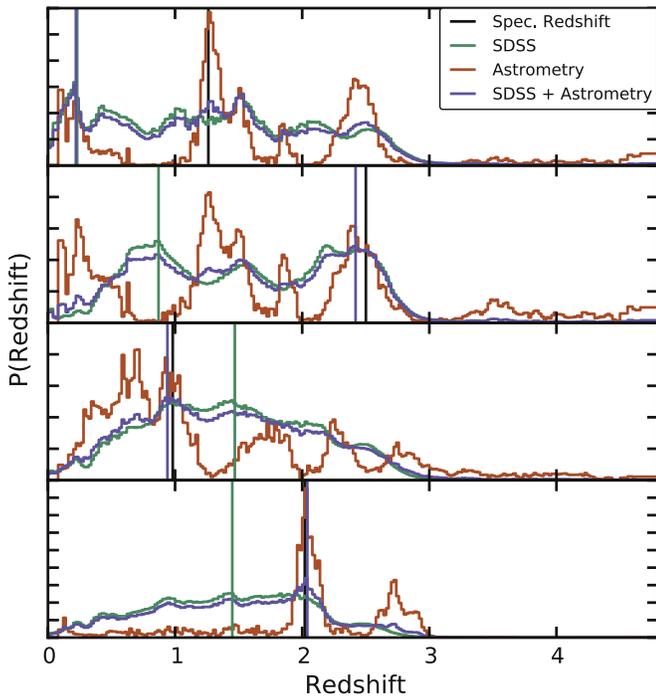


Figure 14. Four selected example quasars demonstrating the photometric redshift probability function using the SDSS colors (green), astrometry (orange), and astro-photometric (purple). The expectation value of the SDSS colors PDF is shown as a vertical green line and the peak of the astro-photometric PDF is shown as a vertical purple line. The spectroscopic redshift is shown as a vertical black line. The colors curve is given five times the weight of the astrometry curve, then the two are added, and finally renormalized to create the purple curve. The top two panels demonstrate how, when the photometric redshifts returned by the colors, are inconsistent with the spectroscopic redshifts, the colors often return the spectroscopic redshift as one of the secondary peaks in the PDF. The astrometric PDF generally has several large peaks or an extended plateau. When the two PDFs are combined it often pulls out the correct peak in the colors PDF. In the top panel, the tertiary peak of the astro-photometric PDF correctly identifies the spectroscopic redshift for a low-redshift quasar where colors alone failed; a different weighting of the colors and astrometry PDFs might have picked up the correct peak. In the second panel, the primary peak of the astro-photometric PDF identifies the spectroscopic redshift for a high-redshift quasar where colors alone failed. The third panel shows how the astrometry PDF helps to identify which peak in the colors PDF is correct. The bottom panel shows how a broad plateau in the colors PDF converges to the spectroscopic redshift by the addition of the astrometry PDF information.

weight of the astrometry curve chosen based on empirical experiments with different weights. The resulting curve is shown in Figure 14 in purple. When the photometric redshifts returned by the colors alone are inconsistent with the spectroscopic redshifts, the correct redshift is generally one of the secondary peaks in the color-based PDF. The astrometric-redshift PDF generally has a plateau at one end of the redshift range or several large peaks. When the two PDFs are combined, it pulls out the correct peak in the color-based PDF as the best estimate of the redshift. The 13,028 training set quasars in Stripe 82 with spectroscopic redshifts and astrometric values are shown in Figure 15 (bottom left panel). Of these objects, 6717 (51.6%) have a calculated astro-photometric redshift within 0.1 and 10,010 (76.8%) are within 0.3, as seen in Figure 16.

Finally, for the 17,321 quasar candidates with matches to the VHS catalog (about 48%) (see Section 6.2) the photometric

redshift was calculated using the SDSS and VHS adjacent colors ($u - g$, $g - r$, $r - i$, $i - z$, $z - J$, $J - H$, $H - K$). The 9244 quasars on Stripe 82 with spectroscopic redshifts and matches to VHS data are shown in Figure 15 (bottom right panel). Of these objects, 4951 (53.6%) have a calculated photometric redshift within 0.1 of the spectroscopic redshift and 7250 (78.4%) are within 0.3, as seen in Figure 16.

Figure 16 demonstrates that adding either NIR colors or astrometric information significantly improves the redshift estimates over using only optical colors. Comparison of the continuously determined redshifts versus the discrete redshift binning from Section 5.2, suggests that the binning method is somewhat more accurate (in terms of having fewer outliers), but not as precise as the astro-photometric redshifts or optical +NIR photometric redshift.

We graphically summarize the quality of the photometric redshifts in Figure 17 by showing the distribution of true redshifts within a given photometric redshift bin. The photometric redshift bins were chosen to match those of the Richards et al. (2006) quasar luminosity function (QLF). It will be necessary to correct for such photometric redshift errors before determining the QLF in Section 8.3. We find that objects with photometric redshifts of $z \sim 1.25$ and $z \sim 3.25$ are particularly robust, whereas the $z \sim 0.85$ objects are often mistaken for $z \sim 2.2$. This is caused by degeneracies in color-redshift space. As shown in Figure 1 of Richards et al. (2001), the colors of particular quasars can fall within the 1σ distribution of the color-redshift relation at many redshifts. Using all four SDSS colors decreases the areas of degeneracy and adding IR colors or astrometry decreases them still further. The degeneracies found in this work are similar to those described in Section 3.4 of Richards et al. (2001).

Overall, we find that optical+NIR magnitudes can improve the photometric redshift accuracy; however, with astro-photometric redshifts we can surpass the improvements due to NIR data alone.

7. CATALOG

From the classification test set, described in Section 2.5, we present a FITS catalog of the 36,569 objects classified as quasars in either Section 5.1 or 5.2. The number of objects and their origin (5.1 or 5.2) is summarized in Table 6 and a description of the columns in the binary FITS catalog table are provided for reference in Appendix. The catalog is available online.

Another Bayesian selection method using optical and mid-infrared (MIR) colors (Richards et al. 2015, submitted) was able to clean out contaminating bright stars using some simple color cuts. We similarly use MIR color cuts to clean bright stars out of our final candidate list. To do so, we matched the quasar candidate catalog to the *WISE* ALLWISE data release.¹⁵ Of our candidates, 19,720 (53.9%) had matches in both W1 and W2 (AB magnitudes). For these objects, we made the following cuts:

$$i < 19.5 \quad (10)$$

$$i < (-5.5(W1 - W2) + 19.5) \quad (11)$$

following Richards et al. (2015, submitted) and using the coadded i magnitude. This process identified 573 candidates that are flagged as likely stellar contaminants in the catalog as

