Optimizing Photometric Redshift Estimation for Large Astronomical Surveys Using Boosted Decision Trees

by

Adam Joseph Sypniewski

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Physics)
in the University of Michigan
2014

Doctoral Committee:

Professor David W. Gerdes, Chair
Associate Professor Dragan Huterer
Professor Timothy A. McKay
Assistant Professor Christopher John Miller
Professor Mark E. Newman

For my family

Acknowledgements

As it turns out, I don't live in a vacuum. Without the influence, guidance, and support of countless people, I would never have reached this point in my life. Although enumerating these sources in an "acknowledgements" section will never repay my debt, I will pay what tribute I can.

First, and always foremost, I want to thank my family. Particularly my wife, Catalina Cecilia, who has persevered with me through the trials and joys of life, both within academia, but more importantly outside my studies. I would never have attained so many achievements, or properly appreciated the countless blessings, in my life without her to share them with. I am truly honored to have such a wonderful woman at my side for support and love. My daughter, Katherine, has certainly not lent focus to my work, but always brightens my life every evening when I arrive home. My parents, Joe and Karen, have personified encouragement itself, always eager to participate in every walk of my life. My siblings, Lauren, Michael, Joey, and Patrick, continue to provide a steadfast loyalty which I see so rarely in the world.

There are countless friends who have inspired me, given me great insight into many fascinating topics, and supplied a continuous source of banter and entertainment. I have been particularly blessed to have so many people in my life who fall into this category that a full enumeration would be hopeless. However, I want to call out two specific individuals that have gone "above and beyond": Robert Lynch and Joshua Gevirtz. They have influenced my perception and interpretation of the world in drastic ways.

Of course, I have had many good mentors within the academic world, too. Particularly, I would like to thank my advisor, Dave Gerdes, for being precisely the kind of advisor I wanted: insightful, understanding, down-to-earth, and an all-around enjoyable person to work with. Tim McKay, also, was a uniquely fascinating and inspirational individual, who brought a fresh perspective to any topic. Brian, Jax, Tomasz, and Scott were particularly sympathetic fellow students who reliably provided distraction when it was most needed.

And finally, I am unfathomably grateful for all of the blessings and graces that God has bestowed on me. I pray that I can do His will in all things.

Table of Contents

List of Figures

List of Appendices

Appendix

Abstract

Upcoming large-scale sky surveys will obtain photometric data for over 10^8 galaxies. The unprecedented size of such data sets make full spectroscopic followup impossible. Therefore, placing precision constraints on cosmological parameters—such as dark energy—will require accurate redshift estimates based on imaging data alone. In this thesis, we describe a method for estimating photometric redshifts (photo-zs) using boosted decision trees (BDTs), which we call ArborZ. We validate ArborZ and test its performance using simulated galaxy catalogs. After showing that ArborZ is robust with respect to variations between the training and evaluation sets, we apply it to data from two major astrophysical surveys: the Sloan Digital Sky Survey (SDSS) and the Dark Energy Survey (DES). We then develop a method for applying ArborZ to estimate the redshifts of galaxy clusters. We test this in simulated data and then apply it to real data from an XCS-DES cluster catalog.

Chapter I

The Development of Cosmology

For all of its existence, humankind has stared into the night sky and pondered its origin. In an effort to address such existential questions as, "Where has the universe come from?" and, "What is the fate of our universe?" people have developed the branch of natural science known as *cosmology*. Cosmology deals with understanding our universe on its very largest scales. The sheer magnitude of these scales has encouraged the countless explanations and creation stories produced by nearly every civilization. This fusion of existentialism, science, and anthropology makes cosmology a unique and exciting field of study. In this chapter, we review this rich history and development of cosmology, concluding with a summary of our current understanding of our universe.

1.1 Early Cosmology

Even the earliest historical records are teeming with examples of humans trying to understand their universe, or trying to use the universe to understand their environment. Some of the earliest evidence is the Egyptian pyramids at Giza, erected sometime in the middle of the third millennium BC (Spence 2000). These pyramids are remarkable in many ways, perhaps the most astounding of which is the pyramids'

meticulous alignment. By mapping out the circumpolar transit of two stars, the Egyptians oriented the pyramids along true north to within three arcminutes. Over the next millennium, the ancient druids constructed, and subsequently improved, Stonehenge in England. Although the purpose of the site is not definitively understood, a common hypothesis is that it was used for predicting astronomical eclipses (Meadows 1978). Later, in the last millennium BC, the ancient Babylonians produced star catalogs which documented stars, constellations, and planetary motions (Rogers 1998); these catalogs are likely the earliest records we have which evidence a systematic organization of astronomical data. Another example of mankind using the stars for guidance is the proverbial "Three Wise Men" of biblical fame, who, guided by a star, sought out the predicted birth of Christ (*The New American Bible*, Matt. 2.1–12). It wasn't long after Biblical times that the Mayan civilization used the orbit of Venus to develop its famous calendar, known for both its complexity and its long-term accuracy (Thompson 1974). Clearly, even in these early and sometimes prehistoric eras, mankind was ascribing importance to the regularity and precision of the universe.

Most of these examples, however, illustrate various cultures' attempts to understand the cosmos for the purpose of answering astrological questions. The first systematic efforts to apply philosophical and mathematical rigor to the study of astronomy and cosmology were effected by the ancient Greeks. Plato, driven by the simplicity and purity of circles, posited that the orbits of planets must be fundamentally circular and geocentric (Hetherington 2006). His students—such as Eudoxus and Aristotle—tried to realize this by postulating a universe of concentric, rotating spheres. Apollonius and Hipparchus modified this model by introducing epicycles and eccentric orbits which further complicated the motions. These developments were eventually compiled by Ptolemy in his opus *Almagest*, which explained many of the observed planetary motions against a fixed background of stars using concentric

circles and epicycles; this work dominated much of astronomical thinking until the Renaissance (Pannekoek 1961).

One other Greek astronomer deserves notice: Aristarchus of Samos. His work, as described by Archimedes, is the earliest record we have of the development of a heliocentric universe, in sharp contrast to the geocentric view held by the rest of ancient Greece. He estimated the ratio of the diameter of the Earth to the diameter of the Sun and concluded that the Sun must be larger than previously thought; indeed, the Sun would have to be much, much larger than Earth itself! He then proposed a heliocentric universe, with the Earth orbiting around the Sun, with fixed background stars orbiting the Sun as well, but at a tremendous (perhaps infinite) distance (Heath 1913). Aristarchus' universe was much larger than that put forth by other Greek astronomers and mathematicians. This, coupled with the counter-cultural concept of a heliocentric universe, led the heliocentric theory to languish in comparison with more popular geocentric models.

For over a thousand years, our understanding of the universe was in stasis. It wasn't until the Renaissance that Nicholas Copernicus, in 1543, proposed a heliocentric universe (Copernicus 1543). This time, the theory took firm root in the minds of scholars, partially due to the model's success in accurately predicting religious holidays and improving stellar navigation (Erickson 2013; Kuhn 1957). This knowledge proved to be a catalyst, driving the study of astronomy forward at an unprecedented rate. Johannes Kepler, despite his love for the Greeks' obsession with mathematical beauty and the Platonic solids, was able to quantify the motions of planets in his three rudimentary laws (Dreyer 1953). Shortly afterwards, building on Kepler's laws, Isaac Newton published his work on gravity (Newton 1687) and endowed further physical meaning to the planetary motions. The French philosopher René Descartes broadened the scope of the cosmic problem, proposing a universe composed of interlinking solar systems (Descartes 1677). With such a burst of knowledge and interest, a scientific

revolution had most certainly begun.

Perhaps the most revolutionary contributions to astronomy and cosmology prior to the twentieth century were due to William Herschel, the self-taught, "amateur" astronomer who became the "Father of Observational Astronomy" (Mullaney 2007). A Knight of the Royal Guelphic Order and Fellow of the Royal Society, his life is a fascinating apology of the value of the liberal arts. A talented musician, he not only played oboe, harpsichord, and organ, but he also composed twenty-four symphonies during his life. But more famous still are his astronomical discoveries made using his hand-built telescopes. In 1781, Herschel discovered the planet Uranus, and six years later he detected two of its major moons: Titania and Oberon. During his career, he cataloged 2500 nebulae (extended objects) and 806 double stars (Clerke 1908). His catalog was edited by John Dreyer, who supplemented it with other astronomical observations and published it as the *New General Catalogue* (NGC), which is still in common use today (Mullaney 2007). In addition, Herschel was responsible for the discovery of infrared radiation (Herschel 1800). But for cosmologists, Herschel's most important contribution to science was his realization that stars cluster; he is the first to conceptualize the modern notion of a galaxy (Holden 1881). With this discovery, the groundwork for modern cosmology was complete.

1.2 Modern Cosmology

1.2.1 A Dynamic Universe

Given that the human understanding of cosmology had, at the dawn of the twentieth century, taken thousands of years to progress to such rudimentary foundations, it seems astonishing to consider that nearly all of modern cosmology was developed in the course of about eighty years. The breakthroughs began with Einstein's theory of general relativity (Einstein 1916). Einstein realized that space-time and mass-energy

are fundamentally intertwined: as John Wheeler put it, "mass-energy tells space-time how to curve, and curved space-time tells mass-energy how to move" (Ryden 2003). Einstein summarized this succinctly in his eponymous field equations (Einstein 1915; Einstein 1916):

$$G_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu} \tag{1.1}$$

where $T_{\mu\nu}$ is the stress-energy tensor and $G_{\mu\nu}$, called the Einstein tensor, is

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} \tag{1.2}$$

Here, $R_{\mu\nu}$ is the Ricci tensor which encodes information about the curvature of space-time, $g_{\mu\nu}$ is the metric tensor which describes the space-time interval, and R is the Ricci scalar which is simply the trace of the Ricci tensor: $R = g^{\mu\nu} R_{\mu\nu}$.

However, there was a notable problem with this formalism as it has been written, *viz.* that it described an unstable universe, an unusual concept at a time when a steady-state model of the universe was held by most philosophers and scientists. To prevent such a perceived affront to philosophy, Einstein inserted an additional term into the equation, called the cosmological constant Λ, noting that its presence did not violate any physical laws (Einstein 1917). The modified equation became:

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu} \tag{1.3}$$

Einstein eventually recanted this modification, calling it "the biggest blunder of his life" (Gamow 1970).

There were, on the other hand, several scientists who were not quite so hesitant to challenge traditional preconceptions about the nature of the universe. One of the first was Alexander Friedmann, a Russian physicist who began to investigate possible solutions to Einstein's equations for universes of different geometries, and suggested that the universe may indeed have dynamical properties (Friedmann 1922; Friedmann

1924). The startling up-shot of his work was that the geometry and evolution of the universe depends on the energy densities of its constituents, which may themselves evolve with time. This provided a framework in which many questions could finally be scientifically discussed. Was the universe infinite or finite in extent? Would it exist forever? Would it eventually collapse under its own gravity? It suddenly became important to precisely measure the contents of the universe.

Another novel thinker of the time was Georges Lemaître, a Belgian priest and professor, who took this idea of a evolving universe a step further. He realized that any expansions or contractions of space-time would imprint an apparent Doppler shift on propagating photons. In fact, he was the first to experimentally verify this phenomenon (which became known as Hubble's Law) and to discover an expanding universe (Lemaître 1927). But this discovery brought with it a strange implication: if the universe is currently expanding, then long ago it must have been in a smaller, hotter state which he dubbed the "Primeval Atom" and which would eventually be popularized as the Big Bang Theory of the universe (Lemaître 1931b). Unfortunately, the journal in which Lemaître published his findings was not well-read outside of Belgium. Although he eventually published a translation in English (Lemaître 1931a) with the help of Sir Arthur Eddington, his discussion concerning Hubble's Law was inexplicably absent—perhaps because Eddington, though he found Lemaître's work a "brilliant solution" to Einstein's unstable universe (Eddington 1930), was uncomfortable with the philosophical notion of a cosmic origin.

Thus it was that only two years after Lemaître's foundational—but quiet—discovery, another scientist stumbled across the same phenomenon. Edwin Hubble, who had recently gained fame for his definitive discovery of galaxies beyond the Milky Way (Hubble 1926), performed a very similar experiment to Lemaître's. He plotted the recession velocity v of galaxies against their distances r from Earth, as shown in Figure 1.1, using 46 galaxies (Hubble 1929). His result, like Lemaître's, evidenced the

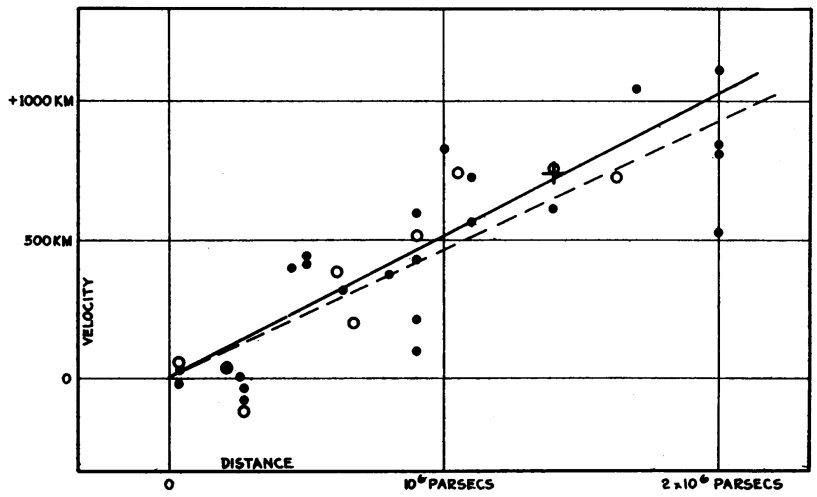

Figure 1.1: The velocity-distance relationship which proved Hubble's Law. The solid points and line are Hubble's fit treating all galaxies with known distances individually. The open points and dashed line are Hubble's fit grouping certain galaxies together. The "+" symbol represents 22 galaxies whose distances could not be estimated individually. Image taken from Hubble 1929.

relationship:

$$v = H_0 r \tag{1.4}$$

where H_0, now known as the Hubble constant, is currently measured to be 67.4 ± 1.4 km/s/Mpc (Ade et al. 2013).

1.2.2 The Discovery of Dark Matter

Another frontier of astrophysical inquiry also sprung up in the 1930s. In 1932, Jan Oort was measuring the velocities of stars in our galaxy when he realized that the mass required to keep them in orbit around the galaxy was larger than the observable mass of galaxy (estimated from the luminosity of stars). He called this unobserved mass "dark matter" and estimated that there had to be at least as much dark matter as regular (baryonic) matter (Oort 1932). The next year, Fritz Zwicky applied the virial theorem to the galaxies in the Coma Cluster and found a similar result to Oort's: the observable mass in the cluster was not large enough to keep the members

gravitationally bound to their cluster (Zwicky 1933). In Zwicky's estimate, the total mass of the system was over 400 times as much as the observable, luminous mass. Zwicky also suggested a clever means of estimating the total mass of a cluster, both observed and dark, by observing the distortion of light rays from background galaxies as they propagate past a massive cluster (Zwicky 1937), a technique known as gravitational lensing.

Despite these findings, did "luminous mass" equal the "dynamical mass" of a system; that is, was it actually possible that a galaxy or cluster have mass which didn't give off light? It would certainly be odd, since this would mean that this dark matter did not interact electromagnetically, unlike most forms of matter known at the time. It wasn't until 1970 that the existence of dark matter was soundly proven by Vera Rubin using rudimentary dynamics. By simply measuring the rotational velocities of stars in a galaxy, once could infer the underlying mass profile $M(r)$ of the galaxy by equating centripetal force F_c and gravitational force F_g:

$$G\frac{M(r)m}{r^2} = F_g = F_c = m\frac{v^2}{r} \tag{1.5}$$

$$v = \sqrt{\frac{GM(r)}{r}} \qquad \Longleftrightarrow \qquad M(r) = \frac{v^2 r}{G} \tag{1.6}$$

where v is the velocity of a start at a distance r from the center of galaxy, $M(r)$ is the total mass of the galaxy within a radius r, and m is the mass of the star. Since $\lim_{r\to\infty} M(r) \equiv M_{\text{gal}}$ is finite, the velocity distribution v should drop off once most of the galaxy's mass is enclosed in a certain radius r. However, instead of the expected drop off, Rubin observed (Rubin and W. Kent Ford 1970) v to flatten out at a non-zero velocity for moderately large values of r before finally dropping off to zero, far after all observed mass was enclosed (see Figure 1.2). She estimated that the total mass must be at least three times as large as that inferred from luminous matter alone (Rubin, Burstein, et al. 1985). This was the strongest evidence for the existence of

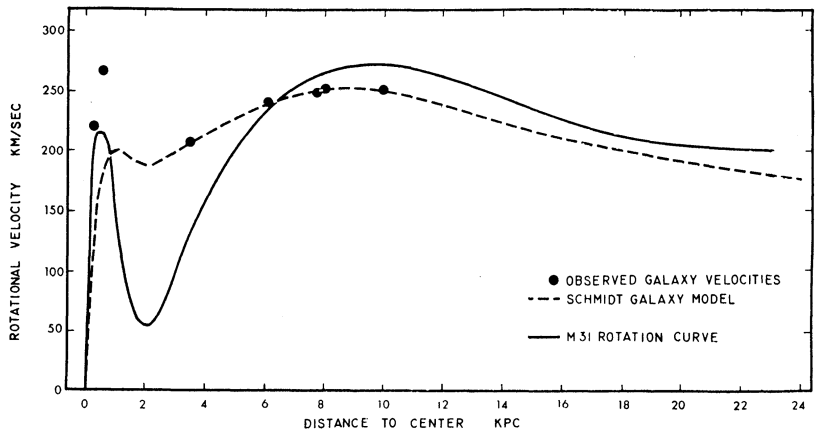

Figure 1.2: The distribution of rotational velocities of stars in galaxy M31 (Andromeda). Rather than quickly dropping off, the velocity profile extends outward, demonstrating the existence of galactic matter even at large radii. Image taken from Rubin and W. Kent Ford 1970.

dark matter yet discovered.

In fact, dark matter very conveniently explained an outstanding problem in galaxy formation. Ostriker and Peebles, while studying early simulations of galaxy evolution, noticed that galaxies seemed "grossly unstable to barlike modes" (Ostriker and Peebles 1973). They showed that if the ratio of a galaxy's total kinetic energy to its total gravitational energy is too large, then mass over-densities will drive the stars to collapse rapidly into a bar. The critical value for this ratio, called the Ostriker-Peebles criterion, is approximately 0.15. If, however, this ratio can be reduced by adding additional gravitational energy in the form of a halo, galaxy formation remains stable. Dark matter provided the perfect candidate for this halo, solving the problem of galaxy formation.

There have been several candidate explanations for dark matter. One theory is Modified Newtonian Dynamics (MoND), which posits that Newton's law of gravity breaks down and requires modification at large scales (Milgrom 1983a; Milgrom 1983b). Another possibility is that the dark matter is simply neutrinos, a naïvely plausible explanation, since neutrinos don't interact electromagnetically. Such theories are called "hot dark matter" theories, since neutrinos are so light as to be

relativistic. However, this, in turn, prevents them from effectively collapsing into halos and facilitating galaxy formation (M. Davis et al. 1992; S. D. M. White et al. 1993). The generally accepted explanation is called "cold dark matter," which posits that some weakly-interacting, massive particle (WIMP) is responsible, with candidate particles often drawn from supersymmetric theories (Harper 2013).

1.2.3 The Cosmic Microwave Background

It was at a similar time in history that one of the most important discoveries in modern cosmology occurred quite serendipitously. Penzias and Wilson, two young radio astronomers, were calibrating out noise sources on the 20ft horn antenna (Crawford, D. C. Hogg, and Hunt 1961) at Bell Labs. However, they were met with an unexpected and irreducible noise source at a temperature of $3.5 \pm 1\,\mathrm{K}$ (Penzias and Wilson 1965), which they measured to be constant across the sky.

Meanwhile, a group of four physicists at nearby Princeton University were developing novel ways to test for the existence of Lemaître's Big Bang. They reasoned that the universe at its earliest epoch must have been extremely hot and in thermal equilibrium characterized by a black-body spectrum. As the universe expanded and cooled, this black-body radiation would necessarily cool as well and should still be observable as a microwave background radiation. When news of Penzias and Wilson's discovery reached Princeton, it was hailed as the long-awaited evidence of a Big Bang (Dicke et al. 1965), and now goes by the name of the cosmic microwave background (CMB) radiation.

The CMB has been been re-measured many times since its initial discovery in order to improve the spatial resolution of the temperature field. Probably the most influential re-measurement was taken by the Cosmic Background Explorer (COBE), launched in 1989. COBE tightened the measurement of the mean (monopole) temperature to $2.728\,\mathrm{K}$. This black-body spectrum, shown in Figure 1.3, is one of the

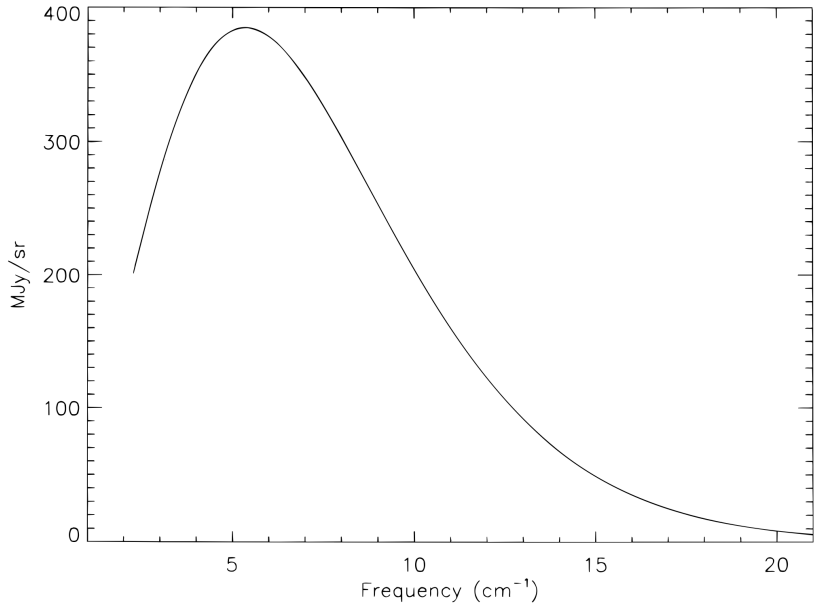

Figure 1.3: The CMB black-body spectrum, measured by COBE, with a temperature of 2.728 K. It ranks among one of the greatest examples of agreement between theory and experiment in the history of science: the errors are a small fraction of the line width. Image taken from Fixsen et al. 1996.

best experimental fits to theory in the history of science! COBE, with its spatial resolution of 7°, also detected a dipole term on the order of $\Delta T/T \approx 10^{-3}$ corresponding to our peculiar velocity through the CMB rest frame. More importantly, however, COBE was able to detect quadrupole anisotropies in the CMB spectrum on the order of $\Delta T/T \approx 10^{-5}$ (Smoot et al. 1992). These fluctuations were evidence for primordial density fluctuations, driven by quantum mechanics, in the dark matter power spectrum (Bardeen, Steinhardt, and Turner 1983; Bond and Efstathiou 1987). As the universe expanded, these fluctuations—called the *baryon acoustic oscillations*—would grow, seeding the large-scale structure we see in our universe today. Thus, this discovery was proof that our understanding of the Big Bang and growth of structure was on the right track. It earned Smoot and Mather the Noble Prize in Physics in 2006, with the committee citing the discovery as "the inception of cosmology as a precision science" (Nobel Media AB 2006).

The discovery of CMB anisotropies sparked a new front of experiments. One of

these experiments is the Wilkinson Microwave Anisotropy Probe (WMAP), launched in 2001 (Spergel et al. 2003). This satellite was critical in the development of our modern concordance cosmology, called ΛCDM. It was 45 times more sensitive than COBE and its 13′ resolution drastically improved our measurements of the CMB anisotropies. The most recent space-based CMB experiment was Planck, which was launched in 2009. Planck's sensitivity was 10 times greater than WMAP and its angular resolution was 5′. It has provided some of the strongest experimental constraints on the nature of our universe in the history of cosmology, improving our estimate of the temperature of the CMB to $2.7255 \pm 0.0006\,\mathrm{K}$ (Ade et al. 2013).

The discovery CMB also offered challenges to cosmologists. Particularly troubling were the so-called *horizon problem* and *flatness problem*. The horizon problem asked: how could the CMB be in thermal equilibrium if different parts of the sky were only now coming into causal contact? The flatness problem was slightly more subtle: Friedmann's equations would suggest that a universe with a curved geometry would quickly grow more curved. We therefore have a fine-tuning problem: in order for the observable universe to be flat today—as suggested by CMB experiments—it must have been flat to within one part in 10^{-60} in the past.