¹⁵ wise2.ipac.caltech.edu/docs/release/allwise/

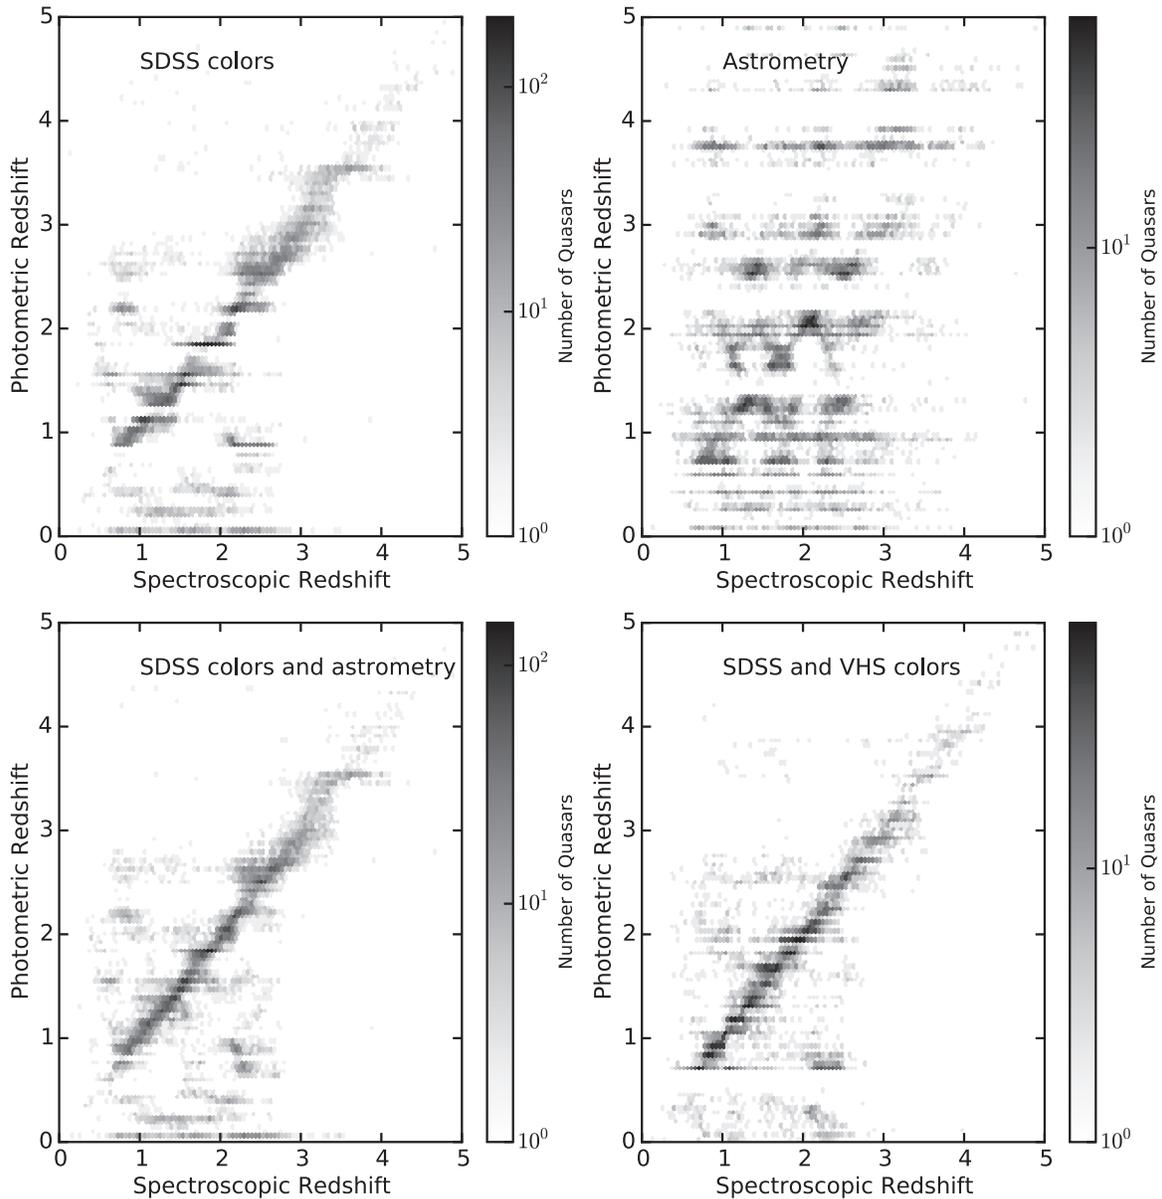


Figure 15. Spectroscopic redshift vs. photometric redshift in hex bins with logarithmic gray scale, using (top left panel) SDSS colors (both single epoch and coadded, when available), (top right panel) astrometry, (bottom left panel) SDSS and astrometry PDFs combined, and (bottom right panel) SDSS and VHS adjacent colors. This illustrates those redshifts where the algorithm has the largest error rate (either due to degeneracy between distinct redshifts or smearing of nearby redshifts). The bottom left panel demonstrates that when the photometric redshifts returned by the SDSS colors are inconsistent with the spectroscopic redshifts, the addition of the astrometry PDF often pulls out one of the secondary peaks in the SDSS PDF as the spectroscopic redshift. The bottom right panel demonstrates how optical+IR magnitudes can similarly improve the photometric redshift accuracy. However, with the addition astrometry we can surpass the improvements due to IR data.

noted in Table 7. The majority of these objects have colors that are consistent with the stellar locus and have a mean i magnitude of 16.8.

Most white dwarf contaminants are below *WISE* detection thresholds. Thus, to eliminate these contaminants we made the following optical color cut, guided by the SDSS white dwarf catalog of Pietro Gentile Fusillo et al. (2015):

$$(r - i) < (-0.62(g - r) - 0.37). \quad (12)$$

We used the coadded magnitudes and confirmed that this cut would remove none of the spectroscopically confirmed quasars from our training set. It removes 48% of the known white dwarfs in Pietro Gentile Fusillo et al. (2015) and identified 178 quasar candidates as possible white dwarfs. These candidates

are flagged as likely white dwarf contaminants in the catalog as noted in Table 7. These possible white dwarfs are all in the bluest corner of $g - r$ versus $r - i$ color space and have a mean i magnitude of 21.7.

All together, after the ALLWISE and white dwarf cuts, there are a total of 35,820 “good” quasar candidates in Stripe 82. (Perform the following query to retrieve these objects from the catalog: `WISEcut_label == 0 & WDcut_label == 0 & candidate_label == 1`.) These candidates are used in the analysis that follows.

Classification over the whole redshift range (as described in Section 5.1) returned 33,240 quasar candidates, or 3.63%, of the 916,587 objects in the test set—roughly consistent with the prior of 5%. Of the 13,221 spectroscopically confirmed quasars

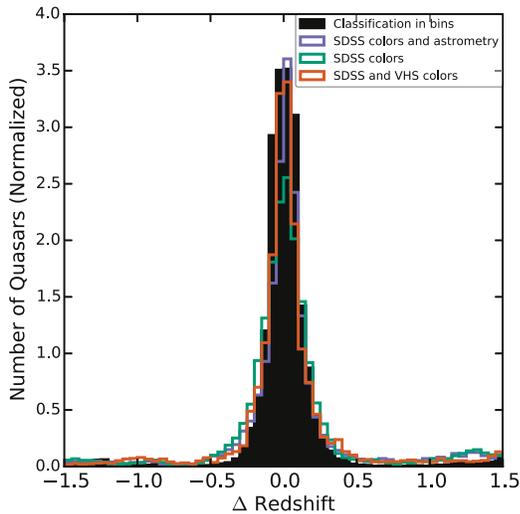


Figure 16. Normalized histogram of the difference between spectroscopic redshift and photometric redshift for quasars. Note how the distribution tightens toward $\Delta z = 0.0$ from the SDSS color photometric redshifts to the astrophotometric redshifts. Shown are SDSS colors (green), SDSS colors and astrometry (purple), and SDSS and VHS colors (orange). Shown in solid black is the histogram of classification in redshift bins from Figure 10.

that could have been returned, we found 12,898 (97.6% completeness). Classification in redshift bins (Section 5.2) returned 31,600 objects as potential quasars. Of the 13,221 spectroscopically confirmed quasars that could have been returned, we found 12,511 (94.6% completeness). Thus, our attempts at simultaneous classification and redshift estimation are somewhat less complete than our efforts to classify quasars regardless of redshift. Using either method, of the 13,221 spectroscopically confirmed quasars that could have been returned, we found 12,953 (98.0% completeness).

Of the candidates, 29,020 (81.0%) were identified by both methods. As shown in Figures 8 and 11, the quasars selected using these two methods show similar distributions. In the bottom panels, we find that the selection in variability parameter space shows no noticeable difference, which is not surprising, as A_g versus γ_g and A_r versus γ_r have no strong redshift trends. However, there are slight differences in color space (top panels). Using the quasar training set in redshift bins we select more $g - r > 1.0$, $u - g < 2.0$, and $z_{\text{phot}} > 3.4$ quasar candidates, many of them potential contaminant stars. Using the full redshift range we select more objects in the bluest corner of $g - r$ versus $r - i$ space, many of them flagged as potential white dwarf contaminants.

As described in Section 2.2, the SDSS I/II quasars were primarily color-selected to $i < 19.1$ for low-redshift and to $i < 20.2$ for high-redshift (Richards et al. 2002), but the target selection on Stripe 82 was deeper, initially going to $i = 19.9$ for low-redshift and $i = 20.4$ for high-redshift; later to $i = 20.2$ for low-redshift sources and $i = 20.65$ for radio sources; and later to $i < 21.0$ for variable sources (Adelman McCarthy et al. 2006). As such, when we consider the completeness of previous spectroscopic observations on Stripe 82, it is important to consider the magnitude of the objects. The “good” quasar candidates are shown in Figure 18. Note the change in character of the new quasar candidates at $i \sim 20.0$.

According to Vanden Berk et al. (2005), the completeness of the SDSS quasar selection algorithm for sources with $i < 19.1$ is $C_q = 94.9^{+2.6}_{-3.8}\%$ at the 90% confidence level. We will

consider the completeness of existing quasar spectroscopy on Stripe 82 both brighter and fainter than this limit. Our region of selection extends beyond the region of uniform spectroscopic follow-up by SDSS: $-10^\circ < R. A. < 50^\circ$, therefore in order to do this comparison, we must limit our examination to this region. This includes 12,107 of the 22,867 “good” quasar candidates in the catalog that are not spectroscopically confirmed. There are 1090 (3183) spectroscopically confirmed quasars brighter than a coadded i -band magnitude of 19.1 (19.9) and we find 61 (192) additional quasar candidates. Assuming that all of our new “good” candidates are real, this completeness of 94.7% (94.3%) agrees well with Vanden Berk et al. (2005). However, we might have expected it to be higher given the additional spectra taken on Stripe 82 since 2005 as part of the BOSS program.

Fainter than this limit, it could be that quasars are not being targeted or that there simply have not been enough fibers devoted to quasar candidates to find all of the objects that we consider to be valid quasar candidates. There are 4591 spectroscopically confirmed quasars dimmer than a coadded i -band magnitude of 19.9 and with a redshift $z < 3.0$. To this we add 9536 quasar candidates with astro-photometric redshift $z < 3.0$. There are 561 spectroscopically confirmed quasars dimmer than a coadded i -band magnitude of 19.9 and with a redshift $z > 3.0$. To this we add 576 quasar candidates with astro-photometric redshift $z > 3.0$.

Figure 19 shows the completeness and new quasar selection as a function of redshift. The left panel shows the quasars and candidates for $i < 19.9$ and right panel shows $i > 19.9$. In short, we have shown that current methods, (only colors, only variability, and other techniques used for Stripe 82 target selection) still are incomplete. Next-generation surveys like LSST will have to adopt more sophisticated methods, of which ours is just a pilot example, to fully exploit the data.

While we find new quasars in Stripe 82, the catalog also includes 466 objects that were *not* selected by our algorithms as quasar candidates, but that are spectroscopically confirmed quasars. This incompleteness demonstrates where there is room for improvement beyond our pilot project. For the sake of completeness, to illustrate where we may be less sensitive, and to make it easier to compute the completeness corrections for our catalog without needing another data source, these quasars are included in our catalog. They are indicated by `candidate_label == 0`. In general, they are in the densest part of the stellar locus and have very small γ_g and γ_r values. More than 50% are between redshifts $2.2 < z < 3.2$ and more than are third are $i > 21.5$ compared to 5% and 9% of the quasar training set as a whole, making these objects particularly difficult to distinguish as quasars.

8. DISCUSSION

We will now explore the quality of the quasar catalog by comparing to other cuts and catalogs, in addition to evaluating it for remaining contaminants. In Section 8.1 we use the quasar variability selection box from Schmidt et al. (2010) as a comparison in time-domain classification. In Section 8.2, we evaluate how well our algorithm recovers quasars from BOSS DR10 and DR12 quasar catalogs. Finally, in Section 8.3, we evaluate completeness and contamination of the candidate quasars using number counts and luminosity function analysis.

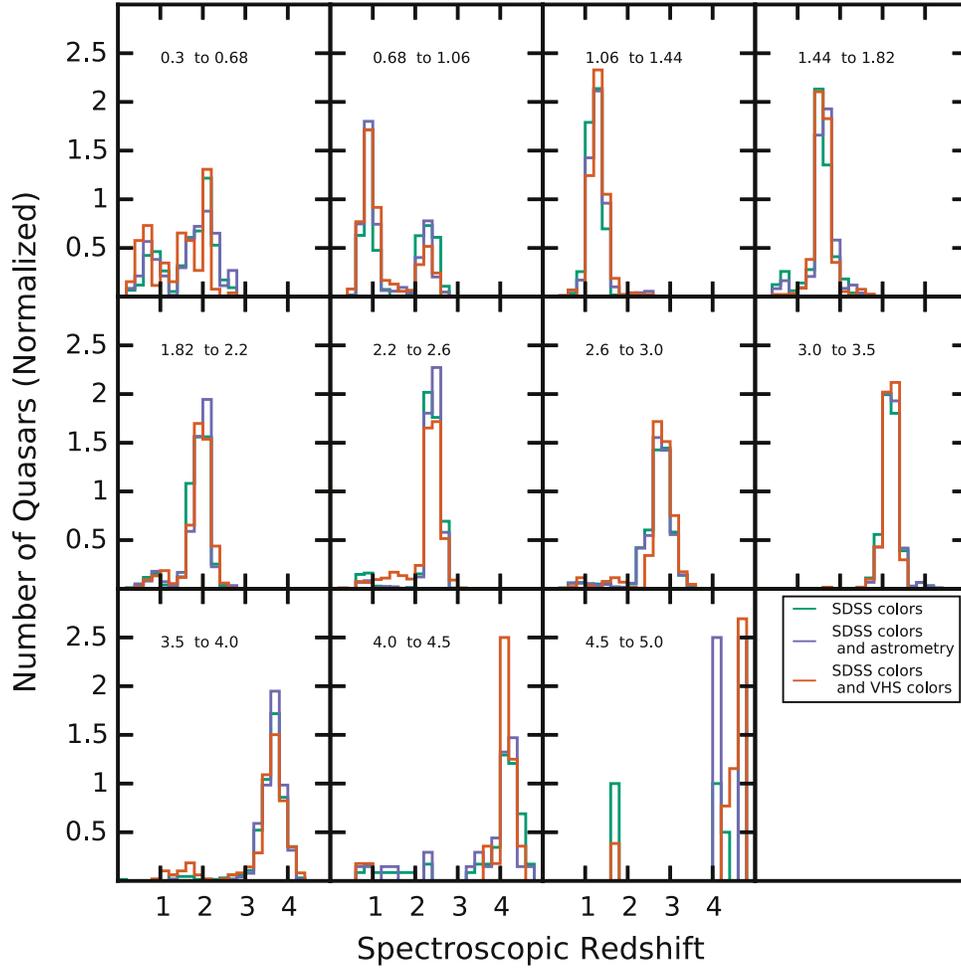


Figure 17. Normalized histogram of spectroscopic redshift in panels based on bins of photometric redshift from 0.3 to 5.0 in the same bins as the luminosity function in Section 8.3. These panels demonstrate which photometric redshift ranges are most unreliable and most reliable. Photometric redshifts were calculated using SDSS colors (green), SDSS colors and astrometry (purple), and SDSS and VHS colors (orange). In particular, note the bimodal distribution at $0.68 < z_{\text{phot}} < 1.06$ compared to the precision at $1.06 < z_{\text{phot}} < 1.44$ and $3.0 < z_{\text{phot}} < 3.5$. This bimodality is caused by degeneracies in color-redshift space. We correct for photometric redshift errors when calculating the quasar luminosity function in Section 8.3.

Table 6
Quasar Candidates

Data Set	Candidate Quasars		w/Spectra		w/o Spectra				
	Total	Fraction	Total	Completeness	Total	$i < 19.9$	$i > 19.9$		
All Candidates	36569	0.040	12953	0.980	23616	1570	0.066	22046	0.934
Whole Redshift Range	33673	0.037	12898	0.976	20775	1048	0.050	19727	0.950
Redshift Bins	32108	0.035	12511	0.946	19597	1282	0.065	18315	0.935
Both Methods	29212	0.032	12456	0.942	16756	760	0.045	15996	0.955
After <i>WISE</i> and WD Cut	35820	0.039	12953	0.980	22867	991	0.043	21876	0.957

8.1. Comparison to Other Variability-based Selection

First we compare our results to the performance of the (variability-based) quasar selection box (in A and γ space) defined in Equations (7)–(9) of Schmidt et al. (2010):

$$\gamma_r = 0.5 \log(A_r) + 0.50 \quad (13)$$

$$\gamma_r = -2.0 \log(A_r) - 2.25 \quad (14)$$

$$\gamma_r = 0.055. \quad (15)$$

Using Stripe 82 data, Schmidt et al. (2010) achieve a completeness of 90% and an efficiency of 96% with this box. Applying the same cuts to our own training sets, as shown in Figure 20's left panel, results in 87% completeness and 74% efficiency. We achieved very different results because we have very different quasar and non-quasar data sets. Schmidt et al. (2010) used quasars with $15.4 < i < 22.0$ with a mean of 19.5 and only 5000 bright F/G-star colored objects with $0.2 < g - r < 0.48$ and $14.0 < g < 20.2$. We used quasars with $15.9 < i < 22.7$ with a mean of 20.2 and 72,680 non-

Table 7
Column Names

Index	Name	Description
1	id	SDSS Coadded ParentID
2–3	R.A., decl.	coadded right ascension, declination
4–8	u, g, r, i, z	single epoch magnitude
9–13	uErr, gErr, rErr, iErr, zErr	single epoch magnitude error
14–18	coadd_u, coadd_g...	coadded magnitude
19–21	J, H, KS	VHS magnitude
22–24	JErr, HErr, KSErr	VHS magnitude error
25–26	W1, W2	WISE magnitude
27–28	W1Err, W2Err	WISE magnitude error
29	zspec	spectroscopic redshift, if none, the value is -9999 .
30	candidate_label	1 if selected as a quasar candidate, 0 otherwise.
31	MQC_label	1 if in MQC, 0 otherwise.
32	DR10Q_label	1 if in DR10Q, 0 otherwise.
33	DR12Q_label	1 if in DR12Q, 0 otherwise.
34	WISEcut_label	1 if cut by Equations (10) and (11), 0 otherwise.
35	WDcut_label	1 if cut by Equation (12), 0 otherwise.
36–45	A_u, gamma_u, A_g, gamma_g...	variability parameters, if none, the value is -9999 .
46–47	auPar, agPar	astrometry parameters, if none, the value is -9999 .
48	qso_prob	probability of being a quasar
49	star_dens	star density
50	qso_dens	quasar density
51	qso_prob_bins	vector—probability of being a quasar
52	star_dens_bins	vector—star density
53	qso_dens_bins	vector—quasar density
54	qso_prob_max	maximum value of <i>qso_prob_bins</i> vector
55	qso_prob_max_bin	redshift bin of maximum value of <i>qso_prob_bins</i> vector
56	photoz_ugriz_pdf	vector—full photo-z PDF, SDSS colors
57	photoz_ugriz_low	low redshift end of the peak in photo-z PDF, SDSS colors
58	photoz_ugriz_best	peak of photo-z PDF, SDSS colors
59	photoz_ugriz_high	high redshift end of the peak in photo-z PDF, SDSS colors
60	photoz_ugriz_prob	probability of photo-z, SDSS colors
61	photoz_astrometry_pdf	vector—full photo-z PDF, astrometry
62	photoz_astrometry_low	low redshift end of the peak in photo-z PDF astrometry
63	photoz_astrometry_best	peak of photo-z PDF astrometry
64	photoz_astrometry_high	high redshift end of the peak in photo-z PDF astrometry
65	photoz_astrometry_prob	probability of photo-z astrometry
66	photoz_added_pdf	vector—full photo-z PDF, SDSS colors and astrometry
67	photoz_added_max	max of photo-z PDF, SDSS colors and astrometry
68	photoz_added_max_bin	redshift bin of maximum value of photo-z PDF, SDSS colors and astrometry
69	photoz_ugrizJHK_pdf	vector—full photo-z PDF, SDSS and <i>JHK</i> colors
70	photoz_ugrizJHK_low	low redshift end of the peak in photo-z PDF, SDSS and <i>JHK</i> colors
71	photoz_ugrizJHK_best	peak of photo-z PDF, SDSS and <i>JHK</i> colors
72	photoz_ugrizJHK_high	high redshift end of the peak in photo-z PDF, SDSS and <i>JHK</i> colors
73	photoz_ugrizJHK_prob	probability of photo-z, SDSS and <i>JHK</i> colors
74	gi_sigma	$g - i$ color offset from the mean color
75	SDSSSPECMATCH	1 if the object had a spectrum from the original SDSS, 0 otherwise
76	BOSSSPECMATCH	1 if the object had a spectrum from BOSS, 0 otherwise
77	DR12QSOMATCH	1 if the object is visually inspected as a quasars in the DR12Q, 0 otherwise
78	ZSDSS	pipeline redshift from SDSS
79	CLASSSDSS	pipeline classification from SDSS
80	ZBOSS	pipeline redshift from BOSS
81	CLASSBOSS	pipeline classification from BOSS
82	DR12QSO_Z_VI	redshift of the quasars if included in the DR12Q

(This table is available in its entirety in FITS format.)

quasars (not just F/G stars) with $14.8 < g < 25.5$ and a mean of 20.6.