The solution to both problems was a period of phenomenally rapid expansion in the early universe, called *inflation*, originally proposed by Guth (1981). The inflationary era was driven by a particle, called the *inflaton*, which had a slow-rolling potential (Linde 1982); that is, the inflaton scalar field moves only very slowly along its nearly flat potential. This would cause an exponential expansion of the universe, driving the universe rapidly toward a flat geometry, and removing causal contact between regions previously in thermal equilibrium. The inflaton would then decay into known Standard Model particles, ending the inflationary era.

1.2.4 Dark Energy

The twentieth century witnessed the most incredible progress ever made in cosmology, but there was still one more secret to be discovered. Already there were hints of something more to the structure of the universe. Until the end of the 1990s, the common belief was that our universe was dominated by matter. Moreover, our universe appeared to have a flat geometry. According to Friedmann's equations, a matter-dominated, flat universe would be 9.3 Gyr old; however, globular clusters—among the oldest objects in the universe—had already been discovered whose ages were estimated at 12 Gyr (Hansen et al. 2002). A new understanding of, or discovery in, cosmology was necessary to explain this contradiction.

A simple, yet elegant, mechanism for testing our understanding of cosmology was already understood. Since light travels at a finite speed, observing more distant astronomical objects is equivalently viewing the universe at an earlier epoch: telescopes are time machines. Probing the physical properties of galaxies at different distances from us is therefore tantamount to examining the evolution of the universe. Since, as Friedmann discovered, the dynamical properties of the universe are intimately tied to its contents, we can begin to understand the constitution and fate of the universe by measuring distances.

As it turns out, however, and as we will discuss in greater detail later, estimating the distances to astronomical objects is an extremely difficult and error-fraught field. Distances have traditionally been extrapolated using the *cosmic distance ladder*, a method of successive estimation techniques, each "rung" of which is calibrated using previous methods on the ladder. Systematic errors are therefore propagated through each rung of ladder, muddling the strength of the cosmological signal.

There was, however, an alternative method. It had been known since the 1960s that type Ia supernovae (SNe Ia) were approximately standard candles; that is, SNe

Ia all explode with the same intrinsic luminosity (Kowal 1968; Pskovskii 1967). This implied that measurements of SNe Ia could yield accurate distance measurements, constraining cosmological parameters and perhaps granting insight into the contents and fate of the universe. Moreover, relative distances could be measured without ever needing to determine the absolute luminosity of SNe Ia, reducing errors and nuisance parameters. On the other hand, supernovae were notoriously hard to discover at the time: there are, on average, only twenty SNe per millennium in our Galaxy, only four of which are SNe Ia (Cappellaro et al. 1997; Türler 2006). Widening the search to extragalactic SNe requires imaging the same area of the sky every few days to detect the birth of these rare and elusive SNe.

Despite the challenges, two audacious, independent teams dared to attempt to measure SNe Ia *en masse*: the High-Z Supernovae Search Team and the Supernovae Cosmology Project. By careful scheduling of telescope time with follow-up imaging and spectroscopy, each team was able to collate about forty SNe Ia. The results were startling: the more distant SNe appeared dimmer than one would expect in a matter-dominated universe (S. Perlmutter 2003; S. Perlmutter et al. 1999; Riess et al. 1998). This could be explained if—contrary to popular belief at the time—the universe were not, in fact, expanding and decelerating (due to the attracting force of gravity). Rather, the universe's expansion would have decelerated in the past, but started accelerating in its most recent epoch. This accelerating expansion would need to be driven by some form of energy with a negative equation of state (a negative pressure), which Michael Turner nicknamed *dark energy* (D. Huterer and Turner 1999).

To add to the surprising discovery of dark energy, the analyses of the two supernovae teams supported a cosmology where dark energy composed an astonishing 75% of the total mass-energy density of the universe. The simplest explanation for dark energy would be Einstein's cosmological constant. Physically, this would imbue space

itself with a non-zero vacuum energy density, and so as the universe expanded from the Big Bang, more space would appear, bringing with it more dark energy. The dark energy, in turn, would continue to drive expansion. Other viable explanations exist, however, such as quintessence (Zlatev, Wang, and Steinhardt 1999), though most experimental constraints suggest the more mundane explanation.

1.3 Concordance Cosmology

Cosmologists have not had a lot of time to consider the new developments outlined in the previous section, but the result of these deliberations is a *concordance* cosmology, that is, the generally accepted understanding of the universe's content and past evolution. In order to give perspective to the discoveries of the past century, I present a concise timeline of the known history of the universe. As in most discussions of cosmology, there are two common alternative representations of time. The first is called *redshift*, denoted z, and defined as the fractional shift in the wavelength of light caused by cosmic expansion:

$$z = \frac{\Delta\lambda}{\lambda} = \frac{\lambda_{\text{observed}} - \lambda_{\text{emitted}}}{\lambda_{\text{emitted}}} \tag{1.7}$$

The other representation is the *scale factor*, $a(t)$, which measures the relative size of the universe, normalized to be 1 today.

1.3.1 Big Bang and Inflation

Approximately 13.7 Gyr ago, the universe and space-time as we know it were in a hot, dense state. All particles were in thermal equilibrium and their density fields were almost perfectly uniform. In these first instants of the universe's existence, inflation occurred, driving a tremendous, exponential expansion of the universe. At the end of inflation (which lasted only microseconds), the expansion of the universe began to

decelerate under the action of gravity.

1.3.2 Big Bang Nucleosynthesis

During the first three minutes after the inflation, the universe continued to expand and cool. This cooling allowed the formation of ionized nuclei (in addition to hydrogen). The nuclear cascade of reactions was already understood by Alpher and his student Gamow in the 1940s (Alpher, Bethe, and Gamow 1948). Nucleosynthesis predicts the relative abundance of chemical elements with surprising precision[1], with ${}^{1}_{1}\mathrm{H}$ constituting 75% of the mass abundance of baryonic matter; ${}^{4}_{2}\mathrm{He}$ constituting most of the remaining 25%; ${}^{2}_{1}H$ (deuterium) contributing 0.01%; and trace amounts of lithium and beryllium.

1.3.3 Baryon Acoustic Oscillations

As we stated before, at this point in its existence the universe was remarkably smooth and uniform. Quantum mechanical fluctuations, however, upset this uniformity, producing small over- and under-densities in the matter power spectrum (Bond and Efstathiou 1984; D. J. Eisenstein and W. Hu 1998; Weinberg et al. 2012). The evolution of these density fluctuations depended strongly on the physical interactions of each component in the perturbation.

Initially, the dark matter remains in the initial perturbation, since it will only interact gravitationally and has little intrinsic kinetic energy. On the hand, neutrinos, begin too hot and fast to be gravitationally bound to the perturbation, begin free-streaming away. The baryonic gas is hot enough to be completely ionized, and thus is strongly coupled to the photons. The resulting baryon-photon fluid will produce a large overpressure and, in an attempt to equalize, begin expanding as a spherical

[1]Big bang nucleosynthesis only produces these lightest elements. All heavier elements in the universe are created by stars.

sound wave. Finally, since the perturbation itself is dense, it will slowly begin to grow as additional dark matter falls into the perturbation under the influence of the gravitational instability. See Figure 1.4a.

As the baryon-photon fluid contains to expand, the size of the density perturbation will grow, widening the density profile of the dark matter (Figure 1.4b). As the fluid expands, it cools, and eventually the baryons can capture electrons. This is the era of *recombination*, approximately $240,000$ years after the Big Bang, at a redshift of $z \approx 1370$. Once the baryons become neutral, the photons free stream away (the *last scattering*), producing the cosmic microwave background, which occurs at $z \approx 1100$, approximately $350,000$ years after the Big Bang (Figure 1.4c).

Meanwhile, without the baryon-photon fluid providing an overpressure, the sound wave has stalled, leaving a shell of baryons at a characteristic distance from the initial density perturbation (see Figure 1.4d). With time, the interaction between the baryons and dark matter smooths out this discrepancy; however, the characteristic baryonic and dark matter overdensities are never complete erased (Figures 1.4e and 1.4f).

These overdensities amplify the growth of large-scale structure, and continue to be visible today in the form of a peak in the galaxy *two-point correlation function*, the excess probability of two galaxies clustering as a function of separation. This peak (Figure 1.5a) is called the *baryon acoustic peak*, and when viewed in Fourier space (i.e., in the matter power spectrum) are seen as a series of peaks called the *baryon acoustic oscillations* (BAO; see Figure 1.5b). By measuring the scale of this characteristic separation, constraints can be placed on the expansion history of the universe and, therefore, on the equation-of-state of dark energy. These perturbations have been observed both in CMB (by the experiments described earlier) as well as in the the statistical distribution of galaxies on the sky (D. J. Eisenstein, Zehavi, et al. 2005), and BAO measurements continue to be an active area of cosmology.

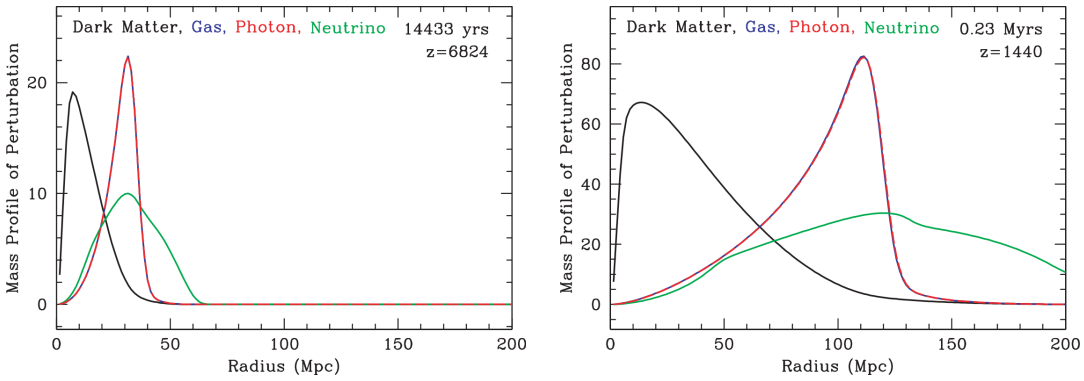

(a) $z = 6824$

(b) $z = 1440$

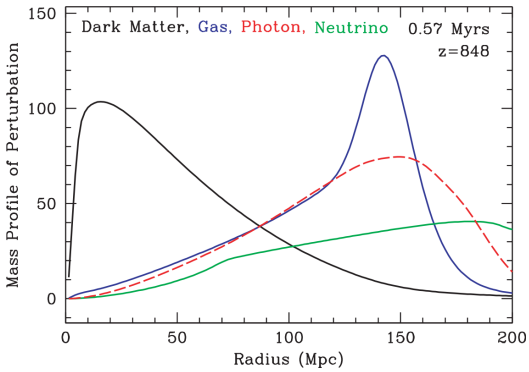

(c) $z = 848$

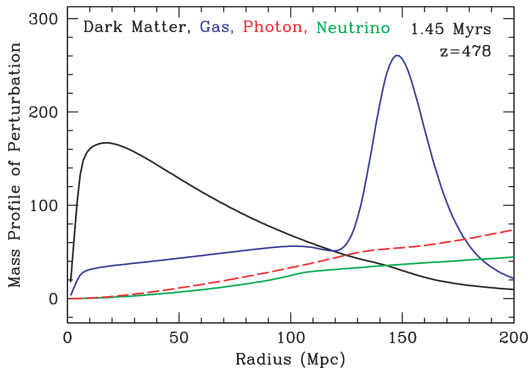

(d) $z = 478$

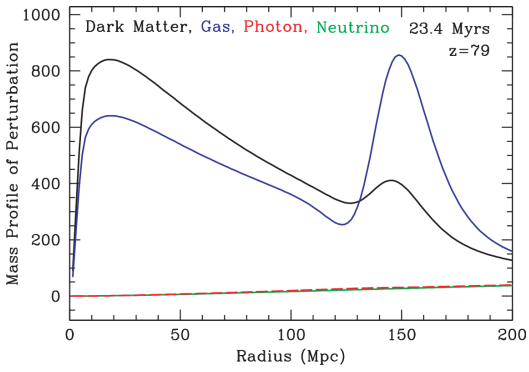

(e) $z = 79$

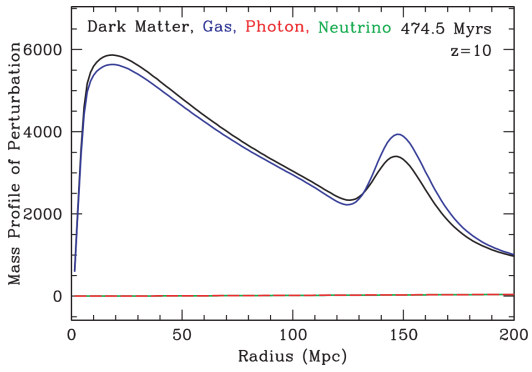

(f) $z = 10$

Figure 1.4: Evolution of the baryon acoustic oscillations (BAO), from early times (high redshift) to more recent times (low redshift). Plots taken from D. J. Eisenstein, Seo, and M. White 2007.

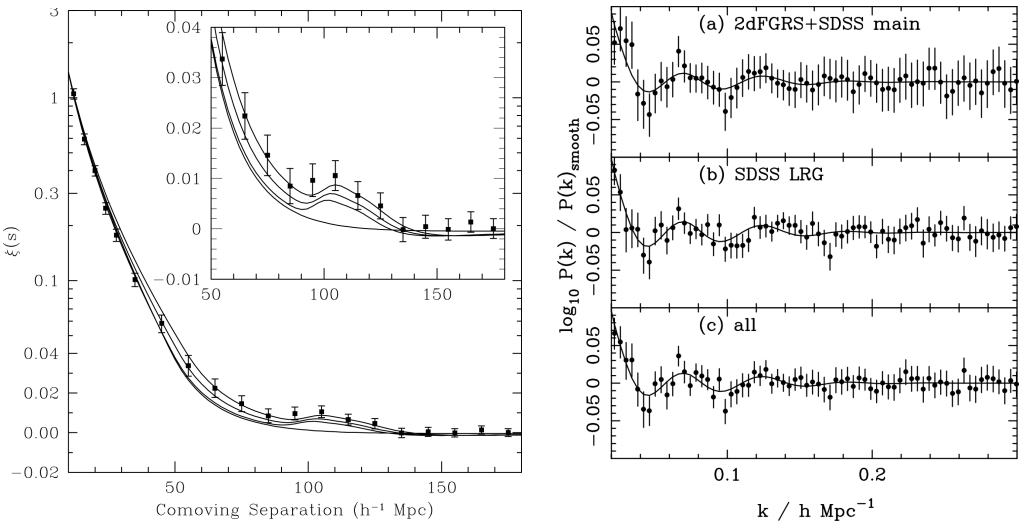

(a) Two-point correlation function (b) Matter power spectrum

Figure 1.5: The BAO peaks manifest themselves as a peak in the galaxy two-point correlation function as well as a series of peaks in the matter power spectrum. Figure 1.5a is taken from D. J. Eisenstein, Zehavi, et al. 2005, and Figure 1.5b is taken from Percival et al. 2007.

1.3.4 The Dark Ages

At this point in its history, the universe entered a "dark age" where it had little visible structure, since stars and galaxies had not formed *en masse*. Inexorably, gravity drove the gravitational collapse of hydrogen, and approximately 100–200 million years after the Big Bang, the universe produced its first stars (Bromm and Larson 2004; Bromm and Loeb 2006) and the dark ages ended.

1.3.5 A Dark Energy Dominated Universe

After the dark ages, the universe continued to be in a matter-dominated state, and density perturbations grew linearly with the scale factor. The large-scale structure formed and evolved, producing the exquisite, complex cosmic web (Bond, Kofman, and Pogosyan 1996) that we observe today (see Figure 1.6).

Meanwhile, dark energy's contribution to the mass-energy density of the universe is increasing, seemingly due to the expansion of the universe. Since dark energy

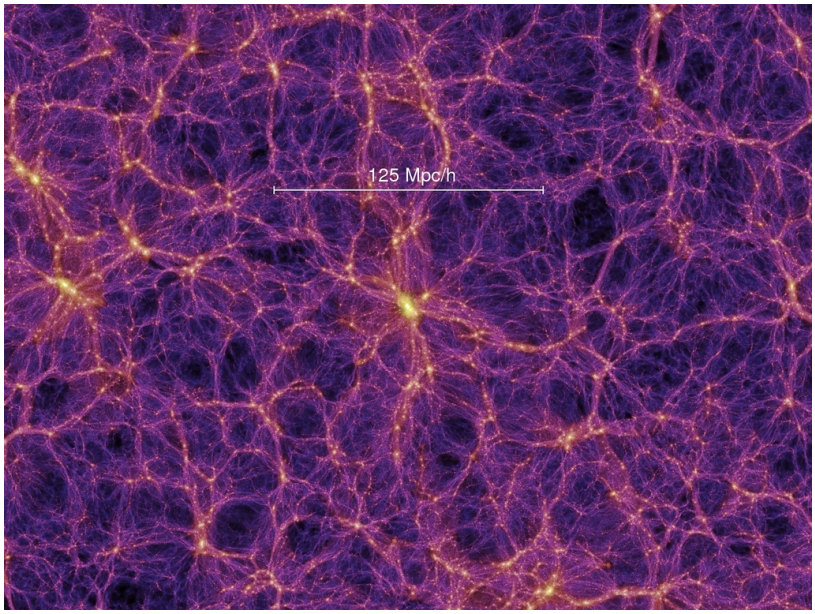

Figure 1.6: The cosmic web, tracing dark matter from an N-body simulation. The cosmic web is the result of structure growth and evolution over billions of years, driven by gravitational instabilities. Image taken from Springel 2004.

acts to drive an accelerating expansion, the universe eventually reached an inflection point where it switched from decelerating (under gravity's influence) to accelerating (due to dark energy). This matter-dark energy equality occurred at $z \approx 0.44$, five billion years ago. We find ourselves, on cosmic scales, very close to the matter-dark energy equality, a curious state called the "coincidence problem" (Carroll 2001). Some attempt to answer this question by positing the *anthropic principle*: if the universe were any other way, we would not be here to observe it. Regardless of one's philosophical bent, however, it appears that dark energy will continue to be an ever-increasing component in the universe. Once it dominates the mass-energy of universe, the growth of structure will shut off entirely.

Chapter II

Cosmic Observables

All of science crucially hinges on the scientist's ability to measure and observe the environment. The concordance cosmology describes a complex universe; how can we best quantify our observations? This chapter introduces *redshift* as one of the fundamental cosmic observables and highlights its importance in modern astrophysical experiments.

2.1 Theoretical Underpinnings

In order to understand our universe, we need to probe its dynamic properties. As we discussed in the previous chapter, one of the most elegant ways for doing this is to realize that, since light travels a finite speed, observing more distant objects is tantamount to probing physics at earlier epochs. Thus, regardless of the details of exactly which phenomena one may want to probe (e.g., clusters, lensing, supernovae), the dynamic properties of the universe can be extracted by measuring how a physical quantity varies with distance. Then the only missing step is the conversion from distance back into time.

2.1.1 Einstein's Field Equations

In order to bridge this gap, we start by turning to Einstein's field equations for general relativity:

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = \frac{8\pi G}{c^4} T_{\mu\nu} \tag{2.1}$$

The first term we address in this equation is the metric tensor. When we look into the night sky, we see that we live in an extraordinarily uniform universe; moreover, there is no known asymmetry which should convince us otherwise. The fact that the universe—on its largest scales—is both homogeneous and isotropic is often called the *cosmological principle*. As a consequence, the universe must have some sort of spherical geometry (the only geometry which preserves such a symmetry), and so the metric tensor (in Cartesian coordinates) must be of the form[1]:

$$g_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{2.2}$$

with coordinates (t, x, y, z). This corresponds to the following four-dimensional space-time interval (t, r, θ, ϕ) in spherical coordinates[2]:

$$ds^2 = -c^2 dt^2 + a^2(t) \left[dr^2 + S_\kappa(r) d\Omega^2 \right] \tag{2.3}$$

where

$$d\Omega^2 = d\theta^2 + \sin^2\theta d\phi^2 \tag{2.4}$$

is the differential solid angle contribution,

[1] Strictly speaking, an overall factor of -1 can be applied to the metric tensor without changing any of the physics. The choice is merely a matter of convention.

[2] Following the convention of physicists, ϕ is the azimuthal angle (angle around $\hat{z}$) and θ is the polar angle / co-latitude (angle from $\hat{z}$).

$$S_\kappa(r) = \begin{cases} R\sin(r/R) & \kappa = +1 \quad \text{(positive curvature; spherical, closed universe)} \\ r & \kappa = 0 \quad \text{(zero curvature; flat universe)} \\ R\sinh(r/R) & \kappa = -1 \quad \text{(negative curvature; hyperbolic, open universe)} \end{cases} \tag{2.5}$$

is the circumferential radial coordinate[3], R is the radius of curvature of the universe (as measured today), and the scale factor $a(t)$ is the ratio of the size of the universe at time t to the size of the universe now ($t = t_0$).

Note that some authors choose to write the space-time interval in terms of the circumferential radial coordinate $x \equiv S_\kappa(r)$. In this case, we have:

$$dx = \left\{ \begin{array}{cc} \cos(r/R) & \kappa = +1 \\ 1 & \kappa = 0 \\ \cosh(r/R) & \kappa = -1 \end{array} \right\} dr \tag{2.6}$$

and

$$r = \left\{ \begin{array}{cc} R\sin^{-1}(x/R) & \kappa = +1 \\ x & \kappa = 0 \\ R\sinh^{-1}(x/R) & \kappa = -1 \end{array} \right\} \tag{2.7}$$

Therefore:

$$dx^2 = \left\{ \begin{array}{cc} \cos^2\left[\sin^{-1}(x/R)\right] & \kappa = +1 \\ 1 & \kappa = 0 \\ \cosh^2\left[\sinh^{-1}(x/R)\right] & \kappa = -1 \end{array} \right\} dr^2 = \left\{ \begin{array}{cc} 1-(x/R)^2 & \kappa = +1 \\ 1 & \kappa = 0 \\ 1+(x/R)^2 & \kappa = -1 \end{array} \right\} dr^2 \tag{2.8}$$

This allows us to write:

$$dr^2 = \frac{dx^2}{1 - \kappa x^2/R^2} \tag{2.9}$$

Thus, the FRW metric can be written as:

[3] r is called the geodesic radial coordinate. This distinction becomes necessary when generalizing spherical coordinates into curved space, since the coordinate r corresponding to intuitive radial distance from the origin is not the same radius which would yield our intuitive measure of circumference ($S_\kappa(r)$).

$$ds^2 = -c^2dt^2 + a^2(t)\left[\frac{dx^2}{1-\kappa x^2/R^2} + x^2 d\Omega^2\right] \tag{2.10}$$

Very often, often authors will define:

$$k = \kappa/R^2 \tag{2.11}$$

and write the metric as:

$$ds^2 = -c^2dt^2 + a^2(t)\left[\frac{dx^2}{1-kx^2} + x^2 d\Omega^2\right] \tag{2.12}$$

Moreover, at this point some authors will choose to write r instead of x, so the reader must be aware that such an r is what we are calling $x = S_\kappa(r)$. Again, both are valid radial coordinates, but they have very different physical meanings. Still other authors will use the scaled circumferential radius $\rho = x/R^2$ and explicitly use $R^2(t)$ outside the brackets instead of $a^2(t)$. *Caveat scientificus.*

It is also worth clarifying some of our assumptions. First is the cosmological principle: after all, it is obvious that, particularly on small scales, the universe is neither homogeneous nor isotropic. However, on scales greater than approximately 100 Mpc, space-time is fairly uniform, as evidenced by experiments such as CMB measurements.

Now that we have a handle on the metric tensor, we need to return our attention to Einstein's equations. A perfect fluid—one without shear stresses, viscosity, or heat conduction—can be fully described using only its energy density[4] ρ and isotropic pressure p. If we model the large-scale universe as a perfect fluid, then the stress-energy tensor (in Cartesian coordinates) is:

$$T_{\mu\nu} = \begin{pmatrix} -\rho & 0 & 0 & 0 \\ 0 & p & 0 & 0 \\ 0 & 0 & p & 0 \\ 0 & 0 & 0 & p \end{pmatrix} \tag{2.13}$$

[4]The mass density would simply be ρc^2.

What of the remaining terms in the Einstein equation? The Ricci tensor can be written in terms of the Christoffel symbols $\Gamma^{\sigma}_{\mu\nu}$ (the affine connection of metric):

$$R_{\mu\nu} = R^{\gamma}_{\mu\gamma\nu} = \partial_{\gamma}\Gamma^{\gamma}_{\mu\nu} - \partial_{\nu}\Gamma^{\gamma}_{\mu\gamma} + \Gamma^{\gamma}_{\sigma\gamma}\Gamma^{\sigma}_{\mu\nu} - \Gamma^{\gamma}_{\sigma\nu}\Gamma^{\sigma}_{\mu\gamma} \tag{2.14}$$

where the differential geometer may identify $R^{\mu}_{\nu\alpha\beta}$ as the Riemann tensor. Computing the various components of the Ricci tensor is an exercise best for the reader, as it leaves one with as distinct and uncomfortable a memory as Hercules must have experienced for days after cleansing the Augean stables. Meanwhile, the Ricci scalar is merely the trace of the Ricci tensor—$R = g^{\mu\nu}R_{\mu\nu}$—and is a stroll through the park in comparison.

Suffice it to say that once these various quantities are fed back into the Einstein equation and the metaphorical crank is turned, two independent equations drop out:

$$\left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3c^2}\rho + \frac{1}{3}\Lambda c^2 - \frac{c^2\kappa}{R^2a^2} \tag{2.15}$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi G}{3c^2}\left(\rho + 3p\right) + \frac{1}{3}\Lambda c^2 \tag{2.16}$$

where the over-dot ($\dot{a}$, $\ddot{a}$) indicates a time derivative, R is the radius of curvature of space (as measured today), and ρ is the energy density of the universe contained in matter and radiation.

These are the two Friedmann equations[5], and they govern the dynamics of the universe on its largest scales. Immediately we see what Friedmann must have seen back in the 1920s: the evolution of the universe is intimately related to what is in it. In addition, we can link this back to Lemaître and Hubble's work by identifying the Hubble parameter as:

$$H = \frac{\dot{a}}{a} \tag{2.17}$$

[5]The first equation is often referred to simply as the Friedmann equation, and the second equation is sometimes called the Friedmann acceleration equation.

The Hubble constant is then $H_0 = H(t = t_0)$, which is simply the Hubble parameter measured today, when the universe's age is t_0.

It is important to note that, although we have dragged the cosmological constant Λ along through these equations, we are not yet certain if dark energy is, in fact, caused by Λ. If we eventually learn that dark energy has dynamical properties not captured by a vacuum energy, we may still be able to remove Λ and subsume dark energy's contribution to the energy density of the universe into ρ. However, for more complicated theories, such as modified gravity (Milgrom 1983a), the solution may not longer be so simple.

2.1.2 The Continuity Equation

In addition to the two Friedmann equations, there is a third, dependent equation which can easily be derived. By differentiating Equation 2.15 and using it to eliminate $\ddot{a}$ in Equation 2.16, one obtains the *continuity equation*:

$$\dot{\rho} = -3H(\rho + p) \tag{2.18}$$

We now introduce an equation-of-state which relates pressure to density via an equation-of-state parameter w, which may be a function of time:

$$p = w\rho \tag{2.19}$$

To make this more convenient, we can change variables:

$$\frac{d}{dt} \longrightarrow H\frac{d}{d\ln a} \tag{2.20}$$

Then, by combining Equations 2.18 and 2.19, we have:

$$\frac{d\ln\rho}{d\ln a} = -3(1 + w) \tag{2.21}$$

The straightforward solution is:

$$\rho = \rho_0 a^{-3(1+w)} \tag{2.22}$$

where ρ_0 is the energy density measured today. This equation gives us the evolution of a component's energy density given its equation-of-state parameter, which we will use in the next section.

2.1.3 Components of the Universe

We want to utilize the continuity equation in order to manipulate the Friedmann equations into a more usable form. To do this, we begin by breaking the energy density ρ in Equation 2.15 into its matter and radiation contributions:

$$H^2 = \frac{8\pi G}{3c^2}\left(\rho_M + \rho_R\right) + \frac{1}{3}\Lambda c^2 - \frac{c^2\kappa}{R^2 a^2} \tag{2.23}$$

Then, if we choose to treat the dark energy contribution and curvature terms merely as additional energy densities, we have:

$$H^2 = \frac{8\pi G}{3c^2}\left[\rho_M + \rho_R + \rho_{DE} - \rho_K\right] \equiv \frac{8\pi G}{3c^2}\left[\rho_{\text{total}} - \rho_K\right] \tag{2.24}$$

where the energy density contributed by dark energy is[6]:

$$\rho_{\text{DE}} = \frac{c^4\Lambda}{8\pi G} \tag{2.25}$$

and the equivalent energy density that we can consider curvature contributing is:

$$\rho_K = \frac{3c^4\kappa}{8\pi G R^2 a^2} \tag{2.26}$$

Clearly, there must exist an energy density which would require a flat geometry for universe (that is, an energy density which corresponds to $\kappa = 0$). This density is called the *critical density*, and can be read off the above equations:

[6]As we mentioned before, dark energy may not be compatible with a cosmological constant. If this is the case, the previous arguments still hold: we can express the overall energy density explicitly in terms of its constituent contributions, including a dark energy component ρ_{DE}.

$$\rho_{\text{crit}} = \frac{3c^2H^2}{8\pi G} \tag{2.27}$$

Modern cosmology most frequently quotes densities in terms if the critical density, since it defines a convenient scale for comparing densities and provides a unitless quantity to work with. For example, the *density parameter* $\Omega(t)$ is the ratio of the actual density of the universe to its critical density:

$$\Omega(t) = \frac{\rho_{\text{total}}(t)}{\rho_{\text{crit}}(t)} \tag{2.28}$$

Similarly, density parameters can be defined for any component (i.e., matter, radiation, dark energy, or curvature) in the universe:

$$\Omega_i(t) = \frac{\rho_i(t)}{\rho_{\text{crit}}(t)} \tag{2.29}$$

It is common to write these Ω_i without the explicit functional format when referring to quantities as measured today; thus, $\Omega_M \equiv \Omega_M(t = t_0)$, where t_0 is the age of the universe.

We can use the density parameters to quickly learn about the curvature of spacetime. First, we see that we can rewrite Equation 2.24 as:

$$\rho_{\text{crit}} = \rho_{\text{total}} - \rho_K \tag{2.30}$$

Therefore, we can use Equations 2.28 and 2.29 to write a remarkably simple version of the Friedmann equation:

$$1 = \Omega(t) - \Omega_K(t) = [\Omega_M(t) + \Omega_R(t) + \Omega_{DE}(t)] - \Omega_K(t) \tag{2.31}$$

Using the definition of Ω_K and evaluating the previous equation at $t = t_0$ we can therefore write:

$$1 - \Omega = -\frac{c^2\kappa}{H_0^2R^2} \tag{2.32}$$

This equation allows us to calculate the geometry of space time (i.e., the sign of κ) using only the energy density of space and Hubble's constant.

However, since we are more interested in the dynamic properties of the universe, we return to Equation 2.22 and discuss the time-evolution of each component. Each component may have a different equation-of-state, and so we address each in turn.

- **Matter**. The matter density (both baryonic and dark) will scale in the intuitive sense: $\rho_M(t) \propto a^{-3}$; that is, its density is proportional to the volume of space it is in. The equation-of-state parameter for matter must therefore be $w = 0$, corresponding to a pressureless gas.

- **Radiation**. Like matter, radiation density will also scale with volume. However, unlike matter, the wavelengths of light will stretch (redshift) as the scale factor grows. Since the wavelength of light is inversely proportional to photon energy, redshifting introduces an additional factor in the radiation scaling: $\rho_R(t) \propto a^{-4}$. This agrees with the equation-of-state parameter derived from classical statistical mechanics: $w = -1/3$ (Reif 1965).

- **Dark Energy**. If dark energy is a cosmological constant, then it does not evolve with time and $w = -1$; that is, a cosmological constant dark energy has a negative pressure, causing space to expand. However, in order to parameterize our ignorance, we usually leave the density-scaling of dark energy in terms of its equation-of-state parameter: $\rho_{DE} \propto a^{-3(1+w)}$. The equation-of-state parameter, in turn, has been parameterized in several ways. The most common parameterization, proposed by Linder 2003, expands w as a function of $a(t)$ to first order:

$$w(a) = w_0 + w_a(1 - a) \tag{2.33}$$

  In this parameterization, dark energy is described by three parameters: (Ω_{DE}, w_0, w_a).

- **Curvature**. The curvature contribution is only parameterized in terms of a density for convenience, and there is no physical significance to be attached to its evolution. Rather, we can just read off its evolution directly from Equation 2.27: $\rho_K \propto a^{-2}$.

We can then combine these equations with Equation 2.24:

$$H^2 = \frac{8\pi G}{3c^2}\left[\rho_{M,0}a^{-3} + \rho_{R,0}a^{-4} + \rho_{DE,0}a^{-3(1+w)} - \rho_{K,0}a^{-2}\right] \tag{2.34}$$

Pulling out a factor of $\rho_{crit,0}$ from inside the brackets and using Equation 2.29 allows us to express our simplified Friedmann equation:

$$H^2 = H_0^2\left[\Omega_M a^{-3} + \Omega_R a^{-4} + \Omega_{DE}a^{-3(1+w)} - \Omega_K a^{-2}\right] \tag{2.35}$$

In fact, by realizing that the expansion of space-time will cause a fractional shift in the wavelength of light ($\lambda_{\text{obs}} = (1+z)\lambda_{\text{emit}}$), we can relate the scale factor $a(t)$ to the redshift z of a distance object:

$$(1+z) = a^{-1} \tag{2.36}$$

This allows us to rewrite the Friedmann equation in terms of redshift, which is the directly observable quantity:

$$H^2 = H_0^2\left[\Omega_M(1+z)^3 + \Omega_R(1+z)^4 + \Omega_{DE}(1+z)^{3(1+w)} - \Omega_K(1+z)^2\right] \tag{2.37}$$

For simplicity, an auxiliary function is often defined

$$E(z) = \sqrt{\Omega_M(1+z)^3 + \Omega_R(1+z)^4 + \Omega_{DE}(1+z)^{3(1+w)} - \Omega_K(1+z)^2} \tag{2.38}$$

which then lets us express the Friedmann equation succinctly as:

$$H = H_0 E(z) \tag{2.39}$$

This equation is easily one of the most important equations in cosmology. Given a set of cosmological parameters $(H_0, \Omega_M, \Omega_R, \Omega_{DE}, \Omega_K)$, we can evaluate the Friedmann equation at any redshift and begin to probe the expansion history of the universe and calculate many physical quantities (e.g., distance, luminosity). Conversely, we can use the Friedmann equation to constrain cosmological parameters if we have measurements of cosmological quantities.

Experimental evidence, however, can simply the equation even further. All of our modern cosmological measurements suggest that the dark energy equation-of-state parameter is -1, in agreement with a cosmological constant. We also have strong evidence for a flat universe, so $\Omega_K = 0$. And, as mentioned in Chapter 1, computer simulations suggest than dark matter is "cold," that is, non-relativistic at the time of the matter-radiation equality, and probably consists of massive, weakly interactive particles (M. Davis et al. 1992; S. D. M. White et al. 1993). This combined theory is called the ΛCDM model of cosmology, and is the fiducial model used in cosmological theory. Such a model reduces the number of free parameters in the Friedmann equation from five to three. We mention such a simplification for completeness and, where appropriate, will use it, but our theoretical discussion will continue to use the more general parameter set.

When discussing dark energy constraints, a common metric for evaluating the strength of experiment in constraining cosmology is the *figure of merit* (FoM). Defined by the Dark Energy Task Force (Kolb et al. 2006), the figure of merit is equal to the reciprocal of the area of the 95% confidence interval ellipse in the (w_0, w_a) plane. For a survey like the Dark Energy Survey (T. Abbott et al. 2005), which will probe the dark energy equation-of-state in several ways, the forecasted confidence intervals are show in Figure 2.1. We will discuss this in more detail presently.

In all of this discussion, perhaps most insightful is the fact that the Friedmann equation hinges on one cardinal observable: redshift. Redshift is truly one of the most

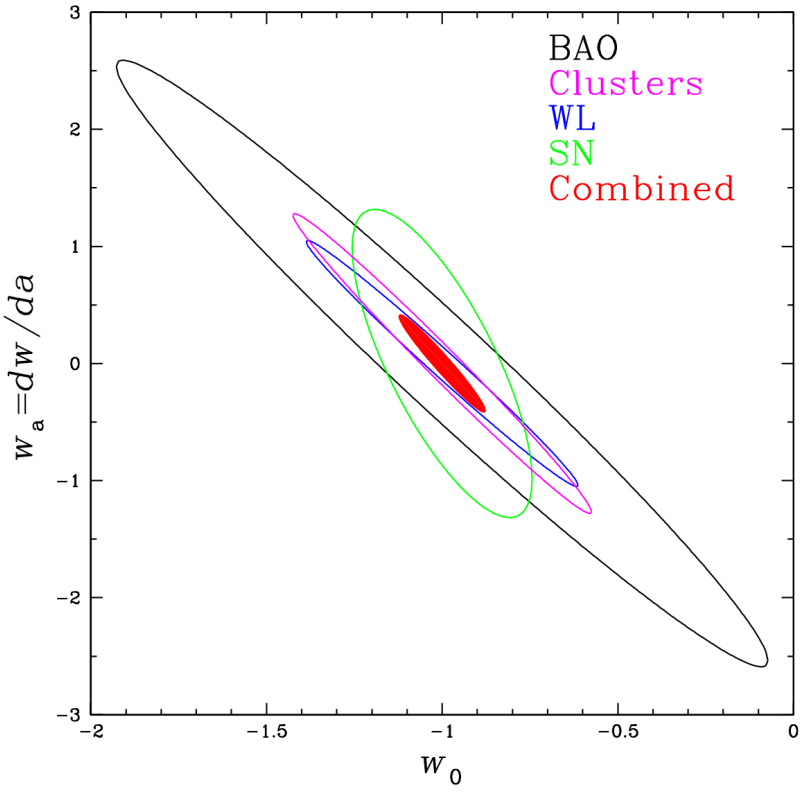

Figure 2.1: 68% confidence intervals forecasted for the Dark Energy Survey (DES) in the (w_0, w_a) plane, marginalizing over other parameters. The DES will probe cosmology on multiple fronts: BAO (black), clusters (magenta), weak lensing (blue), and SNe (green). Combined (red filled), these provide strong constraints on the dark energy equation-of-state. Figure taken from The Dark Energy Survey 2007.

fundamental observables in cosmology. Nearly every modern experiment requires a redshift or redshift estimate to fully constrain cosmological parameters. Generally, this is because the distances between astrophysical objects, or the ability to sort measurements into distance bins (*tomography*), is crucial to a method's ability to extract cosmological signals. For example, type Ia supernovae, being standard candles, require accurate line-of-sight distances to the observer in order to infer the evolution of the intervening space-time (S. Perlmutter et al. 1999; Riess et al. 1998). As another example, weak gravitational lensing considers the redistribution of a source galaxy's surface brightness caused by the presence of a massive gravitational lens, and the strength of the lensing signal depends on the distances between the lens and the source and between the observer and the source (Hoekstra and Jain 2008).

Therefore, we briefly describe the theoretical methodology for measuring distances, and then, since so many techniques rely on accurate determinations of redshift in order to constrain cosmology, we turn our attention to measuring and estimating redshift.

2.2 Distance

There is a subtlety in computing distances in cosmology: since the universe may not be flat, our intuitive notions of distance measure do not always apply. In fact, for an FRW metric there are several different notions of distance. The correct distance measure to use depends on the quantity being measured. The discussion follows the exposition in D. W. Hogg 2000 and D. Huterer 2010.

2.2.1 Proper Distance

The *proper distance* is the physical distance between an observer and a distance object along the line-of-sight, say at radius r. Consider a light ray emitted from the source.

We need to integrate the metric (Equation 2.3) along the path of the ray, from source to observer. Since light moves along a null geodesic, $ds^2 = 0$, and since we are integrating along the line-of-sight, $d\Omega^2 = 0$. Thus, the metric can be rearranged as:

$$c\,dt = a\,dr \tag{2.40}$$

We need to integrate this equation along the line-of-sight, from some time $t_{\rm emit}$ until now, when $t_{\rm obs} = t_0$.

$$\int_0^r dr = c \int_{t_{\rm emit}}^{t_0} \frac{dt}{a(t)} \tag{2.41}$$

We can change variables in the right-hand-side integral from t to a using:

$$\int_{t_{\rm emit}}^{t_0} \frac{dt}{a(t)} = \int_{t_{\rm emit}}^{t_0} \frac{dt}{a(t)}\frac{da}{da} = \int_{a_{\rm emit}}^{1} \frac{da}{\dot{a}a} = \int_{a_{\rm emit}}^{1} \frac{da}{H(a)a^2} \tag{2.42}$$

And now we can switch from a to z via:

$$a^{-1} = 1 + z \qquad \longrightarrow \qquad -a^{-2}da = dz \tag{2.43}$$

Therefore:

$$r = c \int_0^z \frac{dz}{H(z)} \tag{2.44}$$

Identifying r as the proper distance $D_{\rm prop}$ and using the Friedmann equation, we have:

$$D_{\rm prop} = \frac{c}{H_0} \int_0^z \frac{dz}{E(z)} \tag{2.45}$$

Since the quantity c/H_0 appears frequently and has units of distance, it is sometimes called the *Hubble distance*.

2.2.2 Comoving Distance

The proper distance between two objects which are locked in the Hubble flow (i.e., having no peculiar velocity) changes with time as the universe expands. It is sometimes more convenient to imagine a cosmic ruler expanding at the same rate as the universe. This would mean that two objects locked in the Hubble flow would be measured as having a constant separation. This distance is called *comoving distance* and is defined to be equal to proper distance evaluated today. Thus:

$$D_{\text{prop}}(t) = a(t) \cdot D_{\text{comoving}} \tag{2.46}$$

2.2.3 Angular Diameter Distance

In the previous example, we considered the line-of-sight separation between a source and an observer. If, however, we want to know the distance between two objects at the same redshift but separated by an angle $\Delta\theta$ on the sky, we need to change our approach.

Now, in Euclidean geometry, we have our familiar relation between the radius r and arc length s of an arc which subtends an angle $\Delta\theta$: $s = r\Delta\theta$. However, when working in the generalized spherical coordinates of the FRW metric, this no longer holds in its trivial form. The geometric reason is simple: we are using the wrong radius. We now have at our disposal two radii to choose from: a geodesic radial coordinate r and a circumferential radial coordinate $S_\kappa(r)$. When performing calculations along a line-of-sight, we must use the geodesic radius; however, when performing calculations along a circumference (or the surface of a sphere, as the case may be), we must use the circumferential radius in our calculations.

Thus, we define the *transverse comoving distance* $D_\perp$ as the circumferential radius at this redshift in comoving coordinates. So pulling from Equation 2.5 and inserting

the geodesic radius $r = D_{\text{comoving}}$ gives us:

$$D_\perp = \begin{cases} \frac{c}{H_0}\frac{1}{\sqrt{\Omega_K}} \sin\left(D_{\text{comoving}} \cdot \frac{H_0}{c}\sqrt{\Omega_K}\right) & \kappa = +1 \\ D_{\text{comoving}} & \kappa = 0 \\ \frac{c}{H_0}\frac{1}{\sqrt{|\Omega_K|}} \sinh(D_{\text{comoving}} \cdot \frac{H_0}{c}\sqrt{|\Omega_K|}) & \kappa = -1 \end{cases} \tag{2.47}$$

where we have substituted in Ω_K in favor of R.

This is the comoving radius we need to use when calculating arc lengths. So the comoving distance between two objects at the same redshift separated by a angle $\Delta\theta$ is $D_\perp \cdot \Delta\theta$. The physical distance is therefore: $a(t) \cdot D_\perp \cdot \Delta\theta$. Comparing this equation for physical distance to the intuitive Euclidean version $r \cdot \Delta\theta$, we define the *angular diameter distance*:

$$D_A = a(t) \cdot D_\perp \tag{2.48}$$

If you have two objects at the same redshift separated by an angle $\Delta\theta$ on the sky, $D_A \cdot \Delta\theta$ is the physical distance between them.

2.2.4 Luminosity Distance

The *luminosity distance* D_L is the distance one would infer (in a Euclidean geometry) by measuring a bolometric[7] flux F emitted by a source with bolometric luminosity L. The relationship between F and L depends on the area of the sphere over which L is being distributed:

$$F = \frac{L}{A} \tag{2.49}$$

Now, the surface area A of a sphere with centered on the observer is $A = 4\pi D_\perp^2$, where the "radius" we are using is again the transverse comoving radius, since we are measuring perpendicular to the line-of-sight, and not along it. The flux we measure, however, will be reduced by two factors of $(1 + z)$: one factor due to time dilation

[7] A bolometric quantity is one integrated over all frequencies, as opposed to measured through a bandpass filter.

and a second factor because the photon energy will be redshifted to lower energy. Therefore, our equation relating F to L is:

$$F = \frac{L}{4\pi D_\perp^2 (1+z)^2} \tag{2.50}$$

and we can therefore identify the luminosity distance as:

$$D_L = D_\perp (1+z) \tag{2.51}$$

It is precisely this distance—and its dependence on redshift—that was used to discover dark energy using SNe Ia.

Equations 2.47, 2.48, and 2.51 give us the following relationship:

$$D_L = (1+z) D_\perp = (1+z)^2 D_A \tag{2.52}$$

If one could measure any two of these distances at a given redshift, then any discrepancies in this relationship would imply a violation of the homogeneous, isotropic universe described by the FRW metric.

2.3 Redshift

Redshift is a fundamental observable in cosmology. It is, in a sense, our time axis, on which—via distance measurements—we can probe the expansion history of the universe and constrain cosmology. With this in mind, we now turn toward the problem of actually measuring or estimating redshift. Our focus will be on galaxies, as stars are too dim to be seen at the cosmic distances that we are interested in.

2.3.1 Spectroscopic Redshifts

Fundamentally, redshift is observed as a shift in the wavelength of light as in propagates through an expanding space-time. We can therefore determine an object's

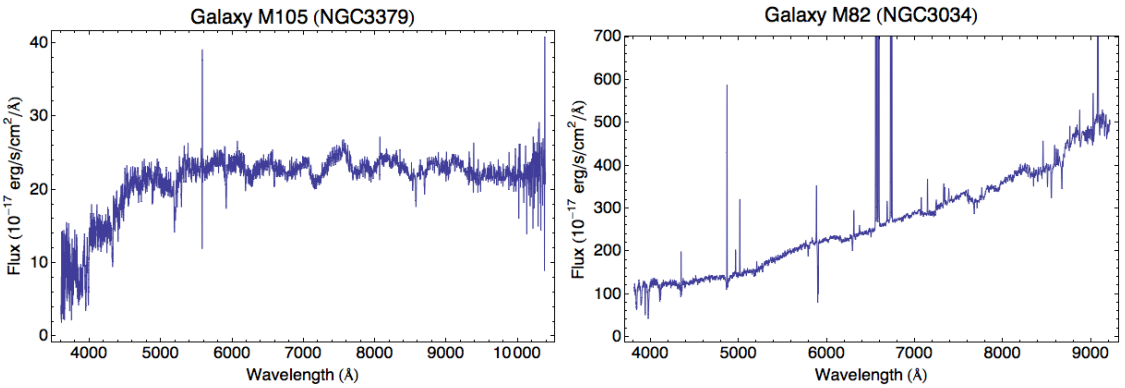

(a) Typical E/S0 Galaxy near $z = 0$ (b) Typical Starburst Galaxy near $z = 0$

Figure 2.2: Typical galaxy spectra, as measured across optical wavelengths. The various absorption and emission lines are apparent as local crests (or spikes) and troughs in the otherwise relatively smooth continuum.

redshift by comparing the observed spectral lines to known spectral features, measuring how far the spectral lines must have shifted. This requires an understanding of the spectral energy distribution (SED) of the distant object, which can be obtained either by understanding the elemental composition of stars, galaxies, etc., or by comparing the SEDs of distant objects to the SEDs of nearby ones.

To drive the point home, we show the spectrum of a typical E/S0 galaxy in Figure 2.2a at $z = 0.003$. Approximately thirty spectral features can easily be detected from the local crests and troughs in the otherwise smooth background. One of the most important spectral features is called the 4000Å break (sometimes called the Ca II break or the HK break). It is the distinct jump in the SED which occurs around 4000Å (in the galaxy's rest frame). Easily identified by eye, its strength depends on the age of the galaxy's stellar population, and tends to be more evident for older, redder galaxies. Of course, galaxies come in many different types. A younger, star-forming galaxy tends to have a stellar population with strong emission lines, as seen in Figure 2.2b. Features such as the 4000Å break become less obvious, and in some cases (e.g., QSOs) are not exhibited at all.

Regardless of its type, the most straightforward experimental method for charac-

terizing a galaxy's SED is to use a spectrograph. Light from the distant object passes through the spectrograph and is spread into, perhaps, hundreds or thousands of independent resolution elements. With sufficiently long exposures, good signal-to-noise spectra can be obtained and used to measure the redshift. Typical spectroscopic surveys, such as the Sloan Digital Sky Survey (SDSS; York et al. 2000), obtain redshift measurements with fractional errors $\Delta z \approx 0.0002$. Such measurements are sufficiently accurate for virtually any cosmological application.