Applying these cuts instead to our full test set, as shown in Figure 20 right panel, gives 49,649 quasar candidates. Of these, 23% are spectroscopically confirmed quasars and another 27% are objects that we identified as quasar candidates in either

Section 5.1 or 5.2 (with the remaining being previously unidentified potential new candidates). If all of our previously identified candidates were actually quasars and the remaining objects identified by these cuts were instead contaminants, then the efficiency of this variability quasar selection box would be 50% and the completeness would be 69%. The majority of the

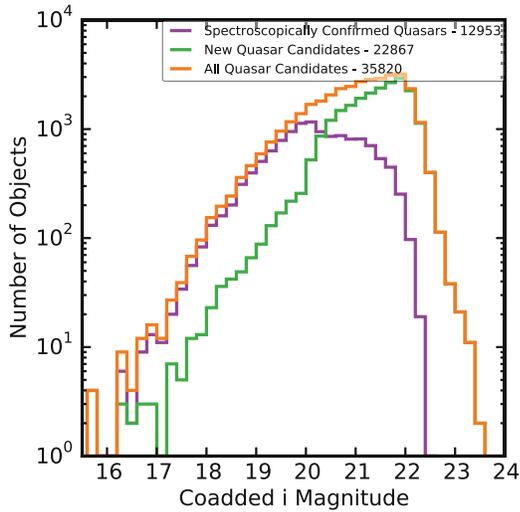


Figure 18. Histogram of coadded i -band magnitude for known Stripe 82 quasars and new quasar candidates. In purple are the previously known, spectroscopically confirmed quasars returned by the selection. The quasar candidates returned by the selection are shown in orange and the new quasar candidates are shown in green.

quasar candidates outside the box are dimmer than a coadded i -band magnitude of 20, where most variability is below the noise level.

This comparison suggests that selection by variability alone, while working well to discriminate between relatively bright F/G stars and quasars, is incomplete when using a realistic sample of non-quasar contaminants, and that our hybrid approach of combining color and variability will yield better results for future surveys.

In Graham et al. (2014) they compare the performance of variability selection using a power law fit to the SF, a DRW fit, and SWV. Using the power-law fit to the SF (SWV) to classify quasars on Stripe 82 they achieve 92% (86%) completeness and 93% (92%) efficiency.

8.2. BOSS Quasar Selection

As described in Section 2.2, in addition to color selection, some of the BOSS quasars on Stripe 82 were targeted using an algorithm based on the same parameterization of variability used herein. We matched our candidate catalog to the SDSS-III/BOSS Data Release 10 Quasar Catalog (DR10Q; Pâris et al. 2014) to see how well we recovered these quasars. These quasars are indicated by `DR10_label == 1`. There are 9590 quasars on Stripe 82 in DR10Q and 7241 were point sources that met the quality cuts to be included in our test set. Of these 7241 known quasars, we recovered 7034 (97.1% completeness) as candidate quasars. The quasars we missed have $i < 22.0$ with a mean of $i = 20.0$ and have $\gamma_g < 0.25$: much less variable than the quasar training set on average.

We found 6562 quasar candidates in the BOSS redshift range ($2.2 < z < 3.5$) based on astro-photometric redshifts. Of these, 49% are training set quasars with spectroscopic redshifts $2.2 < z < 3.5$ ($i < 22.7$ with mean $i = 20.7$) and another 3% are known quasars with spectroscopic redshifts outside this range. Of the remaining 48% (3157 quasar candidates), 1614 are high probability candidates with `qso_prob > 0.8`. These are the objects that are highly likely to be quasars that BOSS has missed, which is consistent with the known incompleteness of BOSS (Ross et al. 2012). Our high probability candidates have $i < 23.0$ with a mean of $i = 21.4$, suggesting that we are able to extend our selection to less luminous objects using the combined color and variability approach.

Since our test set was built, the 12th data release quasar catalog of SDSS-III was made public (DR12Q; I. Pâris et al. 2015, in preparation). Since DR10Q, additional spectroscopic plates were taken on Stripe 82, resulting in 2054 DR12Q quasars on Stripe 82 that are not in the quasar training set, 1162 were point sources that met the quality cuts to be included in our test set. We matched our candidate catalog to DR12Q to see how well we recovered these new quasars. These objects are indicated by `DR12_label == 1`. Of the quasars new in

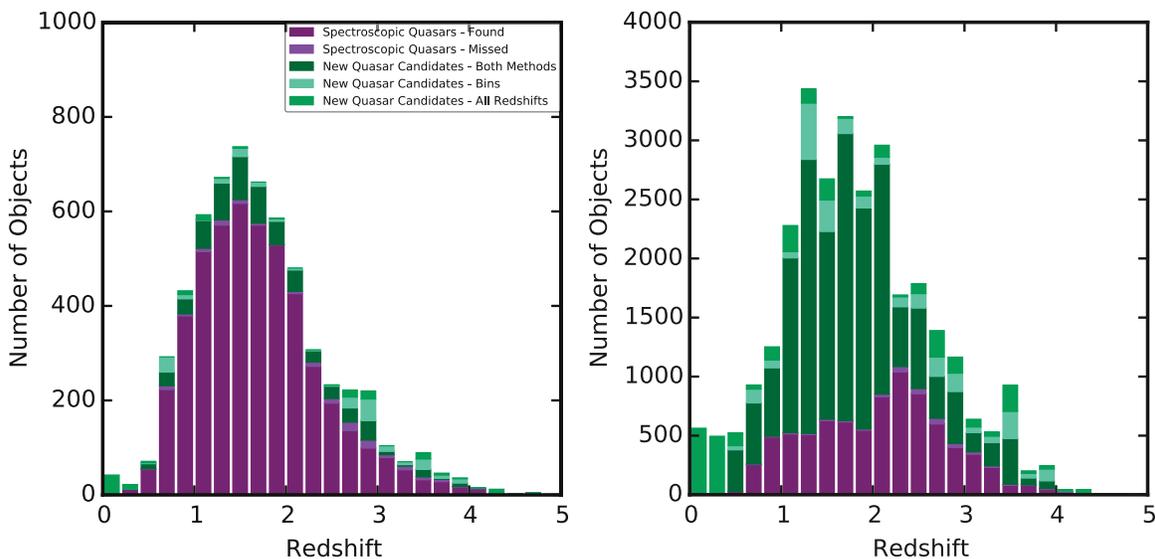


Figure 19. Stacked histogram of redshift for known Stripe 82 quasars and new quasar candidates. The left panel shows the quasars and candidates $i < 19.9$ and the right panel shows $i > 19.9$. Spectroscopic quasars found as quasar candidates and spectroscopic quasars missed are both binned by spectroscopic redshift. Quasar candidates found by both methods and quasar candidates found only using a binned quasar training set are both binned by where the candidate was classified with the highest probability. These bins only span $0.4 < z < 4.0$. Quasar candidates found only using a quasar training set over the full redshift range are binned by the astro-photometric redshift.

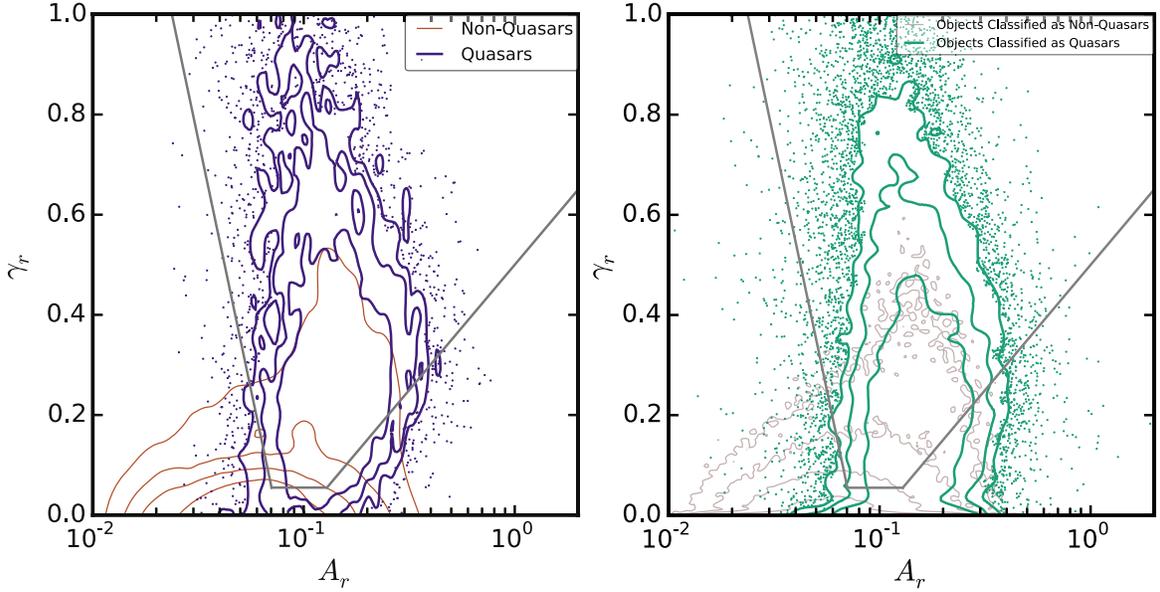


Figure 20. A_r vs. γ_r for the training sets (left panel) and test set (right panel) shown with the Schmidt et al. (2010) variability selection cuts (Equations (13)–(15)) as gray lines. *Left panel:* orange contours show the non-quasar training set and purple contours and scatter points show the quasar training set. *Right panel:* gray contours show all objects in the test set classified as non-quasars and green contours and scatter points show all objects in the test set classified as quasars.