To date, only approximately 3–4 million redshifts have been spectroscopically measured for galaxies. By far, most of these redshifts have been the product of SDSS and the Baryon Oscillation Spectroscopic Survey (BOSS; Dawson et al. 2013). The limiting factors in obtaining spectroscopic redshifts is time and instrument limitations. Spreading light into its constituent wavelengths is necessary for measuring high-resolution spectra; however, this also reduces the signal-to-noise at any resolution element, demanding long exposures for a sufficiently accurate redshift determination. Moreover, each target must be assigned a fiber (to route light to the detector) and there are a limited number of fibers. Each fiber must be manually attached to a plug-plate, which is unique and must be custom drilled for each field to be measured[8]. As a result, there simply aren't enough spectrographs in existence to obtain redshifts at a faster rate.

2.3.2 Photometric Redshifts

The limitation imposed by the relatively meager size of spectroscopic catalogs—and the time constraints in trying to quickly grow existing catalogs—is an obstacle to many modern astrophysical measurements which require large statistics in order to perform precision cosmology. Many such studies could benefit from using substan-

[8]The newest spectrographs, such as those planned to be employed by MS-DESI (Levi et al. 2013; Tarlè 2013), may have robotically positioned fibers which, although still requiring time to position (and limited by the number of fibers), will speed up the throughput of spectral measurements.

tially larger catalogs, even if this implies a degradation of redshift quality. An alternative to using a spectrograph is to use broadband imaging in several bandpasses. By taking several exposures through different broadband filters, a crude approximation of the SED can be produced. If a map from these integrated SED measurements to redshift can be determined, then redshifts can be produced *en masse*. A redshift estimated in such a way using imaging is called a photometric redshift, or photo-z for short.

This idea was first tried by Baum, who used nine bandpasses to approximate the SEDs of six galaxies in the Virgo Cluster (Baum 1957; Baum 1962). The first attempts to apply this technique to galaxy evolution and redshift estimation were limited by two factors: a lack of adequately calibrated photometry, and a shortage of training statistics for determining the magnitude-redshift mapping (Koo 1981; Koo 1985; Loh and Spillar 1986a; Loh and Spillar 1986b). It wasn't until the advent of massive, well-calibrated, multi-band imaging surveys such as SDSS that photometric redshift estimation became a practical and powerful technique (Cunha et al. 2009; Lima, Cunha, et al. 2008; Oyaizu et al. 2008).

Photometric redshift estimation has several distinct advantages over spectroscopic redshift measurements. First, since incoming light passes through only a handful of broadband filters, exposure times can be drastically reduced while still producing the same signal-to-noise ratios. Second, no manual work is needed; plug-plates and fiber position is unnecessary when nothing more than a photograph is being taken. Finally, instead of being limited to the number of fibers in the spectrograph, photometry can be measured for all objects in the field at once. In short, photometric redshifts estimates can be obtained much more efficiently than spectroscopic redshifts; in fact, next generation, large-scale surveys like the Dark Energy Survey (DES; T. Abbott et al. 2005) and the Large Synoptic Survey Telescope (LSST; Ivezić et al. 2008) depend implicitly on photometric redshifts for obtaining their science goals.

In order for photo-zs to be useful in constraining cosmology, the impact that redshift errors have on the uncertainty in cosmological parameters must be understood. Typically, photo-z accuracy is described by two parameters, averaged in redshift bins across the redshift range of interest. The first is the scatter σ_z. The second is the bias $\Delta z = z_{\rm phot} - z_{\rm true}$, which is the mean offset between the photo-z estimate $z_{\rm phot}$ and the true redshift $z_{\rm true}$. By themselves, the mere presence of these statistical errors in photo-z estimates do not contribute significantly to the figure of merit error budget since they can be calibrated out (Amara and Réfrégier 2007; D. Huterer, Kim, et al. 2004; Lima and W. Hu 2007), as shown in Figure 2.3. Rather, it is the uncertainties in these parameters which must be understood and minimized in order to prevent considerable degradation in the dark energy figure of merit (D. Huterer, Takada, et al. 2006; Ma, W. Hu, and D. Huterer 2006).

Figure 2.4 shows how the uncertainty in the dark energy equation-of-state parameters (w_0, w_a) are affected by uncertainty in photo-z parameters. In the absence of any prior on photo-z uncertainty, the parameters in a (Ω_{DE}, w_0) dark energy parameterization are degraded by a factor of 2; for the (Ω_{DE}, w_0, w_a) parameterization, this becomes a factor of 10. More practically, if we instead demand that our uncertainty in the photo-z priors $(\Delta z, \sigma_z)$ do not degrade our measurements of dark energy by more than 50%—assuming a (Ω_{DE}, w_0, w_a) parameterization—then we must constrain our photo-z priors to within 0.003 (Ma, W. Hu, and D. Huterer 2006). The Dark Energy Survey's science requirements are slightly more stringent: $\sigma(\sigma_z) < 0.003$ and $\sigma(\Delta z) < 0.001(1+z)$ in redshift bins of width 0.1 (J. Annis et al. 2010). As long as photo-z estimation techniques can achieve these limits, modern imaging surveys will be able to competitively constrain cosmological parameters.

The degradation of the figure of merit also depends on the number $N_{\rm spec}$ of spectroscopic observations available, since any photo-z algorithm implicitly must use these measurements to determine a mapping from photometric observables to red-

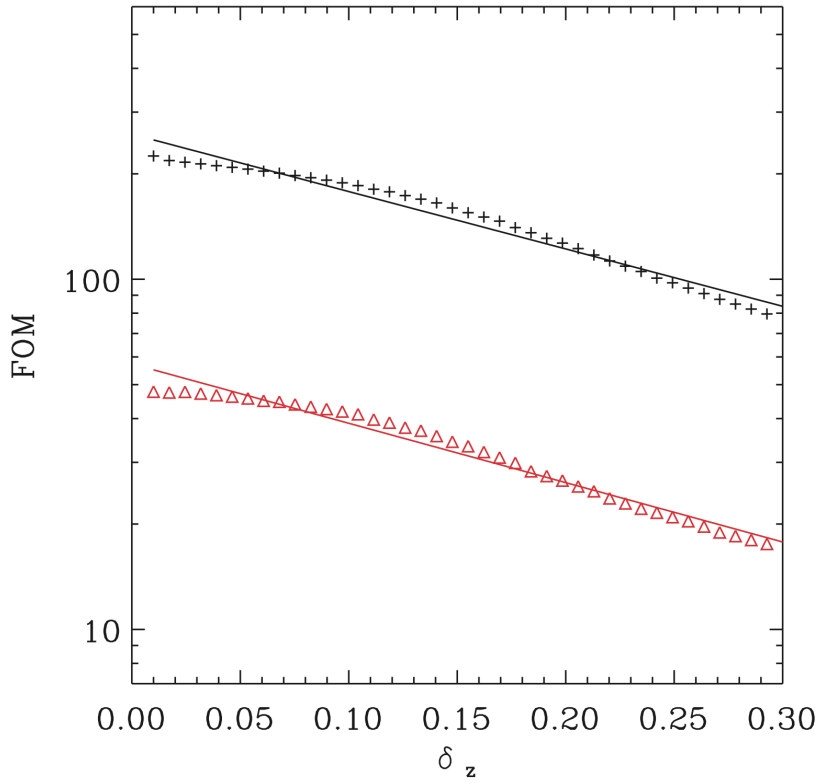

Figure 2.3: Impact of average photo-z error (written here as δ_z rather than σ_z) on the dark energy figure of merit. The black points are for a hypothetical space-based survey with median redshift $z \approx 1.43$, and the red points are for a hypothetical shallower survey with $z \approx 0.9$. Since actual photo-z errors are typically $\sigma_z < 0.1$, we see that we can expect $< 10\%$ degradation in the figure of merit from photo-z errors alone. Figure taken from Amara and Réfrégier 2007.

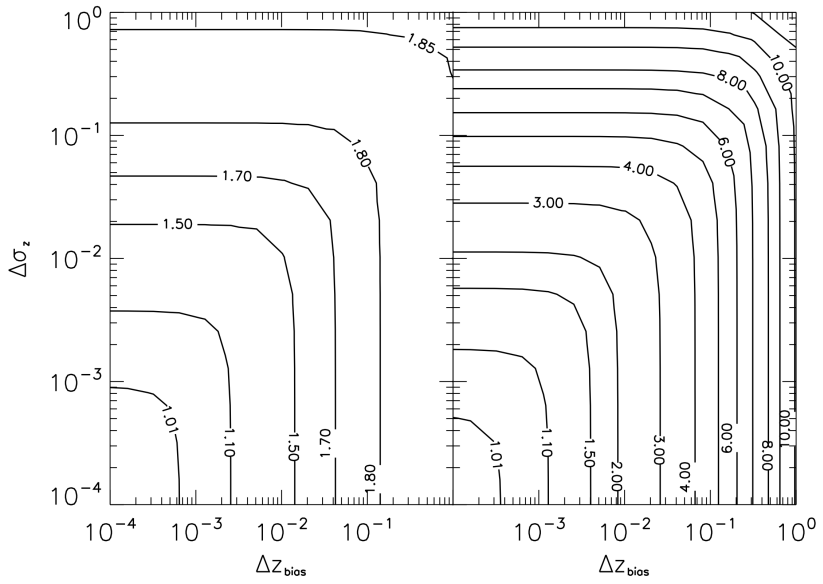

Figure 2.4: Increase in the uncertainty in the dark energy equation-of-state parameters w_0 (left) and w_a (right) as a function of uncertainty in photo-z parameters $(\Delta z, \sigma_z)$. Figure taken from Ma, W. Hu, and D. Huterer 2006.

shift. With more spectroscopic observations (e.g., SED measurements or spectroscopic redshifts), one expects a decrease in the uncertainty of the photo-z priors, and a corresponding increase in the figure of merit. It has been shown (Amara and Réfrégier 2007; Ma, W. Hu, and D. Huterer 2006) that approximately 10^4–10^5 spectra are necessary in order to prevent the figure of merit from degrading by more than 50% (for other figure of merit degradations, refer to Figure 2.5). As we will see in Chapter 4, modern imaging surveys like the SDSS and the DES have these spectra available.

Many different methods have been explored in order to map from observable space to redshift. Broadly, however, they fall into two categories of algorithms: empirical methods and template-based methods. Template-based methods attempt to estimate redshift by comparing photometric observables to a set of known or theoretical SED templates. Although such an approach can work well, the imager must be well-characterized at all wavelengths to be observed, and a set of template SEDs must be available which span the full range of wavelength and spectral types for the galaxies

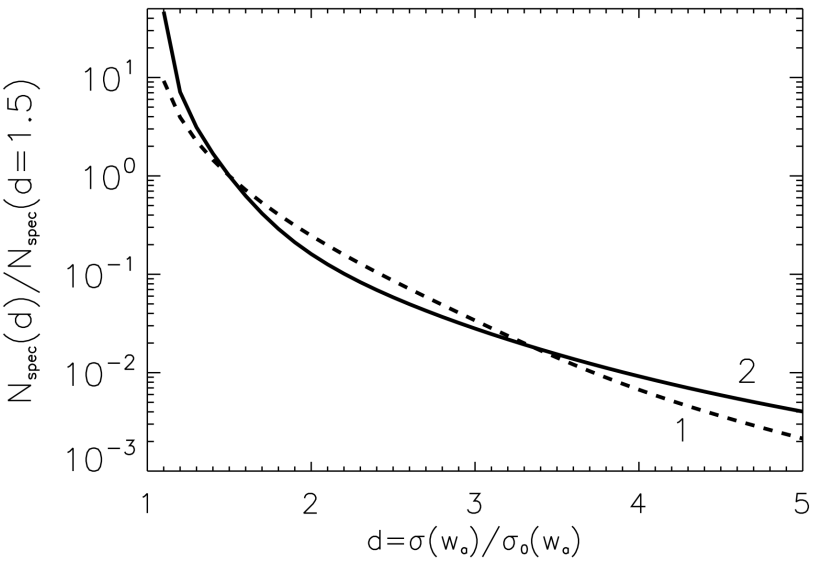

Figure 2.5: The number of spectroscopic observations required to prevent photo-z uncertainty from degrading the dark energy figure of merit by more than a given fraction. For example, if a 100% degradation ($d = 1$) is tolerated, then the number of training spectra can be reduced by a factor of 10. The solid and dashed lines represent two different photo-z prior templates. Figure taken from Ma, W. Hu, and D. Huterer 2006.

being observed. Moreover, the mis-match between the template SEDs and the photometric observables dominates uncertainty measurements, and so errors from such techniques are often underestimated. Nonetheless, template-based methods are likely the only viable option for computing photo-zs in the redshift regime where template SEDs are scare or unavailable and instead need to be synthesized from theories of galaxy evolution. This has made template-based methods ideal for studies in the Hubble Deep Field and Ultra-deep Field (Coe et al. 2006). Examples of successful template-based approaches are: BPZ (Benítez 2000), LePhare (Arnouts et al. 1999; Ilbert et al. 2006), ZEBRA (Feldmann et al. 2006), HyperZ (Bolzonella, Miralles, and Pelló 2000), EAZY (Brammer, van Dokkum, and Coppi 2008), the Bayesian template fitting method of Bender et al. 2001, and a method using linear combinations of SEDs with a Monte Carlo uncertainty estimate developed in Rudnick, Franx, et al. 2001 and Rudnick, Rix, et al. 2003.

However, many of the next-generation, large-sky surveys are challenged not by a

lack of spectroscopic redshifts across the range of interest, but by the sheer number of objects for which photo-zs are required. In many cases, the second category of photo-z estimators—empirical methods—prove to be superb algorithms. These methods use a training set of galaxy observables for galaxies with known redshifts, and then empirically determine the transformation from this multidimensional space to redshift. Such an undertaking is a classic problem in machine learning (Mitchell 1997). The first attempts at empirical photo-z estimation used polynomial fitting to interpolate a smooth transformation between magnitude and redshift (Connolly, Csabai, et al. 1995; Connolly, Szalay, et al. 1997). Although easy to conceptualize, such simplified models can only perform moderately well in an artificially small parameter space of observables. In order to harness greater flexibility and robust estimates, more complex mappings are required. Several techniques have subsequently arisen, such as artificial neural networks (Collister and O. Lahav 2004; Firth, O. Lahav, and Somerville 2003), k-nearest neighbor estimators (Ball, R. J. Brunner, Myers, Strand, Alberts, and Tcheng 2008; Ball, R. J. Brunner, Myers, Strand, Alberts, Tcheng, and Llorà 2007; Cunha et al. 2009; Lima, Cunha, et al. 2008; Zhang et al. 2013), random forests (Carliles et al. 2008), prediction trees (Kind and R. J. Brunner 2013), diffusion maps (Freeman et al. 2009), Gaussian process regression (Way et al. 2009), support vector machines (Wadadekar 2005), ensemble modeling (Way et al. 2009), and k-d trees (Csabai et al. 2003). These methods are also attractive on account of their flexibility in supporting training parameters beyond strict photometric observables (e.g., galaxy shapes). In contrast with template-based approaches, however, empirical methods are not well-suited for extrapolating beyond the parameter space they were trained in.

Of the many possible approaches to photo-z estimation, the author is particularly interested in ArborZ (D. W. Gerdes et al. 2010), which uses boosted decision trees (BDTs) to estimate redshift. This technique not only produces redshift estimates

whose quality is comparable to—and often better than—other algorithms, but also has the distinct advantage of producing a full $p(z)$ probability distribution for the redshift of each galaxy. Although other algorithms exist which can produce $p(z)$ distributions (see, for example, Cunha et al. 2009; Kind and R. J. Brunner 2013; Lima, Cunha, et al. 2008), ArborZ's $p(z)$ results often outperform other methods and have been studied in depth by the author, who was on the team which developed ArborZ.

Therefore, we turn our attention to understanding the intricacies of photo-z estimation using ArborZ, and apply it to several real data sets. With such a tool in hand, we have a strong position from which to constrain cosmology and better understand our universe.

Chapter III

ArborZ

ArborZ is a novel photometric redshift estimation algorithm which uses boosted decision trees (Freund 1995; Schapire 1990) to empirically determine a mapping from galaxy observables to redshift. The technique produces a $p(z)$ probability distribution for each galaxy it is evaluated against, bolstering the strength of the algorithm. Such probability distributions are extremely valuable in many branches of cosmology, and in some cases—such as determining the galaxy source distribution in weak lensing—having an accurate $p(z)$ measurement is as important as accurately determining the redshift itself (Mandelbaum et al. 2008).

A boosted decision tree (BDT) is an algorithm for classifying a vector of observables into one of several disjoint subsets; in the simplest case of two subsets, we could label them "signal" and "background." When there is strong separation (in observable space) between these classes, simple cuts are often sufficient for obtaining a categorization which yields high completeness and high purity. In most realistic situations, however, the relationship between observables and classes is more complex and requires a more technical treatment than simple cuts. Machine learning is ideal for solving such problems, and BDTs are one of most successful techniques to emerge in recent years for solving the classification problem (Hastie, Tibshirani,

and J. Friedman 2009). BDTs have been effectively employed in fields as diverse as handwriting recognition (Howe, Rath, and Manmatha 2005), spam filtering (Drucker, Wu, and Vapnik 1999), and particle identification in high-energy physics (Roe et al. 2005).

Boosted decision trees work by constructing a series of binary cuts on the set of training vectors which maximizes the separation between signal and background classifications. These cuts form a binary tree called a *decision tree* (DT). Mis-classified vectors are reweighed, and the process of constructing DTs repeats iteratively until some error threshold is reached. The overall output of a BDT is then a linear combination of the outputs of the individual DTs, weighted by their misclassification rate.

ArborZ uses BDTs to solve the photo-z estimation problem by dividing the training set into bins in redshift space, and then training one BDT per bin. Each BDT is trained to recognize as "signal" all galaxies whose observables would put it into the respective bin. Galaxies whose true redshifts fall sufficiently far away (so as to prevent overtraining, as we discuss presently) from the BDT's redshift bin are then classified as "background." When presented with a new observable vector, the output of this ensemble can be collated into a $p(z)$. Errors estimates can then be derived from the width of $p(z)$.

In this chapter, we discuss the ArborZ algorithm in detail, starting with the theory behind decision tree construction and working through boosting until we can explain ArborZ in full. We discuss the ideal parameters for constructing BDTs and how ArborZ is applied to galaxy samples in practice. We will use simulated galaxy catalogs to test its performance and show that $p(z)$ distributions are indeed a powerful statistical tool and deserve a standard place in any cosmologist's toolbox.

3.1 Decision Trees

At the core of any boosted decision tree—and therefore at the heart of ArborZ, as well—are the individual decision trees which provide the backbone for classification. We begin with some basic definitions to quantify our discussion. A training set X is a set of n vectors $\vec{x}_1, \vec{x}_2, \ldots, \vec{x}_n$ (often called "events") along with the respective classifications $y_1, y_2, \ldots, y_n$, where $y_i = +1$ for signal events and $y_i = -1$ for background events. Each training vector also is assigned a weight w_i, initially set to unity. Finally, we define the step function:

$$[\![x]\!] = \begin{cases} 1 & x > 0 \\ 0 & \text{otherwise} \end{cases} \tag{3.1}$$

3.1.1 Growing the Tree

The process of constructing a decision tree begins with a root node which contains the entire training set X. The total weight in the node is:

$$W = \sum_{i=1}^{n} w_i \tag{3.2}$$

This allows us to conveniently express the total weight in the node contributed by signal events or background events as:

$$W_S = \sum_{i=1}^{n} w_i [\![y_i]\!] \tag{3.3}$$

$$W_B = \sum_{i=1}^{n} w_i [\![-y_i]\!] \tag{3.4}$$

The algorithm then searches the training vectors to find a cut on the one variable which would provide the best separation between classes. To do this, a measure of separation needs to be introduced. Then the optimal split is that which produces

the largest increase in separation between classes in the parent node and its children (which result from the split)[1].

In order to choose a measure of separation, we first define w_S, the fraction of weight contributed by signal objects:

$$w_S = \frac{W_S}{W_S + W_B} \tag{3.5}$$

where w_B is defined similarly for background galaxies: $w_B = 1 - w_S$. Common separation measures are:

- **Gini index**.

$$G = w_S w_B = \frac{W_S W_B}{(W_S + W_B)^2} \tag{3.6}$$

- **Cross-entropy**

$$E = -w_S \ln w_S - w_B \ln w_B \tag{3.7}$$

- **Misclassification Rate**

$$M = 1 - \max(w_S, w_B) \tag{3.8}$$

- **Statistical Significance**

$$S = \frac{W_S}{\sqrt{W_S + W_B}} \tag{3.9}$$

The optimal cut is the one which maximizes $W_{\text{parent}} S_{\text{parent}} - (W_{\text{left}} S_{\text{left}} + W_{\text{right}} S_{\text{right}})$, where S represents the chosen separation measure applied to the training vectors in the "parent" tree and each of the two resulting sub-trees labeled "left" and "right";

[1]Several authors refer to the optimal split as "maximizing the decrease in impurity," where impurity measures the degree of non-separation (mixing); however, this can lead to confusing a notation when talking about impurity and purity, which would no longer be antonyms.

separation measure is also multiplied by the total weight contained in the subtree beginning at the respective node.

We generally see best results when using the Gini index, with similar or worse results in other measures. We discourage the use of the "statistical significance" since it is monotonically increasing with w_S. This is undesirable, since a node can have good separation for large w_S (containing many signal events) or small w_S (containing many background events). We include it in this list because it has been suggested as an option in the literature (Hoecker et al. 2007). The other separation metrics listed here turn over at $w_S = 0.5$ and approach zero at $w_S \to 0$ and $w_S \to 1$, which fairly consider both signal and background weight.

The algorithm is then recursively applied to each child node. There are several termination conditions for the recursion. First, if the current node contains all signal or all background events ($w_S = 0$ or $w_S = 1$), then the decision tree has done its job of separating classes, and so the recursion terminates. Second, if the number of entries is smaller than some threshold number (which we call a "bucket size"), then we will terminate recursion regardless of the class mixing in the node. The reason is two-fold: we want to be able to characterize these terminal or "leaf" nodes statistically (which cannot be done if we only have one event present), and we want to prevent overtraining (that is, we want to prevent growing a tree around the random noise in our training data). A typical bucket size is on the order of ten events.

Other termination conditions can also be applied to the the recursive algorithm. The most common one is a depth cut. If the path length from the current node to the root node exceeds a certain distance, then recursion stops. Similarly, some authors (Roe et al. 2005) suggest that, of the two child nodes available at each recursive step, only the most promising one (in terms of its separation measure) should be split. In the case that many training events are presented to the DT, these stopping criteria can drastically decrease the computational time and resources required to grow a

tree. Additionally, these criteria can help to prevent overtraining. However, there are two problems with these "premature termination" approaches. First—in the case of the depth limit—there isn't a natural relationship between the number of training events and a maximum tree depth. Second, and more importantly, premature termination may cause the tree to perform sub-optimally, since apparently insignificant cuts can eventually result in significant separation after several cuts. By stopping training early, one prohibits the DT from possibly finding these difficult, but important, opportunities for classification. For example, consider one or more small, pure (all signal or all background) subspaces in an otherwise noisy observable space. It may take many cuts before the subspaces are isolated and correctly classified. If one stops training early, or only grows the most promising branch of the tree, then the statistical signal in these subspaces will be lost. Therefore, we recommend never[2] prematurely stopping the training algorithm.

3.1.2 Using the Tree

At this point, one has a fully trained decision tree (a pedagogical example is shown in Figure 3.1). Each node in the tree is then assigned a classifier h_i, which is defined as $+1$ (a signal node) if $w_S > 0.5$ for the weights in that node[3], and -1 otherwise (a background node). To use a tree to classify a new vector $\vec{x}$, one starts at the root node, traverses the tree (by following the cuts at each branch) until a leaf node is reached, and then assigns the leaf's classifier h_i to $\vec{x}$; that is, we approximate y for the vector $\vec{x}$ as the classifier h_i of the leaf node which would contain $\vec{x}$. An example is shown in Figure 3.2

[2]Of course, "never" is a strong word. In most astrophysical applications, this statement is true. In other data mining fields, however, the sheer size of the data sets can be computationally overwhelming. In such cases, it would be better to explore other ways to quicken training, such as weight trimming (J. Friedman, Hastie, and Tibshirani 2000) or stochastic boosting (J. H. Friedman 2002).

[3]The threshold value 0.5 is chosen to be unbiased with respect to both signal or background events. This is not always desired, as it may be more important to correctly classify signal events

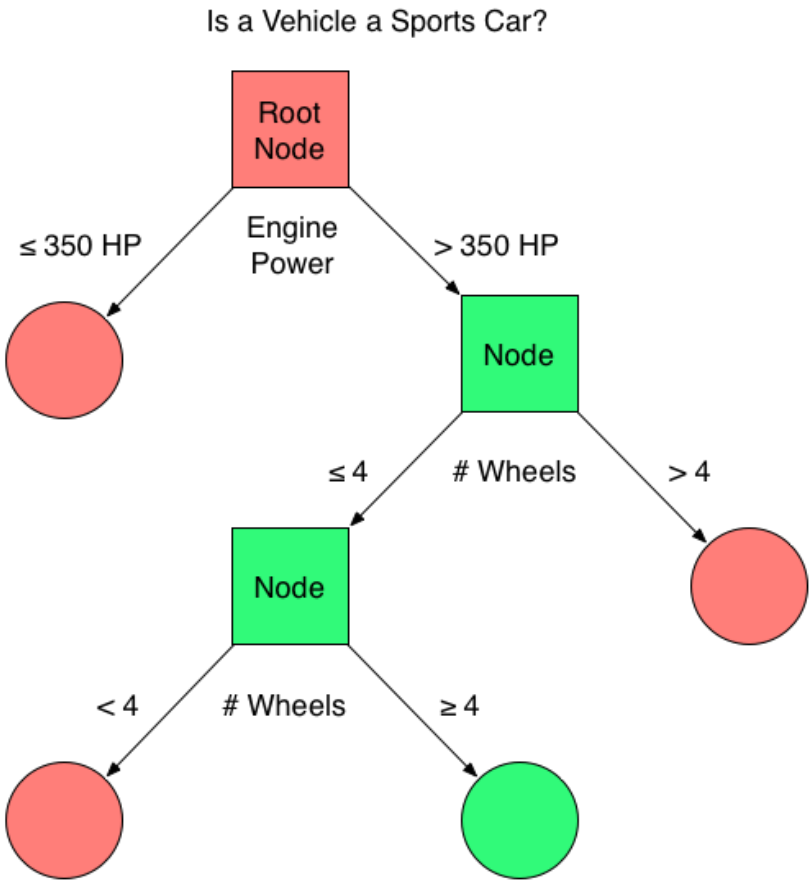

Figure 3.1: An example of a very simple decision tree used to classify cars as sports cars. In this example, engine power is the largest discriminator when determining if a car is a sports car, and so it becomes the first split. The number of wheels on the car was the second (and third) most important variable. Squares indicate branch nodes and circles are used to designate leaf nodes. Green nodes indicate nodes which would be classified as "signal" nodes (nodes in which most of the training vectors are, indeed, sports cars); red nodes are "background" nodes. Since most vehicles are not sports cars, it isn't surprising that the root node is therefore red. On the other hand, most vehicles do have four wheels, and so the other two branch nodes are green. If this example tree continued with splits on color or number of tail-lights, one could conclude that the tree was overtraining.

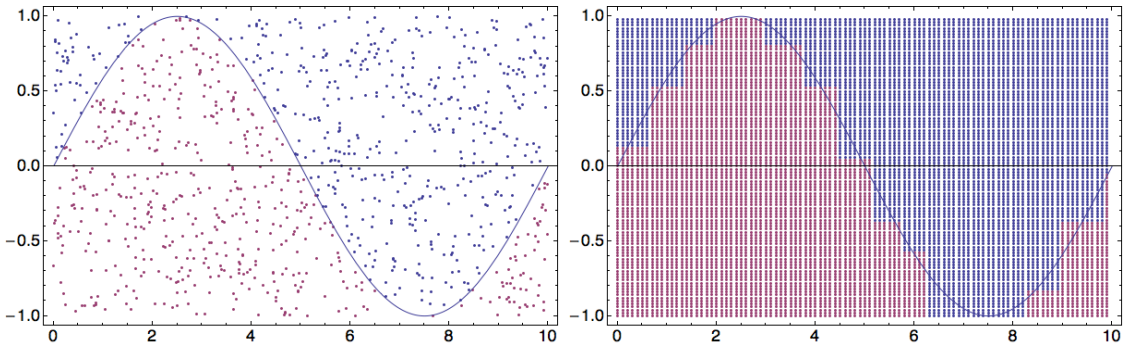

(a) Training Set (b) Evaluation

Figure 3.2: An example of training a decision tree on a simple sine function. On the left we show a training set of 1000 points, defining "signal" to be above the curve (blue points) and "background" to be below the curve (red points). After training the decision tree, we can step through the entire space, querying the tree for its best estimate of the true classification at each point. The results are shown on the right. We see that the decision tree performed excellently. In a handful of places there wasn't enough data in the training set for the decision tree to emulate the function appropriately (e.g., at the trough of the sine wave). We also note that since a decision makes cuts on only one variable at a time, all cuts are orthogonal to the axes, which is clearly visible.

With a fully grown tree and these definitions in hand, we can quantify a tree's performance. The purity p of a decision node is defined as the fraction of the weight in a node N which is correctly classified with respect to the node's classifier:

$$p_N = \frac{1}{W} \sum_{i=1}^{n} w_i [\![h_i y_i]\!] \tag{3.10}$$

where we note that $[\![h_i y_i]\!] = 1$ if event i is correctly classified and zero otherwise. The resubstitution error R of a node N is then:

$$R(N) = 1 - p = \frac{1}{W} \sum_{i=1}^{n} w_i [\![1 - h_i y_i]\!] \tag{3.11}$$

A similar definition for $R(T_N)$ can be defined for the entire subtree T_N beginning at node N, where we use the weights w_i and classifiers h_i in each leaf node under T_N rather than the weights and classifiers of the events associated with node N. Gen-

than background events, etc.

erally, $R(N)$ and $R(T_N)$ will be smaller for better trained (that is, better separated) nodes and trees, respectively. However, these quantities are computed with respect to the training set, and so they are underestimates of typical errors seen when the tree is applied to real data sets.

What R does enable is *pruning*, one of the most commonly performed operations on a grown tree. Pruning attempts to remove the least significant subtrees from the DT. Although it is faster to navigate a pruned tree (since it is smaller than its unpruned version), the primary reason to apply pruning is to prevent overtraining (Breiman et al. 1984). In order to identify the least important subtree, we define the cost complexity ρ for a sub-tree T_N beginning at node N as:

$$\rho(T_N) = \frac{R(N) - R(T_N)}{|T_N| - 1} \tag{3.12}$$

where $|T_N|$ is the number of leaves in subtree T_N. $\rho(T_N)$ is, succinctly, the average decrease in error caused by continuing to grow out the tree T_N rather than simply terminating growth at node N. By iteratively removing the subtree with the smallest cost complexity until no subtrees exist with $\rho < \rho_{\text{prune}}$, we can optimally prune the DT. The choice of pruning strength ρ_{prune} is usually determined empirically.

Despite the apparent advantages of pruning, it is not always appropriate. It is not easy to determine when pruning is discarding useful information or preventing overtraining. If one is using DTs as standalone classifiers (without boosting), pruning may have advantages. More often, however, DTs are used in conjunction with a boosting algorithm. In these cases, boosting will often outperform pruning, and pruning will usually degrade the performance of a BDT. We will discuss the advantages or power of boosting in the next section, but in summary, pruning should be applied with caution.

3.2 Boosted Decision Trees

3.2.1 Boosting

Figure 3.2 is deceiving: in most realistic, complex data sets, particularly higher-dimensional ones, decision trees do not usually perform as well. In fact, DTs fall into a category of learning algorithms called *weak classifiers*, a label given to learning algorithms which are only required to be correct slightly better than 50% of the time (Kearns and Valiant 1994). Thus, individually, DTs are not a viable method for robust classification. However, Schapire 1990 developed a method for combining an ensemble of weak classifiers into a single *strong classifier*, one which can be trained to arbitrary precision. This mechanism is called *boosting*, and is at the heart of boosted decision trees (BDTs).

To understand the principle behind boosting we consider the classic example given in Freund and Schapire 1999. Imagine a gambler who, in an effort to improve his success in horse racing, decides to ask all of his friends which horse they believe will win. Each of his friends will have developed "rules of thumb" to help him make these decisions; for example, one might bet on the horse with the best odds while another always bets on the horse with the best track record. Although crude and often wrong, these rules of thumb still work sufficiently well that they are not outright discarded. By combining his friends' rules of thumb, the gambler will be able to increase his odds drastically. In this analogy, each friend and his rules of thumb are weak classifiers—better than nothing but not that great (in our specific case, the DTs are our weak classifiers). The process of intelligently combining these weak classifiers is called "boosting," and the gambler's successful algorithm which he now uses to place his bets is a strong classifier (in our case, the BDTs).

The original boosting algorithm was AdaBoost (Freund and Schapire 1997), and

it still remains a dominant meta-algorithm for improving the performance of weak classifiers. Other algorithms have also emerged, most of which are variations on the AdaBoost algorithm. Examples are LogitBoost (J. Friedman, Hastie, and Tibshirani 1998), gradient boosting (J. H. Friedman 1999; J. H. Friedman 2002), BrownBoost (Freund 2001), and RobustBoost (Freund 2009). All of these algorithms have one thing in common: they iteratively construct DTs by adjusting ("boosting") the weights of mis-classified objects. When these previously misclassified events are presented to the next DT, their new weights raise the likelihood that the DT splits will correctly classify them. Once the ensemble of DTs (called a "forest") is constructed, each tree is assigned a weight (to be defined presently) based on the purity of its classifications (i.e., how well it solves the classification problem). The linear combination of trees constitutes the BDT. This BDT is the strong classifier used for rigorous classification, considered the "best off-the-shelf classifier in the world" (Hastie, Tibshirani, and J. Friedman 2009).

3.2.2 Constructing Boosted Decision Trees

ArborZ applies AdaBoost to construct its BDTs. For a thorough derivation of AdaBoost, with relevant proofs, see Freund and Schapire 1997. We recount the key steps here to illustrate the procedure. Our equations are optimized for numerical computing, and so differ slightly from the standard notation.

1. Grow the t^{th} tree using the current set of weights $\{w_i\}$. For $t = 1$, set $w_i = 1$ for all i.

2. Calculate the tree's misclassification rate ϵ of the entire tree T_0, defined as the resubstitution error of the tree: $\epsilon = R(T_0)$.

3. Calculate the *boost weight*:

$$\alpha = \frac{1}{2} \ln \left[\left(\frac{1-\epsilon}{\epsilon} \right)^{\beta} \right] \tag{3.13}$$

where β is the boost exponent (typically set to unity).

4. Compute updated weights:

$$w_i \qquad \Longrightarrow \qquad w_i' = w_i \frac{b_i}{Z} \tag{3.14}$$

where the boost factor is:

$$b_i = e^{-\alpha \cdot h_i y_i} \tag{3.15}$$

and Z is a normalization factor:

$$Z = \frac{1}{W} \sum_{i=1}^{N} w_i b_i \tag{3.16}$$

5. Return to the first step and grow a new tree using these new weights.

This iterative process continues until a termination condition is reached. The most obvious termination condition is when the error rate is too low; this happens when the particular set of weights allows a tree to grow perfectly (or as near to perfect as the numerical limits of the computer allow, given that some weights may be extremely small). If this condition is met quickly enough, a limit on the size of the forest is usually imposed. This is sensible, since—due to noise in the data and a finite bucket size—it may not be possible to perfectly classify every event. The error rate reaches an asymptote within the first 30–40 trees, and so 50 trees is a reasonable limit. We will illustrate this phenomenon presently, once the discussion of the ArborZ algorithm has concluded.

3.2.3 Using Boosted Decision Trees

Once the forest is assembled, the BDT classifier must be constructed. For an observable vector $\vec{x}$, the BDT classifier $H(\vec{x})$ is defined as:

$$H(\vec{x}) = \sum_{t=1}^{T} \alpha_t h_t(\vec{x}) \tag{3.17}$$

where the sum runs over the entire forest, and α_t and $h_t(\vec{x})$ are the boost weight and classifier, respectively, for each DT in the forest.

The output of a BDT classifier, sometimes called a "score," differs from that of a DT in that a BDT's score is continuous, whereas a DT is discrete: $+1$ (signal) or -1 (background). A BDT score has only relative—not absolute—meaning: a larger, more positive score is indicative of a signal-like event, and a smaller, more negative score suggests a background-like event. At first glance, this may seem frustrating, since we now have a strong classifier but lack a scale by which to interpret its output, but this is not so. Instead, we need to recast the scores in terms of probabilities.

Bayes' Theorem gives us a natural vantage point from which to interpret these BDT scores. Let D (for "data") be a vector of observables and M_i some particular model or hypothesis. Then Bayes' Theorem states:

$$P(M_i|D) = \frac{P(D|M_i)P(M_i)}{P(D)} = \frac{P(D|M_i)P(M_i)}{\sum_i P(D|M_i)P(M_i)} \tag{3.18}$$

where on the right-hand side we have applied the law of total probability and the sum runs over all possible hypotheses. For those unfamiliar with Bayesian nomenclature, $P(M_i)$ is the *prior*, $P(M_i|D)$ is the *posterior*, and $P(D|M_i)$ is the likelihood. In the particular case of BDTs, we are interested in the probability that an event is signal given its BDT score. Denoting the hypothesis that the event is signal as S and background as B, this is:

$$P(S|H(\vec{x})) = \frac{P(H(\vec{x})|S)P(S)}{P(H(\vec{x})|S)P(S) + P(H(\vec{x})|B)P(B)} \tag{3.19}$$

Now, the probabilities $P(S)$ and $P(B)$ can be fixed by the training set's statistics (just counting the frequency of signal and background events). The conditional probabilities $P(H(\vec{x})|S)$ and $P(H(\vec{x})|B)$ can be calculated by estimating the distribution of scores $H(\vec{x})$ for signal events as well as the distribution for background events. However, we cannot simply use the training statistics for this; we would be choosing highly biased events, since the BDT has already built its classifier around these objects. Instead, we need to reserve a second data set, called a *validation set*, which is similar to the training set in all ways except that the BDT was not presented with its elements for training.

In practice, this validation set is randomly extracted from the full training set database before growing the BDT. The BDT is then presented with the remaining samples for training. Once training is complete, the validation set is evaluated. Since this sample is unbiased and since we know the y_i "truth" values for each vector, we can fairly evaluate Bayes' Theorem, transforming BDT scores into true probabilities. Validation sets are ideally chosen to be half the size of the data set used for training (Breiman et al. 1984).

3.3 ArborZ

ArborZ is the machine learning formalism for applying BDTs to the problem of photo-z estimation. Given a training set of galaxy observables along with their respective spectroscopic redshift, ArborZ divides the training set into redshift bins. A BDT is assigned to each redshift bin and is trained to classify galaxies as belonging to its bin or not. This ensemble of BDTs constitutes ArborZ. When a new set of observables is then presented to ArborZ, each BDT produces a probability that the galaxy belongs

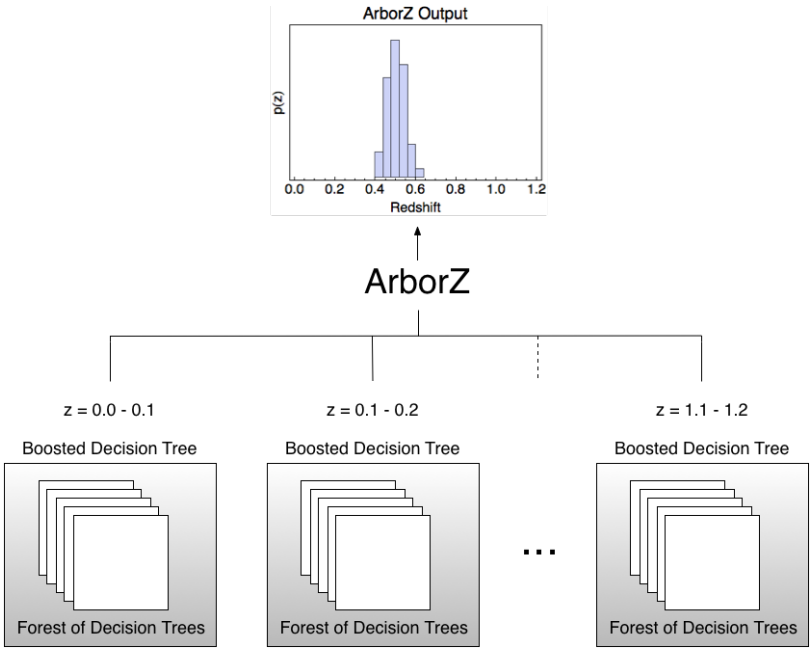

Figure 3.3: ArborZ is an ensemble of BDTs, each of which is assigned to recognize galaxies in its respective redshift bin. Each BDT, in turn, consists of a "forest" of DTs; the BDT unites these weak DT classifiers into a single, strong classifier. After ArborZ is trained, it can evaluated on new sets of observables, producing probability distributions $p(z)$ for each one.

in its bin. The resulting probability distribution is $p(z)$, and from there we can derive other useful metrics (such as an error estimate). A simple schematic describing the relationship between ArborZ, BDTs, and DTs is shown in Figure 3.3.

How many bins should we divide our training set into? Well, we want our bin widths Δz to be on the order of the expected photo-z error. We call this expected error the photo-z resolution, denoted σ_{res}. Any larger and we begin to lose resolution; any smaller, and we risk overtraining the BDTs. If z_{max} is the maximum redshift in the training set, then the optimal number N_{bins} of bins should roughly be:

$$z_{\text{max}} = N_{\text{bins}} \Delta z \approx N_{\text{bins}} \sigma_{\text{res}} \tag{3.20}$$

Good photo-z accuracy is approximately $\sigma_{\text{res}} = 0.02$, and poor performance is approximately $\sigma_{\text{res}} = 0.10$. There is no reason, *a priori*, to assume that a photo-z algorithm ought to perform particularly poorly, so typically we choose values of σ_{res} around

0.01–0.05. For a redshift range of $z_{\rm max} = 1$, this corresponds to a rough estimate of $N_{\rm bins} = 20$–50 bins. As long as sufficient training statistics are available, better redshift number distributions are obtained by erring on the side of more bins, and so $N_{\rm bins} = 50$–80 for $z_{\rm max} = 1$ are common. In the absence of a good prior for $\sigma_{\rm res}$, one can empirically calculate it by minimizing $z_{\rm phot}$ error with respect to $\sigma_{\rm res}$.

An alternative binning method would be one in which each BDT is presented with the same training statistics; that is, the bin widths are permitted to vary in order that each BDT is presented with the same number of events to train on. In principle, this has the advantage of constant signal-to-noise across each BDT, since in a fixed-bin setup, it isn't uncommon to have only a handful (of order ten or so) of training statistics in high-redshift bins. In practice, however, fixed-width bins seem to perform marginally better, and since in many cosmological applications they are easy to work with, we generally suggest using fixed-width bins.

Once the redshift range is divided into bins, we need to train each of our BDTs. This requires that we put discrete labels—signal or background—on each training observable presented to the BDT. We could label galaxies within the bin of interest "signal" and all others "background," but we run the risk of overtraining for galaxies which lie near bin edges. For example, for a bin edges at $z = 0.1$, it is artificial to give galaxies at $z = 0.099$ and $z = 0.101$ different classifications, and doing so would cause BDTs to overtrain on this arbitrary cut-off and ultimately degrade performance. ArborZ's solution is to label galaxies within the bin of interest "signal" and galaxies which fall more than $3\sigma_{\rm res}$ away from the bin's edges are labeled "background"; the remaining galaxies are not presented to the respective BDT during training.

Once we have selected the signal and background samples, we perform a principal component analysis (PCA; Hotelling 1933; Pearson 1901) on the training sample. This helps to reduce correlations between observables, making the DT cuts more efficient (a single cut may suffice where many "steps" would be required to separate

correlated variables). The BDT is then trained on a random two-thirds of this sample; once training is complete, the remaining third is used to determine the probability transformation from BDT scores into probabilities. This process is repeated for each BDT.

Once the training process has completed, we can now present ArborZ with new galaxies. Each BDT outputs a probability that the galaxy belongs in its redshift bin. The ensemble of probabilities forms $p(z)$. Since each BDT outputs a true probability, ranging from 0 to 1, this $p(z)$ is not normalized. One could normalize the $p(z)$, but this has an unintentional side-effect. If no BDT was confident that the galaxy belonged in its redshift bin (i.e., each BDT labeled the galaxy as "background"), then $p(z)$ could be nearly zero across its entire range. By normalizing, we discard this interesting information and artificially boost the galaxy's $p(z)$ contribution. In most applications, therefore, it is prudent to normalize (for convenience), but save enough information to recover the original normalization factor. For each galaxy $p(z)$, we define the peak probability as the height of the maximum bin in $p(z)$:

$$p_{\text{peak}} = \max p(z_i) \tag{3.21}$$

Conveniently, p_{peak} also measures ArborZ's confidence in its redshift measurement. One can then place cuts on p_{peak} to select galaxy $p(z)$s which ArborZ believes are well-estimated.

Regardless of the normalization, there are several quantities that can be derived from a galaxy's $p(z)$. First is the "best-estimate" photo-z, z_{phot}. Typically, this is defined as the mean of the distribution:

$$z_{\text{phot}} = \frac{\int p(z) z \, dz}{\int p(z) \, dz} \tag{3.22}$$

When a single number is desired to characterize the redshift, this is the value we

suggest using. We do, however, caution against excessive use of $z_{\rm phot}$, since any calculations using it are lacking the strong benefits of a full distribution and discarding information in favor of simplicity. To estimate an error σ_z on $z_{\rm phot}$, define σ_z to be half the width of the range containing the central 68% of the area of $p(z)$.

Of course, it is fair to ask whether or not other $z_{\rm phot}$ quantities would better represent the distribution $p(z)$ than its mean. The most obvious candidates are the median and mean. The mode tends to perform poorer than either the mean or median, since it is plagued by the discretized values which result from computing the mode of a binned distribution. These manifest themselves as "spikes" when calculating the redshift distribution $N(z)$ of a sample. The median tends to perform slightly (less than 10%) better than the mean; however, the errors are then overestimated by about 20%. If the error estimates are not needed in a particular application, we suggest using the median; otherwise, the mean will give the best, all-around performance, and so we cast most of our examples in terms of the mean. Unless otherwise specified, $z_{\rm phot}$ refers to the mean of $p(z)$, and $z_{\rm phot}^{\rm med}$ will indicate the median of $p(z)$.

3.4 Characterizing ArborZ

Now that we have thoroughly discussed the underpinnings of the ArborZ algorithm, we will test its performance in simulated (or "mock") galaxy catalogs, where we know the true redshifts for each galaxy. Since our ultimate science goal is produce photo-z estimates for the SDSS and the DES (see Chapter 4), we will use the DES v4.02 mock galaxy catalogs designed to model the color, magnitude, and spatial distributions of galaxies in the DES. The procedure for generating these simulations is described in Appendix A of D. W. Gerdes et al. 2010. The catalog covers $220\,{\rm deg}^2$ out to redshift $z = 1.35$, with magnitude limits of $(26.0, 25.5, 24.8, 24.3, 22.5)$ in $grizY$ bands, respectively (based on the expected 5σ detection limit of DES). It contains

approximately 2.3×10^7 galaxies in ten tiles. We extract 100 000 randomly chosen galaxies from one tile for training, setting aside the remaining 519 116 galaxies in the tile for evaluation.

We train ArborZ using the five g, r, i, z, and Y magnitudes from this training set. We divide the training set into 100 evenly-spaced redshift bins in the range $z \in [0, 1.33]$ using a resolution of $\sigma_{\text{res}} = 0.02$. The training proceeds as discussed in Section 3.3. Each BDT forest consisted of at most 50 DTs. Each DT had a bucket size of 10 and used the Gini index for determining the optimal split.

Example $p(z)$ distributions are shown in Figure 3.4. As we stated earlier, we can use these $p(z)$ to calculate various descriptive metrics, such as the "best-estimate" photo-z z_{phot}, defined as the mean of the distribution, and a photo-z error estimate σ_z, defined as the width of the distribution. To test the bias of $p(z)$, we can measure the fraction f_{pdf} of the area in $p(z)$ which falls below the true redshift:

$$f_{\text{pdf}} = \frac{\int_0^{z_{\text{spec}}} p(z)\, dz}{\int_0^{\infty} p(z)\, dz} \tag{3.23}$$

If $p(z)$ is an unbiased estimator of redshift, then the distribution of f_{pdf} values should be flat. We plot this distribution in Figure 3.5. We see that it is mostly flat everywhere, with a small excess in the center and and small absence at the edges.