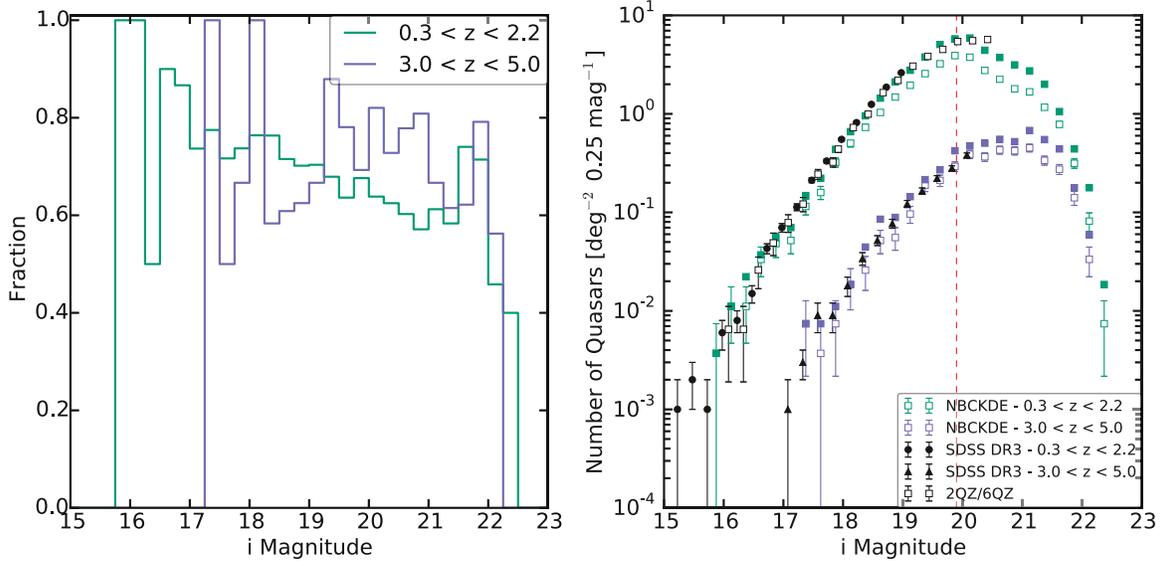


Figure 21. Left panel: ratio of quasars in the MQC on Stripe 82 to those MQC quasars returned by classification using a training set over the full redshift range. This allows us to correct for objects with too few observations to calculate variability parameters, the exclusion of extended sources, and incompleteness in the selection algorithm. The fraction is given as a function of coadded i -band magnitude for two redshift ranges. Right panel: quasar number counts as a function of redshift and i -band magnitude. Black points give the spectroscopic number counts reported in Richards et al. (2009b); circles for $z < 2.2$ and triangles for $3 < z < 5$. The open purple and green squares give the raw number counts (with Poisson error bars) for the candidates reported here. The filled colored squares give the number counts corrected using the left panel. The vertical dashed red line at $i = 19.9$ indicates the target selection depth for low-redshift on Stripe 82.

DR12Q, we recovered 1141 (98.2% completeness). The objects that were missed have $i < 22.1$ with a mean of $i = 21.3$ and have $\gamma_g < 0.33$. Again, they are much less variable than the quasar training set on average.

8.3. Number Counts and the Luminosity Function

In Figure 21 we reproduce the number counts analysis shown in Figure 9 of Richards et al. (2009b), using our candidate quasars. The counts have been corrected for incompleteness as given by the fraction of MQC quasars

recovered as shown in the left panel of Figure 21. In short, the correction is the ratio of known quasars to quasar candidates. This process corrects for: objects with too few observations to calculate variability parameters, the exclusion of extended sources, and incompleteness in the selection algorithm. The right panel shows the number of quasars per deg^2 and per 0.25 mag as a function of coadded i -band magnitude. Open points represent the raw number counts, while the closed points give the completeness-corrected number counts. The turnover at $i = 19.9$ is due to the incompleteness of the spectroscopic sample. This analysis suggests that our selection algorithm is

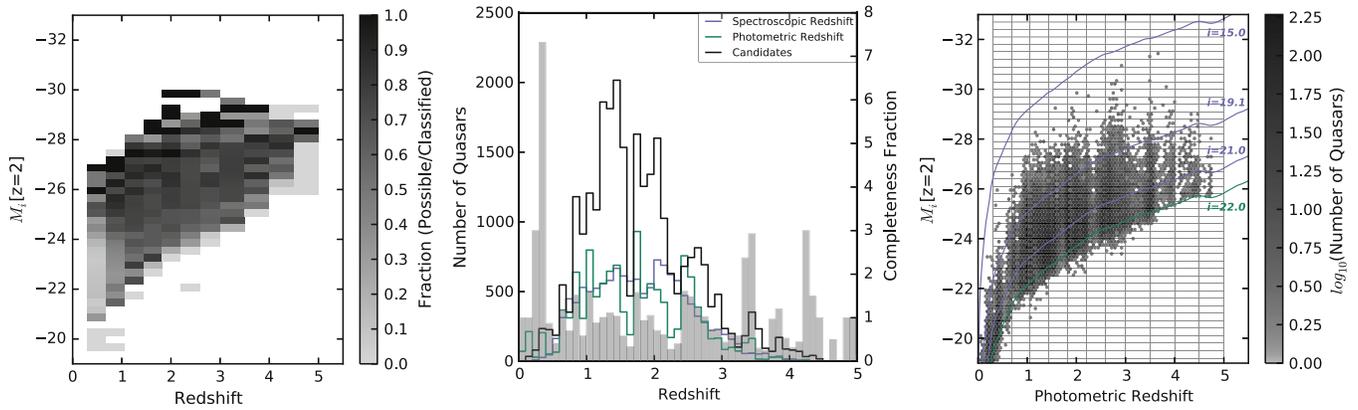


Figure 22. Corrections and cuts used in the QLF in Figure 23. Left panel: completeness fraction, in bins of redshift and absolute magnitude, $M_i[z=2]$, for candidate selection. Similar to Figure 21’s left panel, but in two dimensions. The number of quasars with spectroscopic redshifts on Stripe 82, even if they were excluded from our training set and test set, was divided by all quasars with spectroscopic redshifts that were recovered as candidate quasars. This is to correct for incompleteness from too few observations to calculate variability parameters, the exclusion of extended sources, and incompleteness in the classification algorithm. Center panel: completeness fraction for astro-photometric redshifts. All of the training set quasars are binned by spectroscopic redshift (purple) and astro-photometric redshifts (green). The ratio of the two is shown in gray (right axis). The astro-photometric redshifts of the candidate quasars, after being corrected by the completeness fraction and assuming that objects without spectroscopic redshifts have the same astro-photometric redshift errors as those with spectroscopic redshifts, are shown in black. Right panel: astro-photometric redshift vs. absolute magnitude, $M_i[z=2]$, of all quasar candidates. The green line shows the brightness limit for bins that are used in computing the luminosity function. Purple curves show the $i = 15.0, 19.1,$ and 20.2 mag limits for SDSS spectroscopic follow-up.

neither heavily contaminated (e.g., as might be evidenced by a large excess of bright objects versus known quasars), nor very incomplete—since the corrected counts agree well with the spectroscopic quasar distribution.

Next, we calculate the QLF for our candidate quasars. This QLF calculation was not intended to be a scientific result of this pilot project, as we expected more incompleteness and contamination than shown in Figure 21. However, the result does suggest that accurate determination of the QLF will be possible with photometric selection from LSST and other next-generation surveys.

In order to compare space densities at different redshifts, we must correct our photometry for the effects of redshift on the portion of the spectrum sampled by a given filter. We do this by using a mean K-correction for $z = 2$ in the i -band as described in Richards et al. (2006, Section 5).

As we have seen, and as discussed in Ross et al. (2013, Section 3.4.1), variability selection is less biased than color selection, but we cannot assume variability selected samples are complete and unbiased. Just as with the number counts above, the candidate object QLF must be corrected for the completeness fraction and, additionally, for systematic errors in astro-photometric redshifts. For the QLF, we need to correct for incompleteness in two dimensions: redshift and absolute magnitude (luminosity). The gray-scale $M - z$ bins in the left panel of Figure 22 gives the fraction of MQC quasars recovered. This includes quasars that were not included in our test set so as to correct to the true number of quasars, not just those that met our test set criteria.

Since catastrophic errors in astro-photometric redshifts can distort the QLF, corrections were determined as follows: using bins of $\Delta z = 0.1$, the number of quasars with astro-photometric redshift in that bin was divided by the number of quasars with spectroscopic redshifts in that bin. The resulting ratio is the correction that needs to be applied to objects in each astro-photometric redshift bin to statistically account for errors in the astro-photometric redshift distribution (as opposed to correcting individual values) and is shown in Figure 22’s center panel.

The two corrections are multiplied together and used as a weight for the objects in the QLF.

We compute the QLF by binning the quasars in redshift and absolute magnitude, using the method from Page & Carrera (2000). Figure 22’s right panel shows absolute magnitude as a function of astro-photometric redshift for all quasar candidates. The grid shows the bins within which the QLF is calculated. The edges of the redshift bins are 0.30, 0.68, 1.06, 1.44, 1.82, 2.20, 2.6, 3.0, 3.5, 4.0, 4.5, and 5.0. The M_i bins are in increments of 0.3 mag. The adopted limiting magnitude of $i = 22.0$, is shown as a green line. The resulting i -band QLF is shown as black dots with Poisson error bars in Figure 23.