Previous work has shown that ArborZ z_{phot} estimates perform similarly or better than competing algorithms (D. W. Gerdes et al. 2010), but the real power of ArborZ is in its probability estimates $p(z)$. To illustrate the strength of $p(z)$, we first plot z_{phot} versus z_{spec}, shown in Figure 3.6a. The first thing to notice is the flattening which occurs near $z_{\text{true}} = z_{\text{phot}} = 0.4$. The next obvious artifact is the compressed range of z_{phot}: the distribution does not properly fill out the entire redshift range. Both of these issues plague all photo-z estimation algorithms. The first artifact is caused

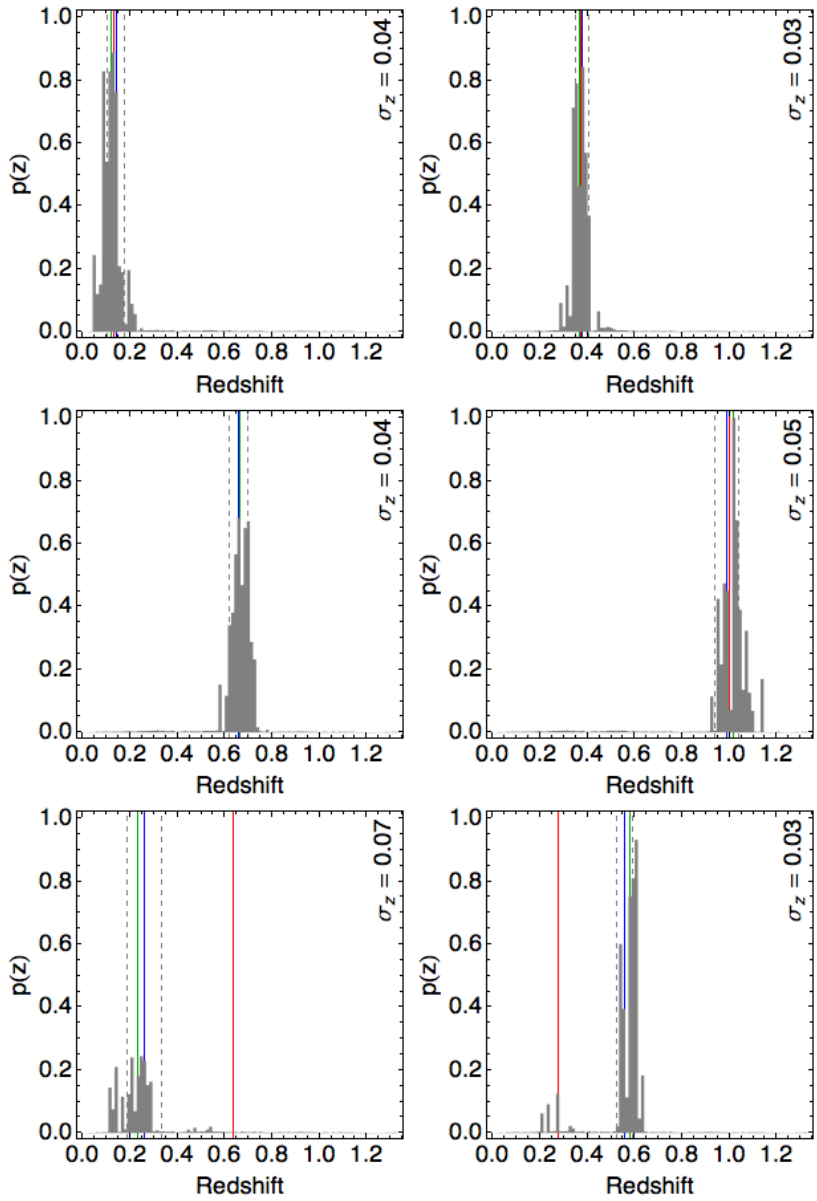

Figure 3.4: Typical $p(z)$ unnormalized probability distributions (probability per redshift bin) produced by ArborZ. In the top four panels we see $p(z)$ distributions at several redshifts. The true redshift is marked by a red line, the mean of the distribution is marked by a blue line, the median is marked with a green line, and the 1σ error range (enclosing half of the the central 68% distribution) is delimited with gray, dashed lines. We see that ArborZ distributions typically have their support over a relatively narrow region in redshift space. In the bottom two panels we show examples of catastrophic failure, where $|z_{\rm phot} - z_{\rm spec}| > 3\sigma_z$. In the bottom left is an example of $p(z)$ with low probability everywhere, which could be removed with a cut on the maximum height $p_{\rm peak}$ of the distribution (discussed later). In the bottom right we show an example of an incorrect $z_{\rm phot}$ estimate, but where the distribution has a secondary peak at the true redshift, illustrating the strength of a $p(z)$ algorithm, since these secondary peaks will contribute to correctly estimating the number distributions $N(z)$ in each redshift bin.

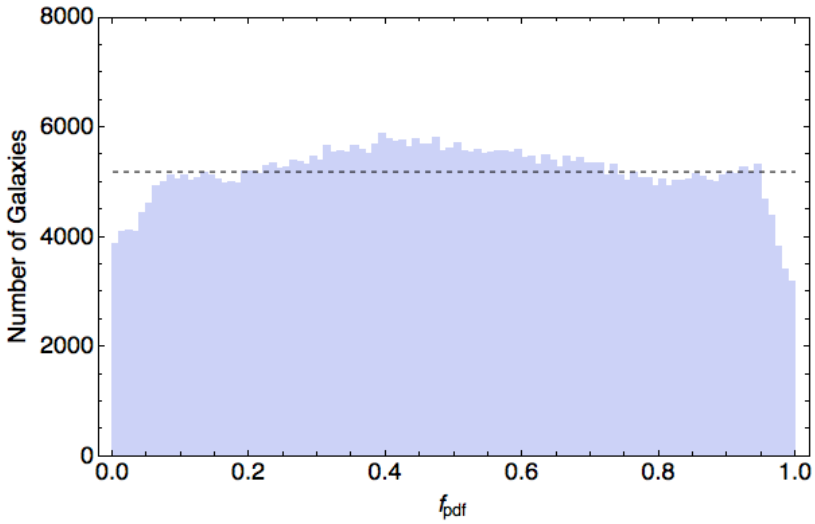

Figure 3.5: The distribution of $f_{\rm pdf}$, the fraction of area in $p(z)$ below $z_{\rm spec}$. An unbiased estimator will yield a flat distribution. The dashed line is the mean of the distribution.

by degeneracies in the galaxy magnitude-color relation; photometric observables, in other words, do not exhibit a one-to-one correspondence with redshift. The second phenomenon is known as *bias*. It is caused by an algorithm being incapable of assigning redshift estimates beyond the domain of its training set. That is, near $z = 0$, it can only overestimate redshift, since negative redshifts cannot be assigned; similarly, at high redshift, it has never encountered redshifts beyond those in its training set, and so it will necessarily underestimate the redshift at this extreme. As long as bias is well-characterized, it can be calibrated out of any cosmological calculations.

We now compare to Figure 3.6b, where we plot $z_{\rm pdf}$ (the sum of individual $p(z)$s) versus true redshift, where $z_{\rm pdf}$ is calculated by summing $p(z)$ distributions of galaxies whose true redshifts lie in the same redshift bin. We see a improved scenario: the degeneracies are gone and bias is all but eliminated. This is because $p(z)$ can properly contribute probability at all statistically meaningful redshift ranges, unlike a best-estimate $p(z)$, which discards information and, being the mean of a distribution, necessarily exhibits bias at the extremes of its range.

To further illustrate the performance of ArborZ, we can reconstruct the entire

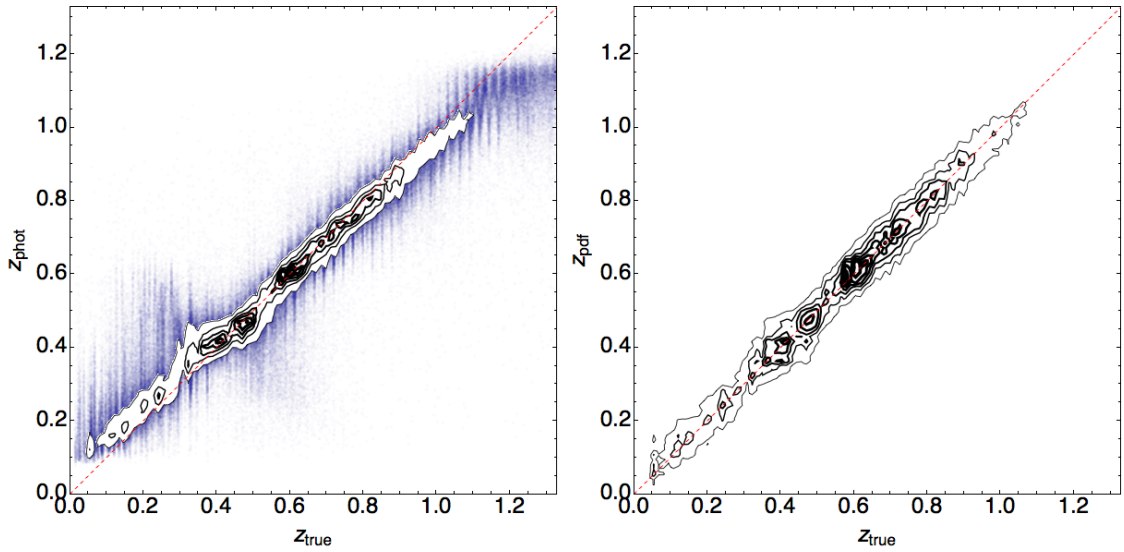

(a) z_{phot} (b) $p(z)$

Figure 3.6: The z_{phot} and $p(z)$ relationship with true redshift. In Figure 3.6a, we see flattening at $z_{\text{true}} = z_{\text{phot}} = 0.4$ due to degeneracies, as well as evidence of bias at the extremes of the redshift range. In comparison, Figure 3.6b (for which a scatter plot cannot be drawn) shows much improved fidelity, with bias nearly removed.

number distribution $N(z)$ of galaxies in the sample. In Figure 3.7 we plot the distribution of z_{phot} as well as the sum of the $p(z)$ distributions, z_{pdf}. Although there is the degeneracy peak at $z = 0.4$ and bias is evident at the extremes of the redshift range, z_{phot} still does a reasonably good job of reconstructing the true redshift distribution. On the other hand, we see that $p(z)$ presents an even better distribution estimate. This allows it to be used for weak lensing or the galaxy-galaxy correlation in baryon acoustic oscillations, both of which do not necessarily require accurate z_{phot} estimates, but rather reliable and unbiased $N(z)$ distributions.

To quantify the fidelity of our estimates N^{set} of the redshift number distribution, we define a goodness-of-fit parameter χ^2:

$$\chi^2 = \sum_i^{N_{\text{bins}}} \left(\frac{N_i^{\text{est}} - N_i^{\text{true}}}{N_i^{\text{true}}} \right)^2 \tag{3.24}$$

where the sum runs over each redshift bin. For the $N(z)$ distribution estimated with z_{phot}, we find $\chi^2 = 25.5$, and for $z_{\text{phot}}^{\text{med}}$ we find $\chi^2 = 13.4$. For the $p(z)$ estimate,

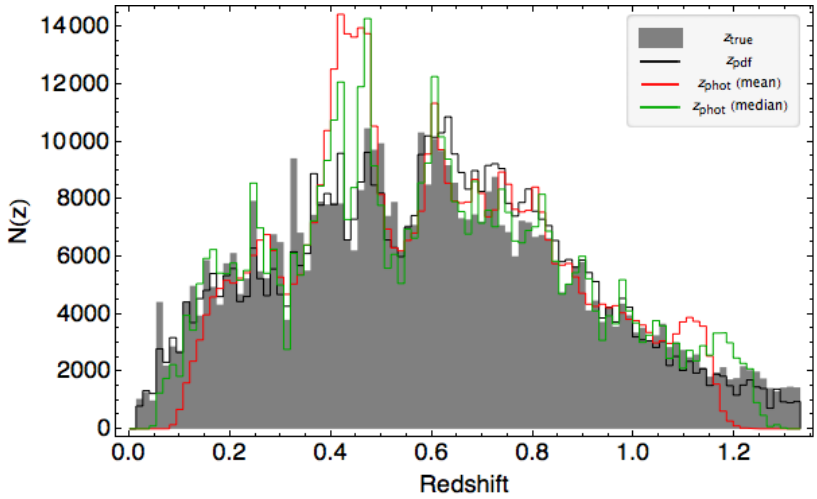

Figure 3.7: The number distribution $N(z)$ of galaxies in the catalog constructed using two methods. The first plots the distribution of $z_{\rm phot}$, the best-estimate photo-z. This does a reasonable job, but exhibits strong symptoms of bias. On the other hand, the $p(z)$ method does a much better job reconstructing the underlying redshift distribution.

$\chi^2 = 3.7$, showing $p(z)$ a better estimator of $N(z)$.

Now that we've seen the effect bias has on photo-z estimates, let's be more quantitative. As we mentioned earlier, a well-characterized bias measurement can be used to calibrate out any effects it may have on cosmology. We estimate bias as a function of redshift by calculating the mean of the residual $z_{\rm phot} - z_{\rm spec}$ in each redshift bin. When characterizing photo-z algorithms, it is useful to consider redshift bins in $z_{\rm spec}$; when calibrating photo-z bias for cosmological application, redshift bins in $z_{\rm phot}$ are used, since $z_{\rm phot}$ is the observable. In Figure 3.8 we show bias using both $p(z)$ contours and $z_{\rm phot}$, plotted against both $z_{\rm true}$ and $z_{\rm phot}$. Additionally, we show the bias calculated with $z_{\rm phot}$ after applying a cut on the height $p_{\rm peak}$ of the largest peak in the unnormalized $p(z)$ distribution. We demanded that $p_{\rm peak} \geq 0.99$, keeping the objects which the BDTs had the highest confidence in. This cut effectively selects objects which were most important to the BDTs during the training process.

Similarly, rather than exploring the mean of the residual distribution, we can look at its width. We can calculate the width in two ways: a traditional standard deviation

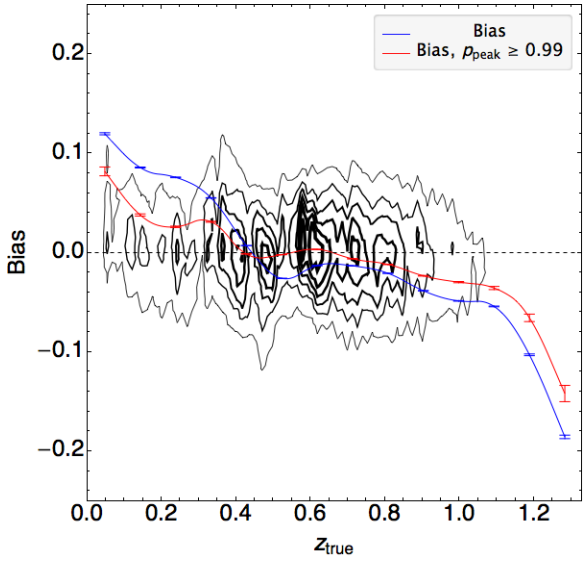

(a) Bias versus z_{true}

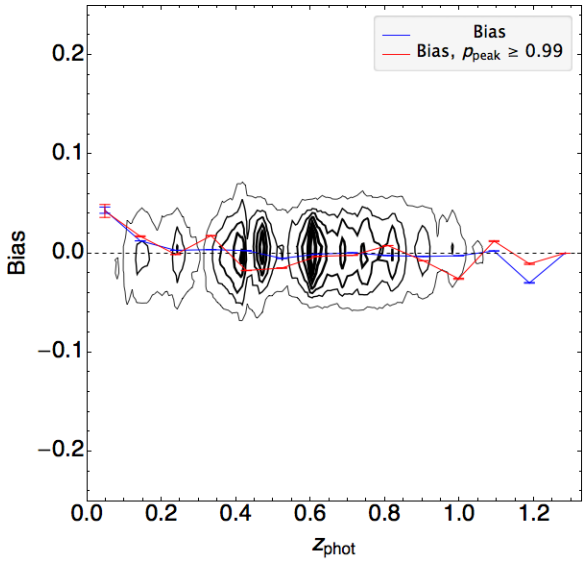

(b) Bias versus z_{phot}

Figure 3.8: Photo-z bias in ArborZ. The black, dashed line represents zero bias. The blue line is the bias calculated using z_{phot}, and the red line is the bias calculated after a cut $p_{\text{peak}} \geq 0.99$ is applied to the $p(z)$ distributions. The black contours are the bias contours calculated using $p(z)$. We see that the $p(z)$ contours contain less biased information than provided by only z_{phot}. Figure 3.8a offers insight into the performance of a photo-z estimation algorithm, but Figure 3.8b encodes the calibration data needed to remove bias in real surveys.

(which will be more sensitive to outliers) or as half of the width of the central 68% of the distribution (which will give us a more accurate estimate of the error for the majority of our sample). We do this in Figure 3.9. Again, we see marked improvement in the error measure when we apply a cut on $p_{\rm peak}$.

As we discussed earlier, ArborZ can also produce error estimates σ_z for each $z_{\rm phot}$ photo-z estimate. If these errors are honest, normally distributed errors, then the normalized errors

$$\sigma_z^{\rm norm} = \frac{z_{\rm phot} - z_{\rm spec}}{\sigma_z} \tag{3.25}$$

should be normally distributed with a mean of zero and a width of one. We show this normalized error distribution in Figure 3.10. With a mean $\mu = -0.04$ and a width of $\sigma = 1.02$, we conclude that these error estimates are, indeed, good estimates of the redshift uncertainty.

As one may expect, there is a relationship between the peak $p_{\rm peak}$ of a galaxy's $p(z)$ probability distribution and its photo-z error σ_z. The cause of this relationship is intuitive: if ArborZ $p(z)$s are correctly representing the probability distribution of the galaxy, then a galaxy which is poorly estimated by all BDTs (i.e., has a low $p_{\rm peak}$) should likewise have a greater uncertainty in its distribution (i.e., a greater width in $p(z)$). This relationship is easy to verify. In Figure 3.11 we plot these two quantities, and see a strong, and relatively tight, relationship. Therefore, we can interpret cuts on $p_{\rm peak}$ as equivalent to cuts on the photo-z error, and so samples with large $p_{\rm peak}$ indeed represent better measured photometric redshifts.

It is important to keep in mind that this discussion is not meant to declare $z_{\rm phot}$ anathema. On the contrary, ArborZ's $z_{\rm phot}$ performs as well or better than many other photo-z algorithms and still is a reasonable estimate of the true redshift, particularly if one calibrates out (or simply avoids) the areas of strong bias at the extremes of the redshift range. Rather, we hope to impress the virtues of $p(z)$ which cannot be

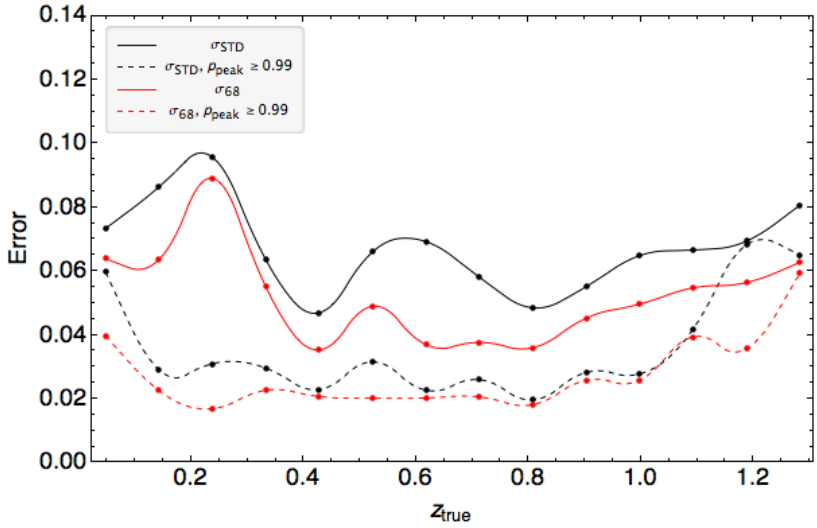

(a) Error versus z_{true}

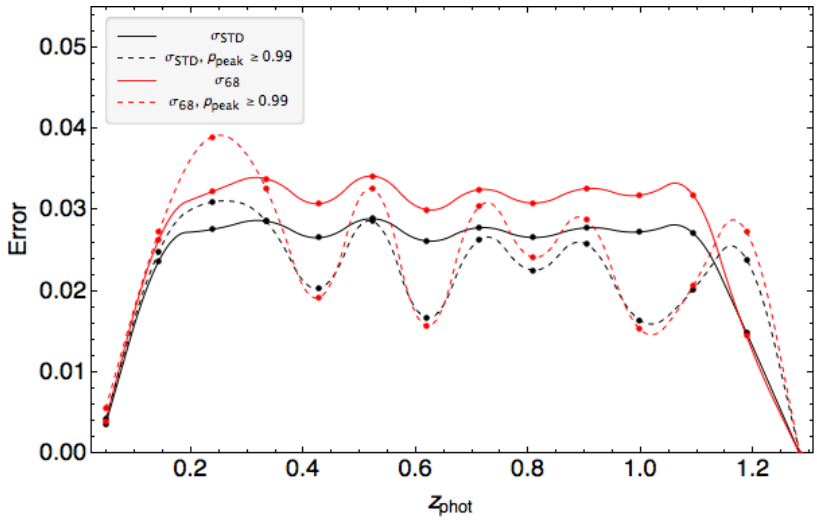

(b) Error versus z_{phot}

Figure 3.9: Width of residuals in ArborZ. We show this width calculated in two ways: the standard deviation and half of the width of the central 68% of the distribution. We also show both measures with and without an additional cut on p_{peak}, the height of the largest peak in the $p(z)$ distribution.

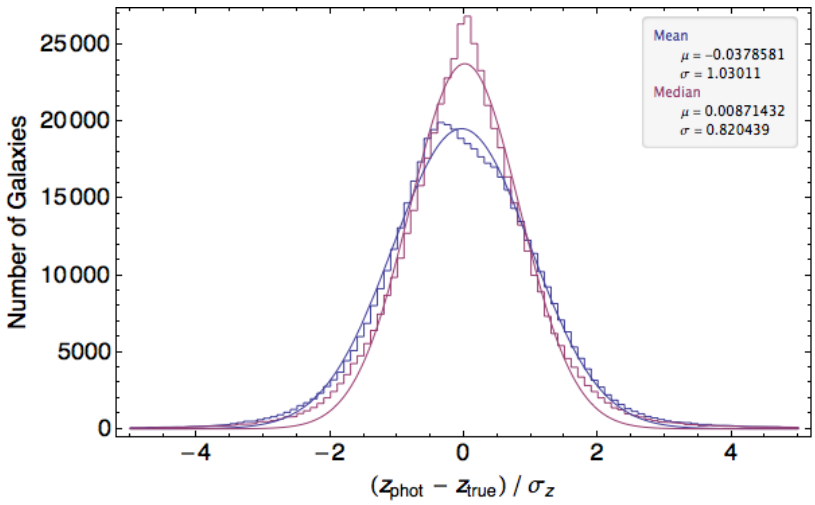

Figure 3.10: The distribution of normalized errors from ArborZ, using both the mean $z_{\rm phot}$ and median $z_{\rm phot}^{\rm med}$ to calculate best-estimate photo-z. If the σ_z error estimates are honest estimates, then this distribution should be Gaussian, with central mean and a width of unity. The calculated errors are good error estimates on $z_{\rm phot}$, but overestimates when applied to $z_{\rm phot}^{\rm med}$.

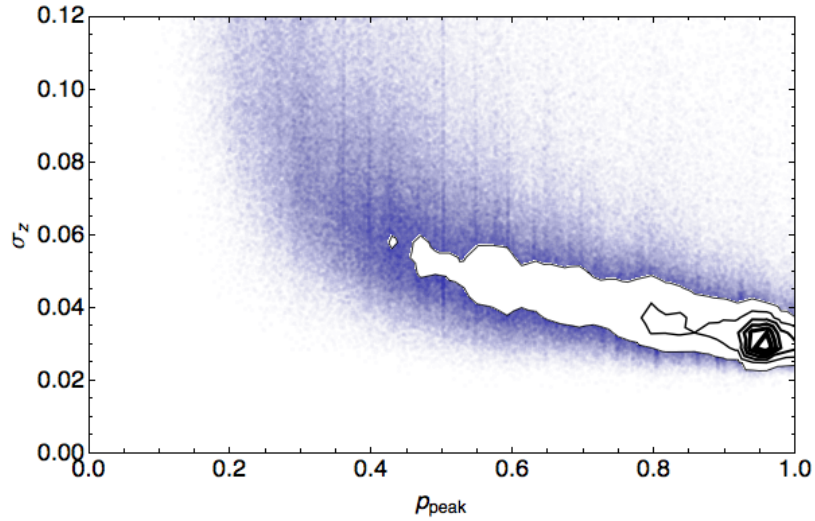

Figure 3.11: The relationship between $p_{\rm peak}$ and photo-z error σ_z. We see a clear correlation between these quantities, allowing us to interpret cuts on $p_{\rm peak}$ as equivalent cuts on photo-z uncertainty.

replicated or emulated by a simple, single-number redshift estimator. Using complete probability distributions reduces bias and provides a cleaner, clearer signal for doing cosmology. Whenever possible one should keep the full distribution or estimate a redshift-dependent quantity $f(z)$ of interest by integrating against $p(z)$:

$$\langle f(z) \rangle = \frac{\int_0^\infty f(z)p(z)dz}{\int_0^\infty p(z)dz} \tag{3.26}$$

We can now quantify the uncertainties in ArborZ's redshift estimates to see how they will impact the dark energy figure of merit. In Figure 3.12, we plot these uncertainties and compare them to the the DES science requirements and other figure of merit degradation limits. We see that the photo-z estimates are within the DES requirements, and nearly every point is within the 10% degradation limit. This gives us confidence that ArborZ will be able to produce quality photo-zs for constraining cosmology.