As with the number counts, the QLF analysis shows relatively close agreement with the space density of known quasars. There is evidence for both incompleteness and contamination in the lowest redshift bin. This is perhaps not surprising given the effects the host galaxy has on quasar colors and apparent variability and the fact that we only include point sources. We show the Richards et al. (2006, Figure 18) and Ross et al. (2013, Figure 11) SDSS spectroscopic QLFs in the $z = 2.4, 2.8,$ and 3.25 bins. This comparison reveals that our QLF agrees better with the Ross et al. (2013) QLF. The Ross et al. (2013) QLF has the smaller corrections of the two spectroscopic QLFs, which suggests that the Richards et al. (2006) QLF was undercorrected. In the three highest redshift bins our QLF suggests a higher space density than the Richards et al. (2006) QLF. This could be a sign of contamination in our catalog, though it could also be true to some extent, given the relatively large completeness fraction for candidate selection needed for the smaller spectroscopic sample from which the Richards et al. (2006) QLF was derived. Most importantly, given the lack of contamination and the dependability of the completeness corrections, this analysis bodes well for our future ability to determine the QLF for faint populations in post-SDSS sky surveys.

9. FUTURE WORK

The purpose of this investigation was to demonstrate that using a combination of optical colors and variability parameters

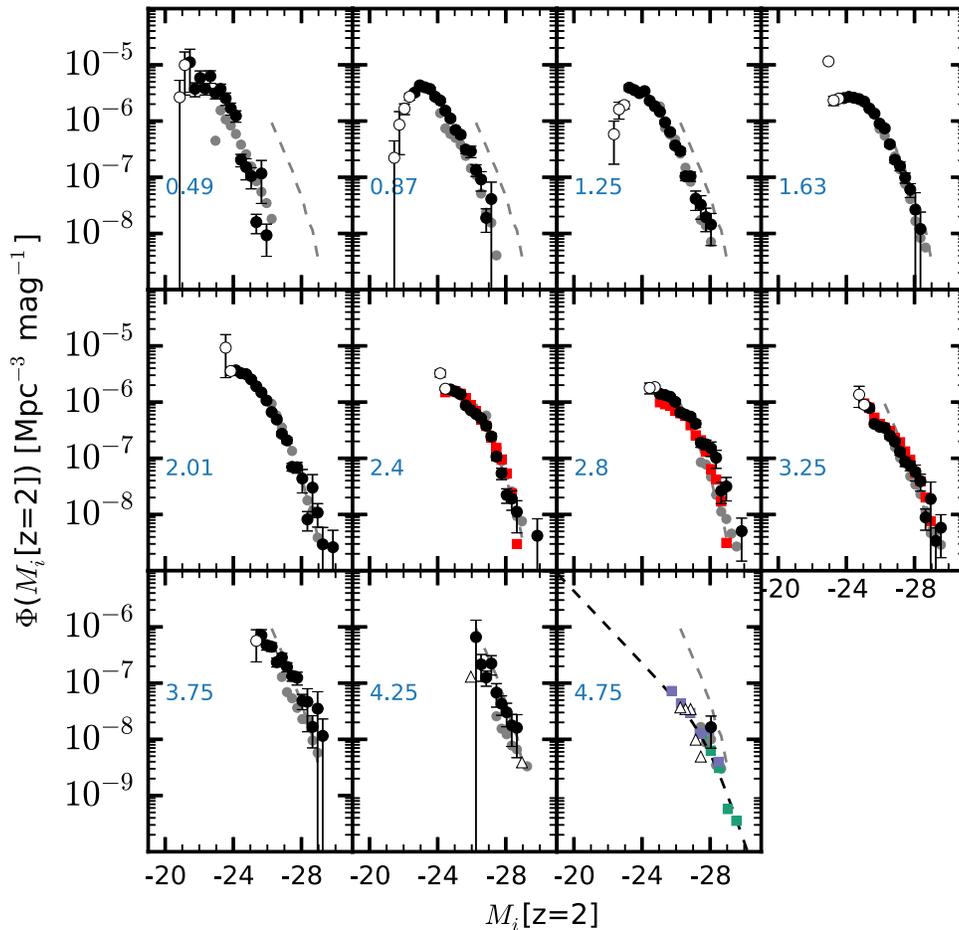


Figure 23. $M_i[z=2]$ binned luminosity function of the sample with astro-photometric redshifts using the method from Page & Carrera (2000) (with Poisson error bars). The mean redshift of each slice is given in each panel. Black filled circles are complete bins, empty triangles indicate the lower limit for complete bins where the completeness fraction (shown in Figure 22’s left panel) is 0, and empty circles are partial bins (a portion of the bin is dimmer than $i=22$). The gray circles show the binned luminosity function and the gray dashed line shows the $z=2.01$ curve both from Richards et al. (2006, Figure 18) for comparison. In the $z=2.4, 2.8,$ and 3.25 panels, the red squares show the binned luminosity function for BOSS quasars from DR9 from Ross et al. (2013, Figure 11). In the 4.75 panel, the green squares, purple squares, and dashed black line show the binned luminosity function at $z=4.9$ for Stripe 82, DR7, and double power law fits from the maximum likelihood analysis from McGreer et al. (2013, Figures 12 and 13).

improves quasar classification efficiency and completeness over the use of colors alone. This is one step toward finding an optimal strategy for photometric quasar selection.

In the future, we hope to use a data set that includes both point sources and extended sources, thus incorporating the variable nucleus with the steady host galaxy. Additionally, we plan to explore alternative parameterizations of quasar variability. The underlying mechanism and most appropriate model for quasar variability remain open questions and there are more sophisticated models (e.g., Kelly et al. 2009; MacLeod et al. 2010; Kasliwal et al. 2015) that merit exploration. Given the large quantity of data expected in future surveys such as LSST, a more computationally efficient approach than the SF may become important; e.g., the Kelly et al. (2009), Kozłowski et al. (2010), and MacLeod et al. (2010) approaches require only $O(N)$ rather than $O(N^2)$ operations to determine the model parameters for a light curve with N data points. As described in Section 2.4, the likelihood method is biased and more robust approaches such as those described in the appendices of Kozłowski et al. (2010) or Hernitschek et al. (2015) should be investigated. Currently, we use variability data from each band separately. We hope to explore the various methods for merging bands together, even

with non-simultaneous observations, as will be the case with LSST.

This work relies on KDE for classification and it is important to explore other methods to see if they will be more successful. In the future, we hope to make use of other types of algorithms (e.g., Feigelson et al. 2003, 2012; Chakraborty et al. 2013), such as random forests (e.g., Gao et al. 2009; Richards et al. 2011; Carrasco Kind & Brunner 2013), gradient boosting machines (Hastie et al. 2001), and Bayesian classification with hash tables (Gupta et al. 2014). Additionally, in the future our catalogs will not have binary classifications, but will simply give probabilities for all objects.

We have used the combination of optical and MIR colors for quasar selection in another paper (Richards et al. 2015, submitted). In the future we will combine optical, IR, and variability data to produce the most complete and efficient catalog possible.

In order to improve the astro-photometric redshift estimations, we will multiply smoothed PDFs instead of adding by weights. Additionally, we will incorporate the clustering redshift estimation of Ménard et al. (2013) and Rahman et al. (2015) and explore photometric and astro-photometric redshift accuracy without u -band observations (to mimic DES and Pan-

STARRS observations). Finally, we will explore how simultaneous color and variability classification performs using other time-domain surveys, including DES (The Dark Energy Survey Collaboration 2005), Pan-STARRS (Kaiser et al. 2010), and LSST simulated data (Connolly et al. 2014).

10. CONCLUSIONS

Using the NBC KDE quasar selection algorithm, we demonstrated that using a combination of optical colors and variability parameters improves quasar classification efficiency and completeness over the use of colors alone. For classification using colors alone, there are redshift ranges with poor completeness where the quasar and non-quasar training sets overlap in color space. Variability alone does not have these redshift trends, but it has a lower efficiency than coadded colors at all other redshifts. The addition of coadded colors to variability information improves classification at all redshifts. Using variability alone, colors alone, and combining variability and colors we achieve we achieve 91%, 93%, and 97% quasar completeness and 98%, 98%, and 97% efficiency, respectively, with particular improvement in the selection of quasars at $2.7 < z < 3.5$, as shown in Figure 7.

We classified quasars and estimated their redshifts simultaneously by limiting the training set to non-overlapping redshift bins from 0.4 to 4.0 with a bin width of 0.2. We successfully classified known quasars into the correct redshift bins with 75% or higher completeness, depending on the redshift bin, as shown in Figure 9.

Overall, we identified 35,820 type 1 quasar candidates in the SDSS Stripe 82 field using the combination of optical photometry and variability either over the full redshift range or within one of the redshift bins. Of the 13,221 spectroscopically confirmed quasars that could have been returned, we found 12,953 (98.0% completeness). Of the 22,867 quasar candidates that are not spectroscopically confirmed, 21,876 (95.7%) are dimmer than a coadded i-band magnitude of 19.9. Figure 18 shows the magnitude distribution of the candidate quasars.

Photometric redshift estimates of these candidates using optical photometry and astrometric parameters are accurate to within $|\Delta z| < 0.1$ for 51.6% of quasars and within 0.3 for 76.8% of quasars. The combination of optical photometry and astrometry makes the photometric redshifts more accurate when colors alone returns the correct redshift as one of the secondary peaks in the PDF. The astrometric PDF pulls out the correct peak in the color PDF, as shown in Figure 14. We find that objects with photometric redshifts of $z \sim 1.25$ and $z \sim 3.25$ are particularly robust.

In Figure 20, our color and variability selection was compared to other cuts in variability space that have been used on Stripe 82. We demonstrated that variability alone is incomplete and that our hybrid approach will yield better results for future surveys. Additionally, we have shown that our selection recovered 97% of the quasars in the DR12 quasar catalog and we selected additional candidates in the BOSS redshift range with high confidence (and at even higher redshift).