In the examples we've discussed so far, we have only used photometric magnitudes to train ArborZ. There is no requirement on using only magnitudes, let alone magnitudes at all. Another option would be to use colors (differences in magnitudes). The benefit of colors is that, for faint galaxies in the photometric target set, the distribution of colors may be more representative of the distribution of colors in the training set. On the other hand, the errors in the magnitude measurements will propagate into the colors, which, being differences of magnitudes, will become relatively large. In practice, training on colors (alone or in addition to magnitudes) yields very similar results, but with slightly larger photo-z errors.

What of the myriad other variables in a galaxy catalog? Certainly we expect position-based correlations, since clustered galaxies are more likely to be at the same redshift than randomly chosen galaxies. On the other hand, sample variance will dominate these correlations in all realistic training sets. Careful thought would need to go into integrating positions into any photometric redshift algorithm. We have also

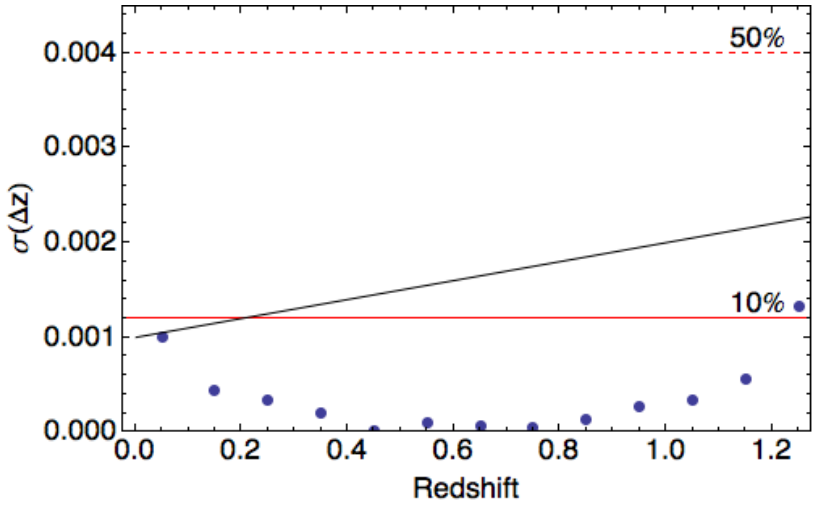

(a) Uncertainty in the photo-z bias

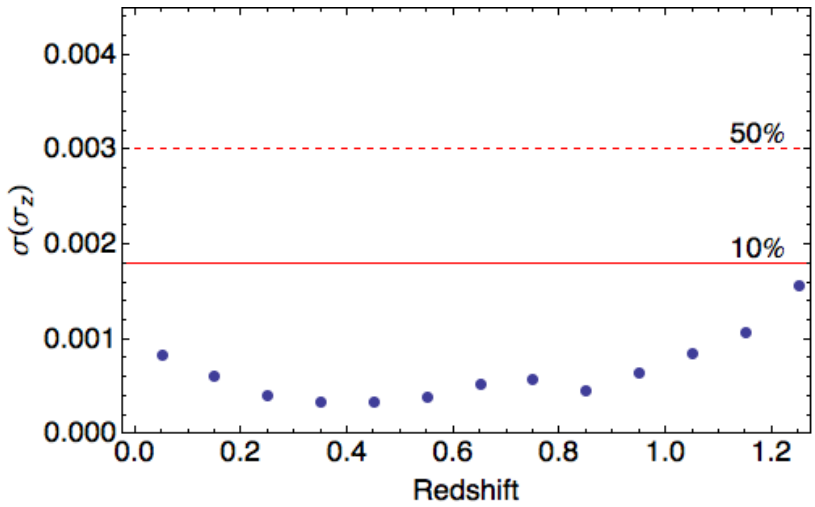

(b) Uncertainty in the photo-z error

Figure 3.12: The uncertainty in ArborZ's photo-z estimates in the mock catalog. The points indicate the uncertainties. The solid red line marks the 10% degradation in the dark energy figure of merit. The dashed red line marks the 50% degradation. In the upper plot, the DES science requirement is the solid black line; in the lower plot, the DES science requirement is the dashed red line, the same as the 50% mark.

tried adding ellipticity and magnitude errors to the mix of input variables without substantial improvement. Rather, we caution against using too many variables, since one will fall immediately into the trap of overtraining. A few, well-chosen variables with known correlations to redshift (e.g., magnitudes or colors) will go further than any haphazard amalgamation.

3.5 Optimization

Having discussed the general performance of ArborZ in a simulated galaxy catalog, we now turn our discussion toward subtleties in the algorithm. There are many different tuning knobs we can tweak, parameters in the training algorithm. Some of these can impact results, whereas others are negligible. In some cases, we can increase the speed of the algorithm without impacting results. Our test machine was a 2.7 GHz quad-core (eight virtual cores) i7 processor with 16 GB of RAM.

We begin by discussing timing—how long it takes ArborZ to perform its calculations. The training algorithm itself has two natural places to "parallelize," that is, divide computational time across multiple processing cores. The first is at the DT growing level. Whenever a split is made, the DT continues to grow left and right branches, containing disjoint datasets. Therefore, new threads (independently managed sets of processor instructions) can be spawned to grow either side of the tree. The second opportunity for parallelizing is at the ArborZ level: each BDT can be constructed in its own thread. In practice, scheduling more threads than cores will not lead to increased performance (indeed, moving shared resources and the time spent scheduling many threads generally degrades performance). Our test machine, therefore, should spawn $N_{\text{threads}} = 8$ threads to be maximally efficient. Of the two parallelizing options, it is best to capitalize on the latter, since the former will exhibit diminishing returns as we delves deeper into the DT; moreover, spawning threads will

take time, and so they should be spawned at the highest levels of the program. Since each BDT will be constructing many DTs, it is better to spawn threads to manage BDTs rather than having many short-lived threads at the DT level. When we are evaluating, rather than training, we can simply divide the target observables among N_{threads} different tasks.

For the datasets described in Section 3.4, training 100 BDTs on 100 000 galaxies required 1087 seconds of wall time (actually time as measured by a watch) and 7 934 seconds of CPU time (the sum of the wall times each CPU spent, but only when scheduled on this task). Evaluation on the 519 116 galaxies in the target sample took 421 seconds of wall time and 3 124 seconds of CPU time. Thus, ArborZ can run quite quickly on any modern computer, even when working with the large catalogs produced by modern surveys.

There is another, non-threaded possibility for improving ArborZ's training speed. When training individual DTs, the optimal cuts are found by stepping across the entire range of each variable. Training vectors will similar values for a particular variable will result in similar cut qualities. Therefore, instead of stepping across each possible cut, we could make larger jumps at the cost of small changes in cut quality. We expect this trade-off to be corrected by the BDT boosting algorithm; if we choose a suboptimal cut, boosting will not only alter the weights to correct for it, but also assign that particular tree a lower weight in the evaluation process.

To test this, we performed the same training and evaluation described previously, but only checked at most 20 different split positions at each branch. Training required 994 seconds of wall time (7 408 seconds of CPU time), an improvement of approximately 8%. We so no decrease in photo-z quality, and so this mechanism provides a simple way to reduce training time slightly.

We now turn our attention away from timing optimizations to result-driven optimizations. There are many parameters we can change would could potentially impact

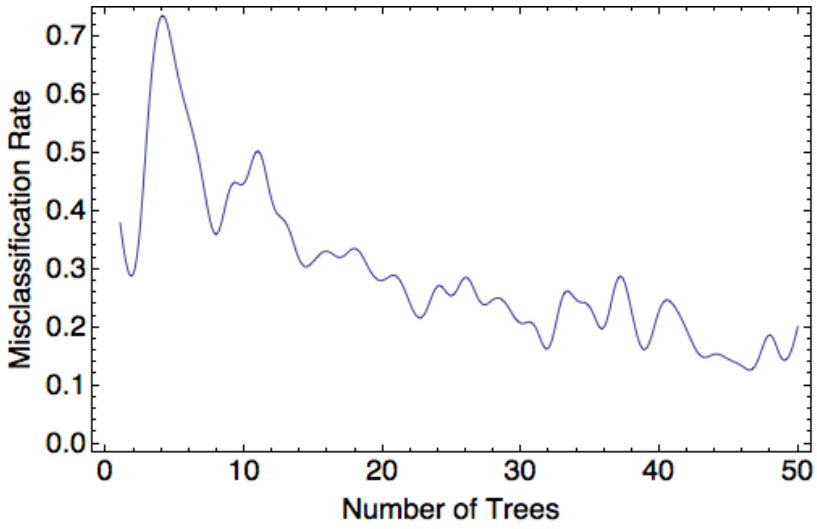

Figure 3.13: Decision tree misclassification error as a function of BDT forest size. For nearly all training sets, the misclassification rate stabilizes near 40–50 trees for forest.

performance. One such example is the initial BDT weights. Typically, they are set to unity for the first DT in the forest. However, some photo-z estimation methods calculate weights for the galaxies in training set depending on how important they are in covering the observable space of the evaluation sample (Cunha et al. 2009; Lima, Cunha, et al. 2008). Similar priors could be applied to the ArborZ training set in the form of initial weights. However, when we calculate and apply such weights, we find no difference in the results; the BDT re-weighting algorithm drives the weights to the same optimal values anyway, making additional priors unnecessary.

Other parameters which affect the training process are the DT bucket size, the DT separation measure, and the signal threshold (usually set to 0.5 as discussed in Section 3.1.2). However, we see little or no difference in the photo-z results when modifying these parameters. Typically, we use 50 DTs per BDT forest. We determine this value empirically; DT misclassification rate is a rapidly decreasing function of forest size (see Figure 3.13), and performance results stabilize with forests of around 50 trees.

Chapter IV

Applications of ArborZ

We have seen that ArborZ produces excellent estimates of a galaxy's redshift; moreover, the $p(z)$ distributions constructed by ArborZ enable better constraints on cosmology through their more faithful representation of redshift and uncertainty. In this chapter, we apply ArborZ to two major surveys—the Sloan Digital Sky Survey and the Dark Energy Survey—and present photometric redshift catalogs for each, paving the way for many cosmological probes.

4.1 The Sloan Digital Sky Survey

The Sloan Digital Sky Survey (SDSS; York et al. 2000) is one of the largest—and certainly one of the most influential—optical imaging surveys ever undertaken. Taking data from 1998 to present, the SDSS has imaged over $14\,000\,\mathrm{deg}^2$ of the northern sky, yielded measurements in *ugriz* bands for over 2×10^8 galaxies. Nearly two million optical galaxy spectra have been measured since its inception, producing in the largest spectroscopic redshift catalog ever obtained (Blanton, H. Lin, et al. 2003; Dawson et al. 2013). The most recent public data release of the SDSS is data release 10 (DR10; Ahn et al. 2013), which is the catalog we use here.

4.1.1 Selecting a Clean Galaxy Sample

We select a clean galaxy sample from DR10 by imposing a number of cuts on the data. There are three different magnitude measurements we use in making these cuts: $m_{\rm psf}$, derived from a best-fitting PSF model; $m_{\rm model}$, which is the better of an exponential profile and a de Vaucouleurs' profile; and $m_{\rm cmodel}$, which is a combination of an exponential and a de Vaucouleurs' profile. Our cuts follow the standard cuts recommended by the SDSS team for selecting a clean galaxy sample, as well as cuts recommended in Sheldon et al. 2012. We summarize these cuts below[1].

- We perform star-galaxy separation by looking at the concentration c of each object:

$$c = m_{\rm psf} - m_{\rm model} \tag{4.1}$$

  For stars, $c = 0$ within errors. Following the SDSS pipeline, we define galaxies as having $c > 0.145$ (with c derived from the summed fluxes in all bandpasses).

- We select unique, resolved objects by checking the RESOLVE_PRIMARY status flag: RESOLVE_PRIMARY & SURVEY_PRIMARY $\neq 0$.

- We require detections in both r and i bands:

$$\text{FLAGS\_[RI] \& (BINNED\_1 | BINNED\_2 | BINNED\_4)} \neq 0$$

- We reject objects which are flagged for any of the following: SATURATED, BRIGHT, DEBLEND_TOO_MANY_PEAKS, PEAKCENTER, NOTCHECKED, NOPROFILE.

- We reject blended objects (BLENDED $\neq 0$) which were not successfully deblended (NODEBLEND $\neq 0$).

[1]We adopt the computer science notation where "&" indicates the bit-wise AND operator and "|" indicates the bit-wise OR operator.

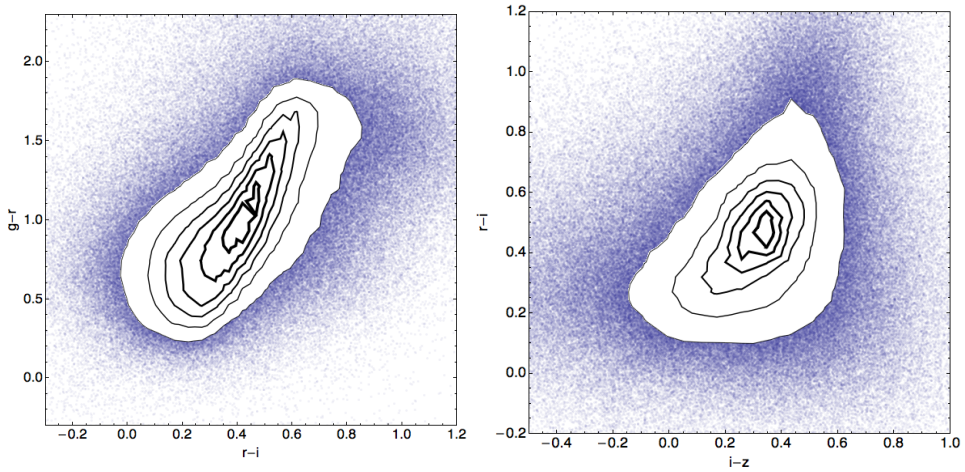

(a) $g - r$ vs $r - i$ (b) $r-$ vs $i - z$

Figure 4.1: Color-color plots showing the distribution of colors in the SDSS DR10 photometry.

- To reduce noise, we place magnitude cuts 21, 22, 22, 20.5, and 20.1 on the extinction-corrected (D. J. Schlegel, D. P. Finkbeiner, and M. Davis 1998) m_{model} magnitudes in $ugriz$ bands, respectively.

- We require that the extinction-corrected r-band m_{model} magnitudes be between 15 and 29. Similarly, we require that the extinction-corrected r-band m_{cmodel} magnitudes be between 15 and 21.8.

- In two runs, the u-band amplifier was experiencing problems. We therefore remove objects from runs 2190 and 2189 which were flagged with NOTCHECKED_CENTER or LOCAL_EDGE.

After applying these cuts, we obtain a imaging catalog of 80 882 179 galaxies. The $g - r$ versus $r - i$ and $r - i$ versus $i - z$ color-color plots are shown in Figure 4.1. The magnitude distributions are shown in Figure 4.3.

4.1.2 Training Catalog

In order to evaluate photo-zs for the galaxies in the imaging catalog, we need to assemble a training set of galaxies with known redshifts. Using the adaptive matching technique described in Appendix B, we can combine redshifts from multiple spectroscopic sources, matching these redshifts to the optically-imaged galaxies in DR10. These redshift sources are described as follows.

- **SDSS**: The vast majority of redshifts were obtained as part of the SDSS survey itself. The catalog consists of three samples: a magnitude-limited ($r_{\mathrm{Petro}} < 17.77$) legacy "main" sample from SDSS-I/II with a median redshift of $z \approx 0.10$ (Strauss et al. 2002), a luminous red galaxy (LRG) sample which is volume-limited to $z \approx 0.38$ but extends out to $z \approx 0.55$ (D. J. Eisenstein, J. Annis, et al. 2001), and the BOSS sample of approximately 9×10^5 redshifts with median redshift $z \approx 0.51$ (Dawson et al. 2013). After applying quality cuts, 1 284 330 redshifts were kept.

- **2dFGRS**: 31 962 redshifts from the 2dF Galaxy Redshift Survey (Colless et al. 2001).

- **2SLAQ**: 5 548 redshifts from the 2dF-SDSS LRG and QSO Survey (Cannon et al. 2006).

- **AGES**: 1 502 redshifts from the AGN and Galaxy Evolution Survey (Kochanek et al. 2012).

- **CFRS**: 178 redshifts from the Canada-France Redshift Survey (Le Fèvre, Crampton, et al. 1995; Lilly, Le Fèvre, Crampton, et al. 1995).

- **CNOC2**: 882 redshifts from the Canadian Network for Observational Cosmology Field Galaxy Survey (H. K. C. Yee et al. 2000).

- **DEEP2 / DEEP3**: 1 729 redshifts from the DEEP Extragalactic Evolutionary Probe 2 (M. Davis et al. 2003; Newman et al. 2012) and Probe 3 (Cooper, Aird, et al. 2011).
- **GAMA**: 44 552 redshifts from the Galaxy and Mass Assembly survey (Driver et al. 2011).
- **MGCz**: 1 741 redshifts from the Millennium Galaxy Catalogue redshift survey (Allen et al. 2006).
- **OzDES**: 13 746 redshifts obtained on the Anglo-Australian Telescope as part of the OzDES program, the primary spectroscopic followup program for the Dark Energy Survey (DES).
- **PRIMUS**: 17 347 redshifts from the Prism Multi-Object Survey (Coil et al. 2011).
- **TKRS**: 190 redshifts from the Team Keck Redshift Survey (Wirth et al. 2004).
- **VIPERS**: 5 923 redshifts from the VIMOS Public Extragalactic Redshift Survey (Guzzo et al. 2013).
- **VVDS**: 3 488 redshifts from the VIRMOS-VLT Deep Survey (Le Fèvre, Vettolani, et al. 2005).
- **WiggleZ**: 34 135 redshifts from the WiggleZ Dark Energy Survey (Drinkwater et al. 2010).
- **zCOSMOS**: 1 004 redshifts from the Cosmological Evolution Survey (COSMOS) field (Lilly, Le Brun, et al. 2009; Lilly, Le Fèvre, Renzini, et al. 2007).

In total, our training set of spectroscopic redshifts matched to DR10 imaging contains 1 448 257 galaxy redshifts. The overall redshift distribution, as well as cumulative

distributions for each source catalog, is shown in Figure 4.2. The entire redshift range is well-represented out to $z \approx 0.8$, with a tail extended to $z \approx 1.0$. The magnitude distributions of this spectroscopic training set is shown in Figure 4.3.

4.1.3 Testing Unrepresentation

When we considering Figure 4.3, we see that the training set has only limited coverage (under-representation) at faint magnitudes, and excess coverage (over-representation) at bright magnitudes. Before attempting to produce a photo-z catalog for the SDSS, it is worth considering the question of observable unrepresentation: can ArborZ produce accurate redshift estimates if the distributions of training set observables differ significantly from the same observables' distributions in the target imaging catalog?

The author has investigated this problem in the past using the SDSS and GAMA surveys as examples, and found that ArborZ still produces good results (Sypniewski and D. W. Gerdes 2011). We briefly discuss the case of unrepresentation in the SDSS DR10 data set here. To clarify, we use the term "unrepresentation" to refer to different observables distributions in the training and target samples, typically where one or more instances of under- or over-representation are present, but where the training set still covers the same range of observable space as the target set. In the case of "non-representation," the target set would contain a region in observable space where no similar observables are present in the training. No empirical algorithm can hope to reliably extrapolate to these regions, and in such cases a template-based approach to photo-z estimation may be preferred.

To test ArborZ's performance, we split the spectroscopic training set discussed in Section 4.1.2 into two disjoint subsets: one containing redshifts from the SDSS survey only, and the second containing the remaining galaxies from the non-SDSS surveys. The redshift distributions of these two samples are shown in Figure 4.4, and the magnitude distributions are shown in Figure 4.5.

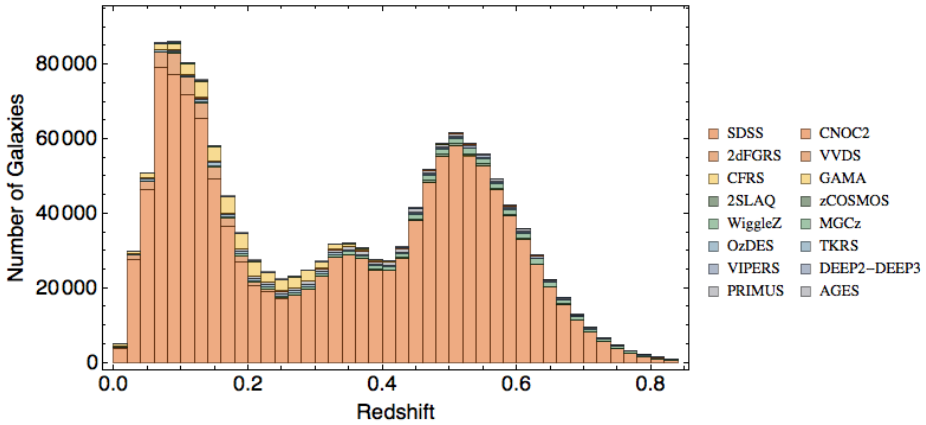

(a) Overall redshift distribution (stacked)

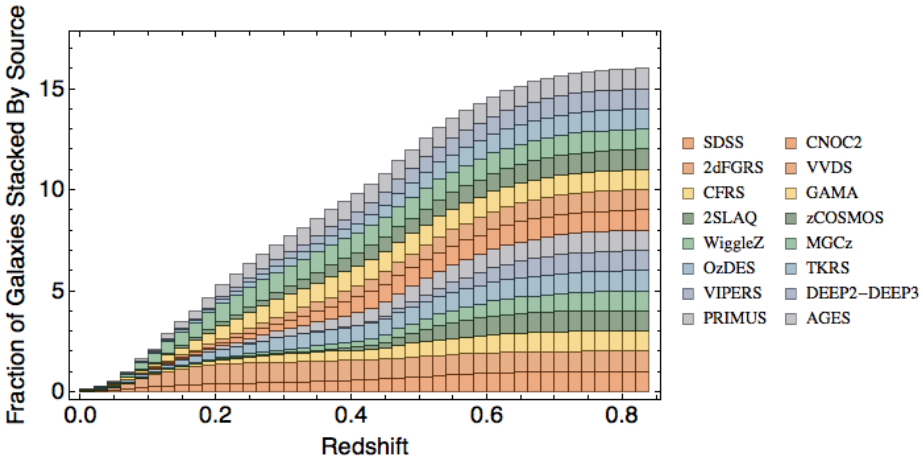

(b) Cumulative distribution functions (stacked)

Figure 4.2: The redshift distribution of the SDSS spectroscopic training set. In Figure 4.2a is the overall redshift distribution, including contributions from all sixteen source catalogs. In Figure 4.2b, we stack the cumulative distribution functions (CDFs) of the redshift distributions from each catalog. In both plots, SDSS is shown on the bottom of the stack, and AGES is shown on top.

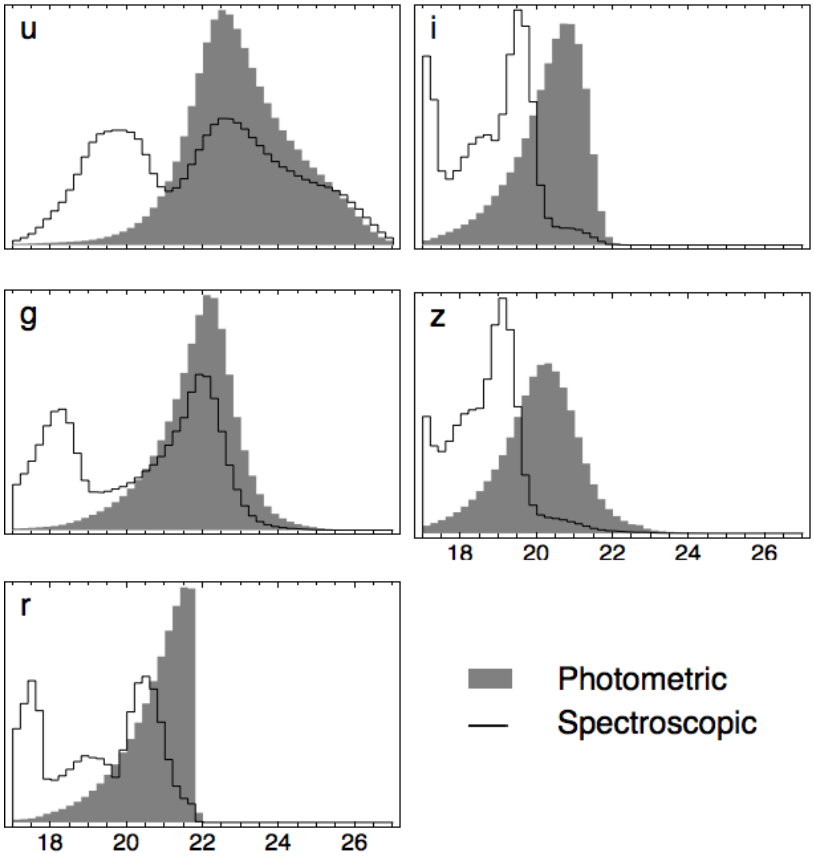

Figure 4.3: The distributions of $m_{\rm model}$ magnitudes in the SDSS DR10 imaging and the spectroscopic training set. Magnitude cuts are placed on both the extinction-corrected $m_{\rm model}$ magnitudes as well as the extinction-corrected $m_{\rm cmodel}$ magnitudes. Both distributions are normalized to more easily compare contributions at each magnitude.