We used MIR color cuts to remove a small number of bright star contaminants from our final candidate list. Our number counts and QLF analyses, shown in Figures 21 and 23, show there is little contamination remaining and that there is

relatively close agreement with the space density of known quasars.

From the NBC KDE classification test set, we present a catalog of known quasars and candidate quasars on Stripe 82. The catalog is available as a FITS file. Future work along these lines will be needed to capitalize on the imaging data produced by Pan-STARRS, DES, and LSST.

This work was supported in part by NASA-ADAP grant NNX12AI49G and NSF grant 1411773. This research made use of Astropy¹⁶ (Astropy Collaboration et al. 2013), TOPCAT,¹⁷ (Taylor 2005), and STILTS¹⁸ (Taylor 2006). All figures made use of matplotlib¹⁹ (Hunter et al. 2007). Figures 1, 3, 8, 11, and 20 were made using densityplot²⁰ (Krawczyk & Peters 2014). Special thanks go to Alex Gray for developing the code for the NBC KDE algorithm and to Erica Smith for carefully reviewing the manuscript.

Funding for the SDSS and SDSS-II has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, the U.S. Department of Energy, the National Aeronautics and Space Administration, the Japanese Monbukagakusho, the Max Planck Society, and the Higher Education Funding Council for England. The SDSS website is <http://www.sdss.org/>.

Funding for SDSS-III has been provided by the Alfred P. Sloan Foundation, the Participating Institutions, the National Science Foundation, and the U.S. Department of Energy Office of Science. The SDSS-III website is <http://www.sdss3.org/>.

APPENDIX CATALOG COLUMNS

In Section 7 we created a catalog of quasar candidates. All of the columns in the catalog are described in Table 7, but a few columns need some extra explanation.

Columns 4–13 are the single epoch magnitudes and magnitude errors used for the single epoch classification. We used a randomly chosen epoch from the observations for each object. The single epoch magnitudes are asinh magnitudes from Lupton et al. (1999). Columns 14–18 are the coadded magnitudes and magnitude errors. The coadded magnitudes are from Annis et al. (2014). Columns 19–24 are the VHS magnitudes and magnitude errors in Vega. Columns 25–28 are the *WISE* magnitudes and magnitude errors in AB.

Columns 30–34 are labels. Specifically, column 30 is the candidate label: if the object was classified as a quasar in either Section 5.1 (over the whole redshift range) or 5.2 (in redshift bins) the value is 1, otherwise it is 0. Column 31 is the MQC label: if the object is in the catalog the value is 1, otherwise it is 0. Column 32 is the DR10Q label: if the object is in the catalog the value is 1, otherwise it is 0. Column 33 is the DR12Q label: if the object is in the catalog the value is 1, otherwise it is 0. Column 34 is the *WISE* cut label: if the object is cut by Equations (10) and (11) the value is 1, otherwise it is 0. Column 35 is the white dwarf cut label: if the object is cut by Equation (12) the value is 1, otherwise it is 0. To retrieve the quasar candidates that pass these cuts (the

¹⁶ astropy.org

¹⁷ www.star.bris.ac.uk/~mbt/topcat

¹⁸ www.star.bris.ac.uk/~mbt/stilts

¹⁹ matplotlib.org

²⁰ github.com/CKrawczyk/densityplot

“good” quasar candidates) perform this query on the catalog: `WISEcut_label == 0 & WDCut_label == 0 & candidate_label == 1`. To limit to the new candidates (not spectroscopically confirmed quasars) add: `& zspec < 0`.

Columns 48–55 are the various classification results. Specifically, columns 48–50 are the results of classifying the test set over the full redshift range as described in Section 5.1. Column 48 is the probability of being a quasar, column 49 is the star density from the KDE ($P(D|M)$), and column 50 is the quasar density from the KDE. If the object was not found to be a candidate over the full redshift range the value is -9999 . Columns 51–55 are the results of classification using redshift bins as described in Section 5.2. Column 51 is the probability of being a quasar, column 52 is the star density, and column 53 is the quasar density. Each is a vector with 18 cells, one for each redshift bin from 0.4 to 4.0. If the object was not found to be a candidate in any bin, all cell values are -9999 . Column 54 is the maximum value of column 51, and column 55 is the center of the redshift bin corresponding to that maximum probability. If the object was not found to be a candidate in any bin these columns will be -9999 .

Columns 56 to 73 are the various redshift estimation results. If we were unable to calculate a redshift estimate for the object the value will be -9999 . Column 74 indicates whether the object’s $g - i$ color is within 1σ (0.68), 2σ (0.95), or 3σ (0.99) of the mean color for quasars at the astro-photometric redshift. Outliers are an indication of either bad estimated redshifts or non-quasar contaminants. Columns 75–82 are the details of matching to all spectra taken on SDSS Stripe 82.

REFERENCES

- Abazajian, K. N., AdelmanMcCarthy, J. K., Agüeros, M. A., et al. 2009, *ApJS*, **182**, 543
- Abraham, S., Philip, N. S., Kembhavi, A., Wadadekar, Y. G., & Sinha, R. 2012, *MNRAS*, **419**, 80
- AdelmanMcCarthy, J. K., Agüeros, M. A., Allam, S. S., et al. 2006, *ApJS*, **162**, 38
- Ahn, C. P., Alexandroff, R., Allende Prieto, C., et al. 2014, *ApJS*, **211**, 17
- Aihara, H., Allende Prieto, C., An, D., et al. 2011, *ApJS*, **193**, 29
- Alam, S., Albareti, F. D., Allende Prieto, C., et al. 2015, *ApJS*, **219**, 12
- Andrae, R., Kim, D. W., & Bailer-Jones, C. A. L. 2013, *A&A*, **554**, 137
- Annis, J., Soares-Santos, M., Strauss, M. A., et al. 2014, *ApJ*, **794**, 120
- Assef, R. J., Stern, D., Kochanek, C. S., et al. 2013, *ApJ*, **772**, 26
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, 33
- Bailer-Jones, C. A. L. 2012, *A&A*, **546**, 89
- Bailer-Jones, C. A. L., Smith, K. W., Tiede, C., Sordo, R., & Vallenari, A. 2008, *MNRAS*, **391**, 1838
- Ball, N. M., Brunner, R. J., Myers, A. D., & Tcheng, D. 2006, *ApJ*, **650**, 497
- Bayes, T. 1763, *Philosophical Transactions of the Royal Society of London*, **53**, 370
- Bovy, J., Hennawi, J. F., Hogg, D. W., et al. 2011, *ApJ*, **729**, 141
- Bovy, J., Myers, A. D., Hennawi, J. F., et al. 2012, *ApJ*, **749**, 41
- Butler, N. R., & Bloom, J. S. 2011, *ApJ*, **141**, 93
- Carrasco Kind, M., & Brunner, R. J. 2013, *MNRAS*, **432**, 1483
- Carrasco Kind, M., & Brunner, R. J. 2014, *MNRAS*, **442**, 3380
- Chakraborty, A., Feigelson, E. D., & Babu, G. J. 2013, *PASP*, **125**, 295
- Connolly, A. J., Angeli, G. Z., Chandrasekharan, S., et al. 2014, *Proc. SPIE.*, **9150**, 14
- Croom, S. M., Richards, G. T., Shanks, T., et al. 2009, *MNRAS*, **392**, 19
- Croom, S. M., Smith, R. J., Boyle, B. J., et al. 2004, *MNRAS*, **349**, 1397
- D’Abrusco, R., Longo, G., & Walton, N. A. 2009, *MNRAS*, **396**, 223
- Davoodi, P., Oliver, S., Polletta, M. d. C., et al. 2006, *AJ*, **132**, 1818
- Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2013, *AJ*, **145**, 10
- de Vries, W. H., Becker, R. H., & White, R. L. 2003, *AJ*, **126**, 1217
- de Vries, W. H., Becker, R. H., White, R. L., & Loomis, C. 2005, *AJ*, **129**, 615
- Dexter, J., & Agol, E. 2011, *ApJL*, **727**, L24
- Eisenstein, D. J., Weinberg, D. H., Agol, E., et al. 2011, *AJ*, **142**, 72
- Elvis, M., Civano, F., Vignali, C., et al. 2009, *ApJS*, **184**, 158
- Fan, X., Strauss, M. A., Becker, R. H., et al. 2006, *AJ*, **132**, 117
- Feigelson, E. D., & Babu, G. J. 2003, *Statistical Challenges of Astronomy* (New York: Springer)
- Feigelson, E. D., & Babu, G. J. 2012, *Modern Statistical Methods for Astronomy: With R Applications* (Cambridge: Cambridge Univ. Press)
- Filippenko, A. V. 1982, *PASP*, **94**, 715
- Frieman, J. A., Bassett, B., Becker, A., et al. 2008, *AJ*, **135**, 338
- Fukugita, M., Ichikawa, T., Gunn, J. E., et al. 1996, *AJ*, **111**, 1748
- Gao, D., Zhang, Y. X., & Zhao, Y. H. 2008, *MNRAS*, **386**, 1417
- Gao, D., Zhang, Y. X., & Zhao, Y. H. 2009, *RAA*, **9**, 220
- Giannantonio, T., Scranton, R., Crittenden, R. G., et al. 2008, *PhRvD*, **77**, 12
- Glikman, E., Helfand, D. J., & White, R. L. 2006, *AJ*, **640**, 579
- Graham, M. J., Djorgovski, S. G., Drake, A. J., et al. 2014, *MNRAS*, **439**, 703
- Gray, A., Richards, G., Nichol, R., Brunner, R., & Moore, A. 2005, in *How to do Bayes-Optimal Classification with Massive Datasets: Large-Scale Quasar Discovery*, ed. L. Lyons & M. K. Unel (Oxford: Imperial College Press), 147
- Gunn, J. E., Carr, M., Rockosi, C., et al. 1998, *AJ*, **116**, 3040
- Gunn, J. E., Siegmund, W. A., Mannery, E. J., et al. 2006, *AJ*, **131**, 2332
- Gupta, P., Connolly, A. J., & Gardner, J. P. 2014, *MNRAS*, **439**, 644
- Guy, J., Sullivan, M., Conley, A., et al. 2010, *A&A*, **523**, A7
- Hastie, T., Tibshirani, R., & Friedman, J. 2001, *The Elements of Statistical Learning: Data Mining Inference, and Prediction* (New York: Springer)
- Hennawi, J. F., Myers, A. D., Shen, Y., et al. 2010, *AJ*, **719**, 1672
- Hernitschek, N., Rix, H.-W., Bovy, J., & Morganson, E. 2015, *ApJ*, **801**, 45
- Hinshaw, G., Larson, D., Komatsu, E., et al. 2013, *ApJS*, **208**, 19
- Huff, E. M., Hirata, C. M., Mandelbaum, R., et al. 2014, *MNRAS*, **440**, 1296
- Hughes, P. A., Aller, H. D., & Aller, M. F. 1992, *ApJ*, **396**, 469
- Hunter, J., Firing, E., Caswell, T. A., et al. 2007, *CSE*, **9**, 90
- Ivezić, Ž., Connolly, A. J., VanderPlas, J. T., & Gray, A. 2014, *Statistics, Data Mining, and Machine Learning in Astronomy* (Princeton, NJ: Princeton Univ. Press)
- Ivezić, Ž., Smith, J. A., Miknaitis, G., et al. 2007, *AJ*, **134**, 973
- Ivezić, Ž., Tyson, J. A., Acosta, E., et al. 2008, arXiv:0805.2366
- Jiang, L., Fan, X., Annis, J., et al. 2008, *AJ*, **135**, 1057
- Jiang, L., Fan, X., Bian, F., et al. 2014, *ApJS*, **213**, 12
- Kaczmarek, M. C., Richards, G. T., Mehta, S. S., & Schlegel, D. J. 2009, *AJ*, **138**, 19
- Kaiser, N., Burgett, W., Chambers, K., et al. 2010, *Proc. SPIE*, **7733**, 77330
- Kasliwal, V. P., Vogeley, M. S., & Richards, G. T. 2015, *MNRAS*, **451**, 4328
- Keller, S. C., Schmidt, B. P., Bessell, M. S., et al. 2007, *PASP*, **24**, 1
- Kelly, B. C., Bechtold, J., & Siemiginowska, A. 2009, *ApJ*, **698**, 895
- Kochanek, C. S., Eisenstein, D. J., Cool, R. J., et al. 2012, *ApJS*, **200**, 20
- Koo, D. C., Kron, R. G., & Cudworth, K. M. 1986, *PASP*, **98**, 285
- Kozłowski, S., Kochanek, C. S., Udalski, A., et al. 2010, *ApJ*, **708**, 927
- Krawczyk, C. M., & Peters, C. M. 2014, densityplot: v1.0.0
- Lacy, M., Storrie-Lombardi, L. J., Sajina, A., et al. 2004, *ApJS*, **154**, 166
- Law, N. M., Kulkarni, S. R., Dekany, R. G., et al. 2009, *PASP*, **121**, 1395
- Leistedt, B., Peiris, H. V., Mortlock, D. J., Benoit-Lévy, A., & Pontzen, A. 2013, *MNRAS*, **435**, 1857
- Lilly, S. J., le Fèvre, O., Renzini, A., et al. 2007, *ApJS*, **172**, 70
- Lupton, R. H., Gunn, J. E., & Szalay, A. S. 1999, *AJ*, **118**, 1406
- MacLeod, C. L., Brooks, K., Ivezić, Ž., et al. 2011, *ApJ*, **728**, 26
- MacLeod, C. L., Ivezić, Ž., Kochanek, C. S., et al. 2010, *ApJ*, **721**, 1014
- MacLeod, C. L., Ivezić, Ž., Sesar, B., et al. 2012, *ApJ*, **753**, 106
- Maddox, N., Hewett, P. C., Proux, C., Nestor, D. B., & Wisotzki, L. 2012, *MNRAS*, **424**, 2876
- McGreer, I. D., Jiang, L., Fan, X., et al. 2013, *ApJ*, **768**, 105
- McMahon, R., Banerji, M., Gonzalez, E., et al. 2013, *Msngr*, **154**, 35
- Ménard, B., Scranton, R., Fukugita, M., & Richards, G. 2010, *MNRAS*, **405**, 1025
- Ménard, B., Scranton, R., Schmidt, S., et al. 2013, arXiv:1303.4722
- Mushotzky, R. F., Edelson, R., Baumgartner, W., & Gandhi, P. 2011, *ApJ*, **743**, L12
- Myers, A. D., Brunner, R. J., Richards, G. T., et al. 2006, *ApJ*, **638**, 622
- Oguri, M., Inada, N., Pindor, B., et al. 2006, *AJ*, **132**, 999
- Oke, J. B., & Gunn, J. E. 1983, *ApJ*, **266**, 713
- Page, M. J., & Carrera, F. J. 2000, *MNRAS*, **311**, 433
- Palanque-Delabrouille, N., Yèche, C., Myers, A. D., et al. 2011, *A&A*, **530**, A122
- Papovich, C., Cool, R., Eisenstein, D., et al. 2006, *AJ*, **132**, 231
- Parejko, J. K., Constantin, A., Vogeley, M. S., & Hoyle, F. 2008, *AJ*, **135**, 10
- Pàris, I., Petitjean, P., Aubourg, E., et al. 2014, *A&A*, **563**, A54

- Peng, N., Zhang, Y., Zhao, Y., & Wu, X. b. 2012, *MNRAS*, **425**, 2599
- Pier, J. R., Munn, J. A., Hindsley, R. B., et al. 2003, *AJ*, **125**, 1559
- Pietro Gentile Fusillo, N., Gaensicke, B. T., & Greiss, S. 2015, *MNRAS*, submitted (arXiv:1501.05309)
- Rahman, M., Mnard, B., Scranton, R., Schmidt, S. J., & Morrison, C. B. 2015, *MNRAS*, **447**, 3500
- Richards, G. T., Deo, R. P., Lacy, M., et al. 2009a, *AJ*, **137**, 3884
- Richards, G. T., Fan, X., Newberg, H. J., et al. 2002, *AJ*, **123**, 2945
- Richards, G. T., Myers, A. D., Gray, A. G., et al. 2009b, *ApJS*, **180**, 67
- Richards, G. T., Myers, A. D., Peters, C. M., et al. 2015, *ApJS*, **219**, 39
- Richards, G. T., Nichol, R. C., Gray, A. G., et al. 2004, *ApJS*, **155**, 257
- Richards, G. T., Strauss, M. A., Fan, X., et al. 2006, *AJ*, **131**, 2766
- Richards, G. T., Weinstein, M. A., Schneider, D. P., et al. 2001, *AJ*, **122**, 1151
- Richards, J. W., Starr, D. L., Butler, N. R., et al. 2011, *ApJ*, **733**, 10
- Richards, J. W., Starr, D. L., Miller, A. A., et al. 2012, *ApJS*, **203**, 32
- Riegel, R., Gray, A., & Richards, G. 2008, in *Massive-Scale Kernel Discriminant Analysis: Mining for Quasars*, ed. C. Apte et al. (Atlanta: SIAM), 208
- Ross, N. P., McGreer, I. D., White, M., et al. 2013, *ApJ*, **773**, 14
- Ross, N. P., Myers, A. D., Sheldon, E. S., et al. 2012, *ApJS*, **199**, 3
- Rowan-Robinson, M., Babbedge, T., Oliver, S., et al. 2008, *MNRAS*, **386**, 697
- Salvato, M., Hasinger, G., Ilbert, O., et al. 2009, *ApJ*, **690**, 1250
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, **500**, 525
- Schmidt, K. B., Marshall, P. J., Rix, H.-W., et al. 2010, *ApJ*, **714**, 1194
- Schmidt, K. B., Rix, H. W., Shields, J. C., et al. 2012, *ApJ*, **744**, 2
- Schneider, D. P., Richards, G. T., Hall, P. B., et al. 2010, *AJ*, **139**, 2360
- Scranton, R., Ménard, B., Richards, G. T., et al. 2005, *ApJ*, **633**, 589
- Sesar, B., Ivezić, Ž, Lupton, R. H., et al. 2007, *AJ*, **134**, 2236
- Stern, D., Eisenhardt, P., Gorjian, V., et al. 2005, *ApJ*, **631**, 163
- Suchkov, A. A., Hanisch, R. J., & Margon, B. 2005, *AJ*, **130**, 2439
- Taylor, M. B. 2005, TOPCAT & STIL: Starlink Table/VOTable Processing Software
- Taylor, M. B. 2006, STILTS—A Package for Command-Line Processing of Tabular Data
- The Dark Energy Survey Collaboration 2005, arXiv:astro-ph/0510346
- Vanden Berk, D. E., Schneider, D. P., Richards, G. T., et al. 2005, *AJ*, **129**, 2047
- Vanden Berk, D. E., Wilhite, B. C., Kron, R. G., et al. 2004, *ApJ*, **601**, 692
- Warren, S. J., Hewett, P. C., & Foltz, C. B. 2000, *MNRAS*, **312**, 827
- Weinstein, M. A., Richards, G. T., Schneider, D. P., et al. 2004, *ApJS*, **155**, 243
- York, D. G., Adelman, J., Anderson, J. E., et al. 2000, *AJ*, **120**, 1579
- Zu, Y., Kochanek, C. S., Kozłowski, S., & Udalski, A. 2013, *ApJ*, **765**, 106