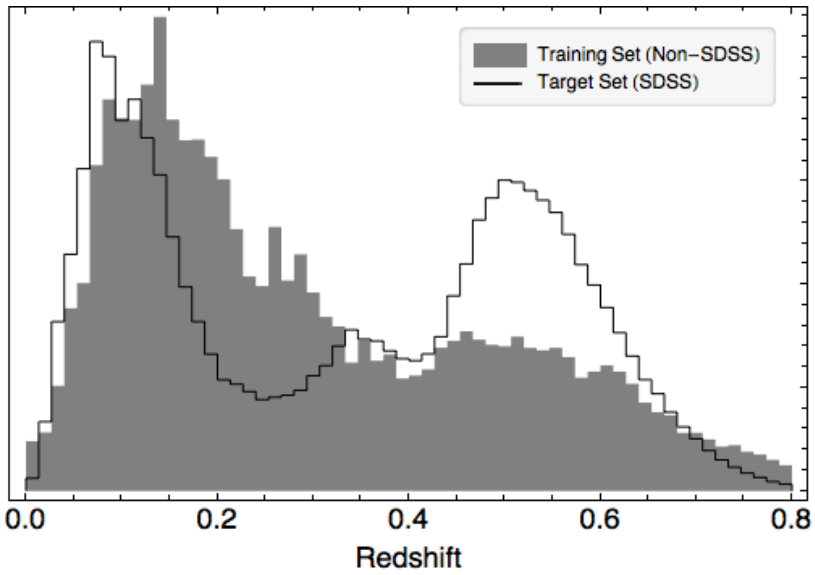

Figure 4.4: The redshift distributions of the training (non-SDSS) set galaxies and target (SDSS) set galaxies for the unrepresentation test.

We train ArborZ on the non-SDSS galaxy sample using 60 bins with $\sigma_{\rm res} = 0.03$. Figure 4.6 shows the stacked $p(z)$ measurements versus the spectroscopic redshift, and Figure 4.7 shows the reconstructed redshift distribution $N(z)$. In both cases, we see that the BDTs can properly weight the training set galaxies, preventing significantly degraded photo-z results when faced with different observable distributions in the training and target samples. As shown in Figure 4.8, even with unrepresentation, the uncertainty in the photo-z bias and error are still well-understood.

4.1.4 SDSS Photometric Redshift Catalog

Before presenting the final catalog, we perform a validation run, using the SDSS spectroscopic training set assembled in Section 4.1.2. We train ArborZ on 100 000 randomly selected galaxies from the large training set, and then evaluate it on the remaining galaxies. This gives us the best opportunity available for testing our training parameters. We used 80 equal-width bins out to redshift 0.8, with photo-z resolution $\sigma_{\rm res} = 0.02$.

In Figure 4.9a we show the reconstructed redshift distribution $N(z)$. Although the

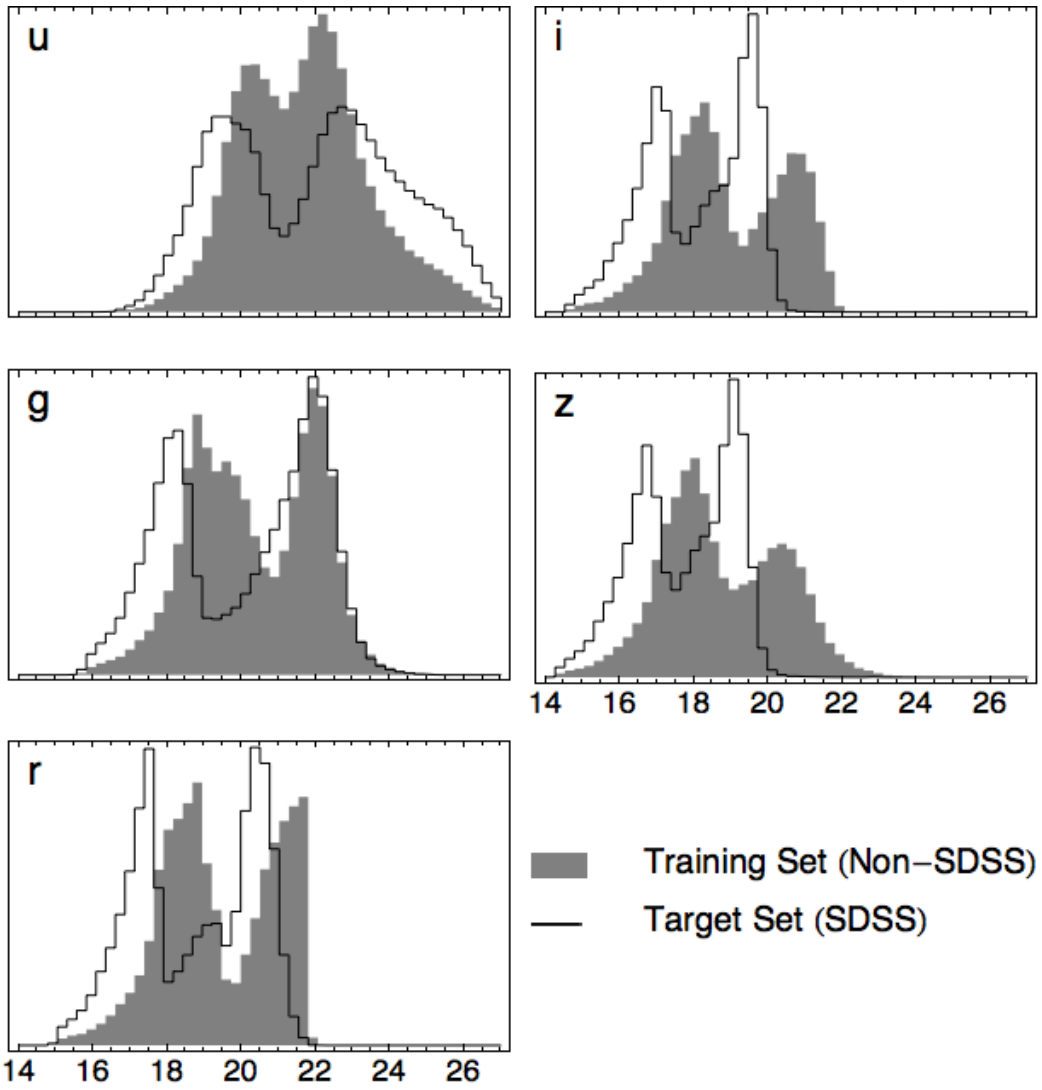

Figure 4.5: The magnitude distributions of the training (non-SDSS) set galaxies and target (SDSS) set galaxies for the unrepresentation test. The magnitude distributions are sufficiently different as to be worth validating ArborZ's performance before blindly trusting its results.

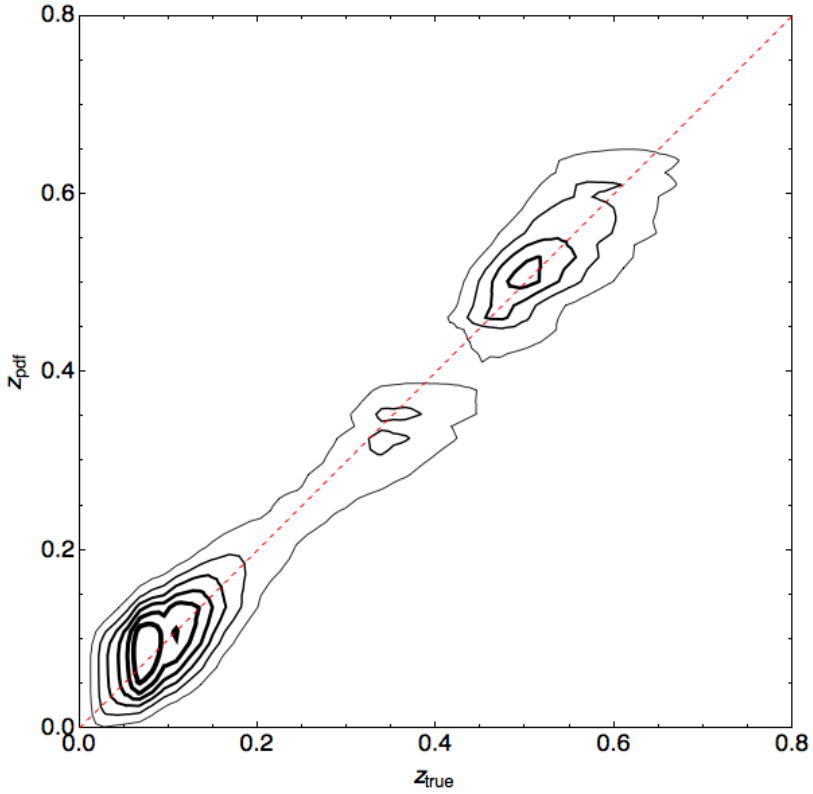

Figure 4.6: ArborZ $p(z)$ results plotted against $z_{\rm true}$, showing that $p(z)$ is still an accurate redshift measure in the case of unrepresentation.

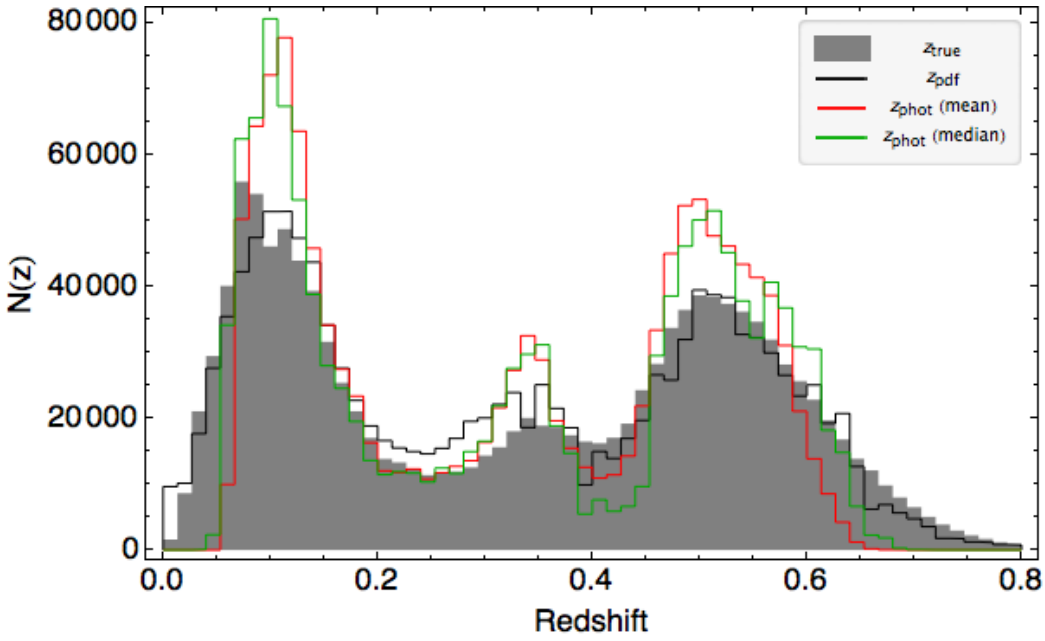

Figure 4.7: The constructed redshift distribution for the unrepresentation test. $p(z)$ shows a small overestimate near $z = 0.3$, but does an overall excellent job of reconstructing the underlying redshift distribution, even when the distribution of observables is different for the training and target sets.

best-estimate photo-z $z_{\rm phot}$ performs well, we see that $p(z)$ displays excellent fidelity with respect to the true redshift distribution. The χ^2 metric defined in Equation 3.24 is 22.3 for $z_{\rm phot}$, 19.2 for $z_{\rm phot}^{\rm med}$, and 3.1 for $z_{\rm pdf}$, demonstrating that $p(z)$ does an excellent job of recovering the true redshift distribution. We show the $p(z)$ versus $z_{\rm true}$ contours in Figure 4.9b.

Figure 4.10a shows the bias of ArborZ across the redshift range. Bias calculated using both $p(z)$ and $z_{\rm phot}$ shows low bias across the entire redshift range. We also note that cuts on the maximum height $\max p(z)$ of the unnormalized $p(z)$ distribution can yield lower bias. The width of the residual distribution is shown in Figure 4.10b. This average photo-z error over the entire redshift range, calculated as the width of the central 68% of the area of the residual distribution, is approximately 0.038. When we compare the photo-z uncertainties to dark energy figure of merit degradation limits, we find that we are well within the limits expected from photo-z estimates (see Figure 4.11).

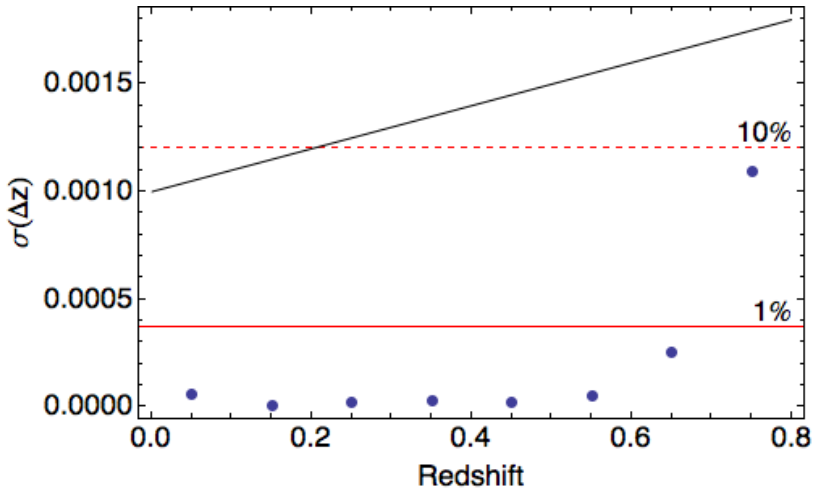

(a) Uncertainty in the photo-z bias

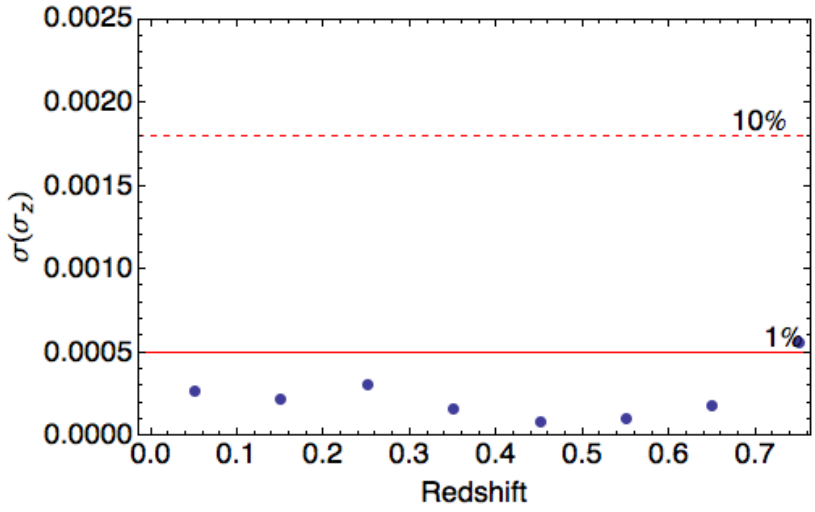

(b) Uncertainty in the photo-z error

Figure 4.8: The uncertainty in ArborZ's photo-z estimates in the SDSS unrepresentation test. The points indicate the uncertainties. The solid red line marks the 1% degradation in the dark energy figure of merit. The dashed red line marks the 10% degradation. In the upper plot, the DES science requirement is the solid black line; in the lower plot, the DES science requirement is above the range of the plot.

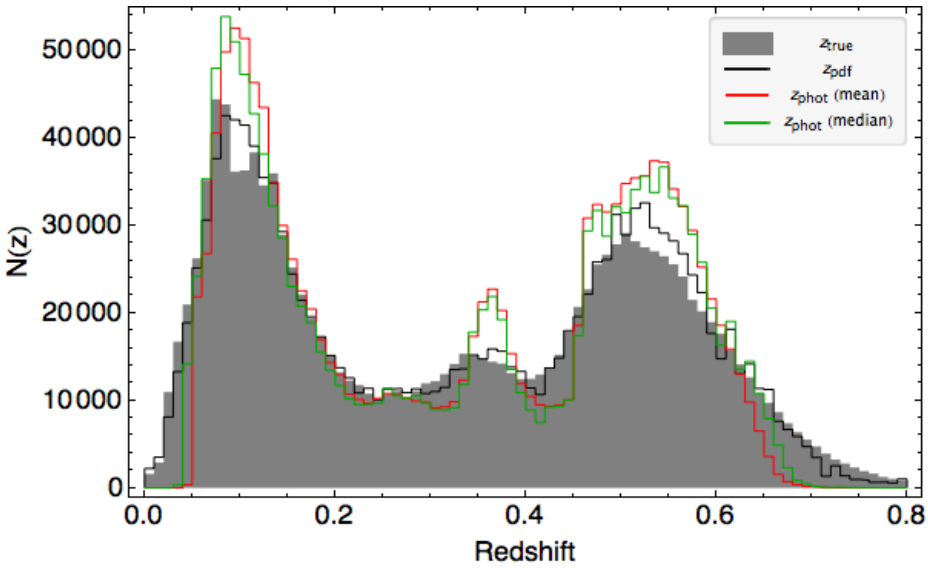

(a) Reconstructed redshift distribution

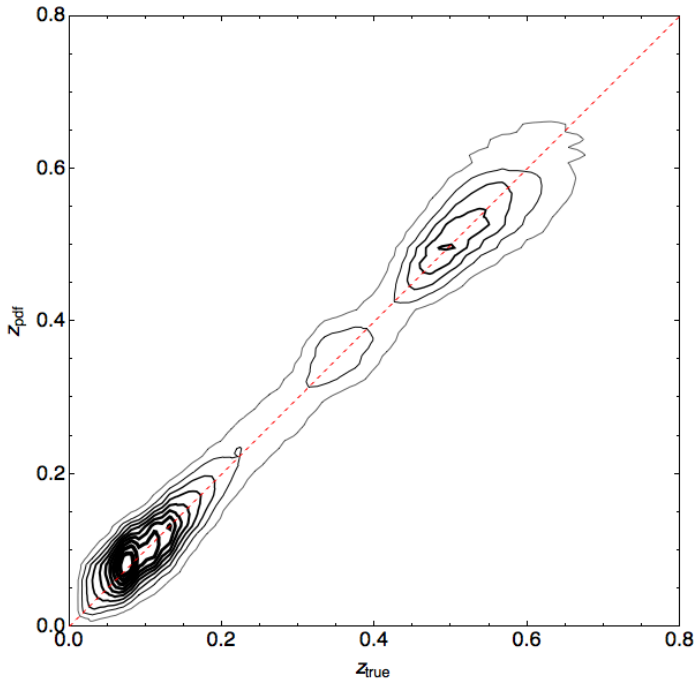

(b) $p(z)$ vs. z_{true} contours

Figure 4.9: The photo-z results in the SDSS validation sample, calculated by training on 100 000 random galaxies in the overall spectroscopic training set, and evaluating on the remainder. Good agreement between ArborZ and the true redshift is seen across the entire redshift range.

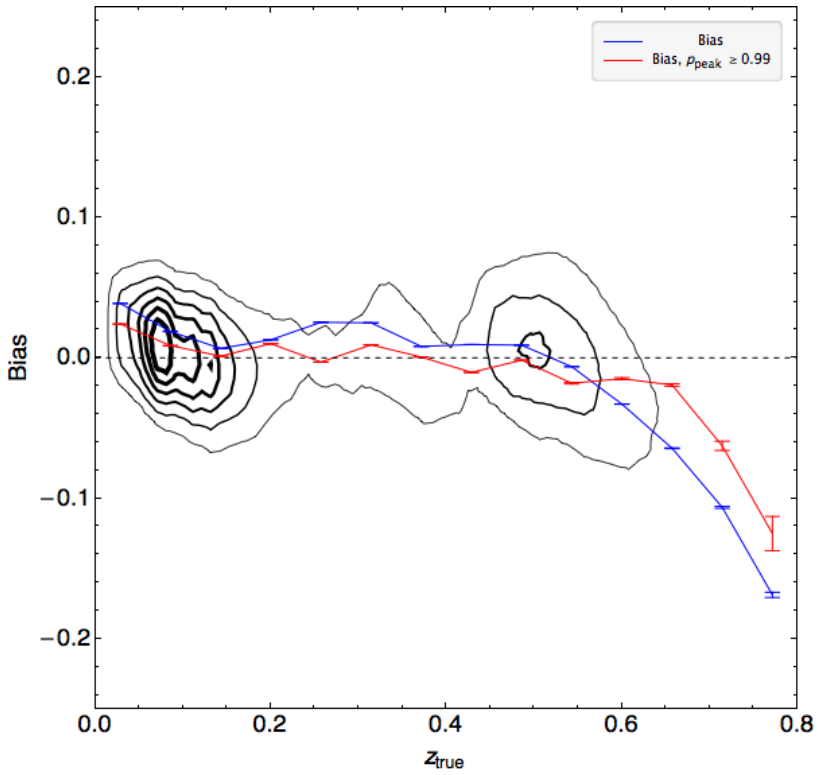

(a) Photo-z bias

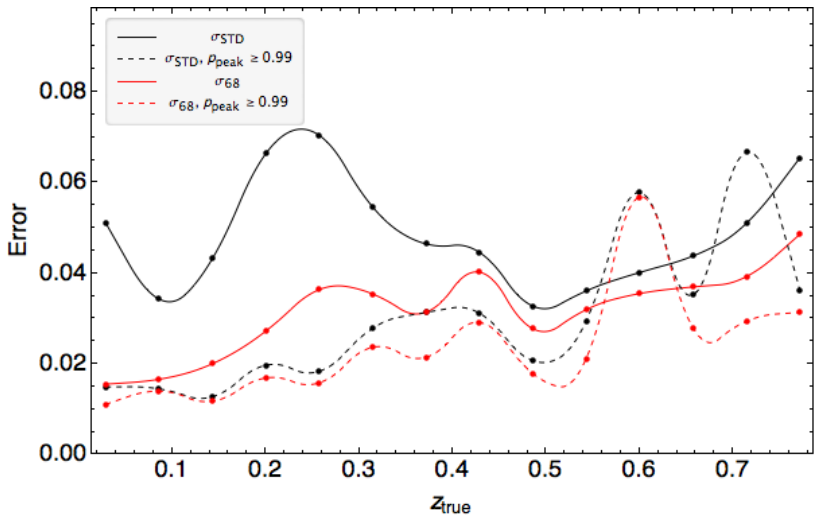

(b) Photo-z error

Figure 4.10: Bias and error plots for the SDSS validation sample, indicating good photo-z results across the entire redshift range. The peaks in the photo-z error in the $p_{\rm peak} > 0.99$ sample near $z = 0.6$ and $z = 0.72$ are due to low statistics.

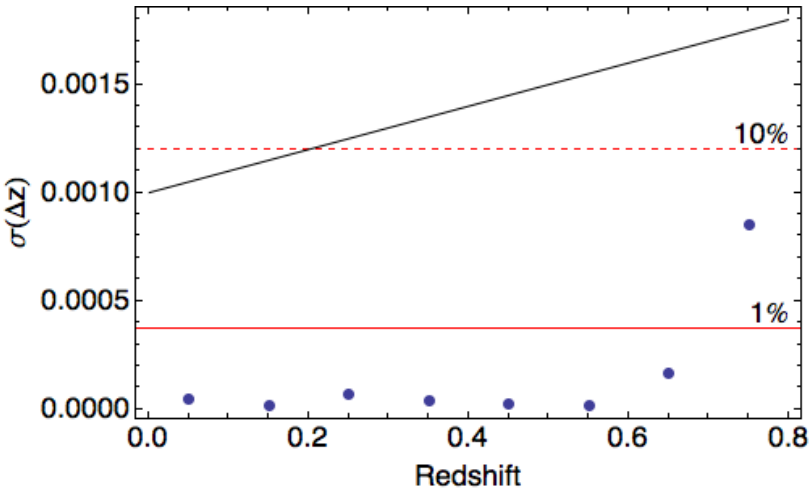

(a) Uncertainty in the photo-z bias

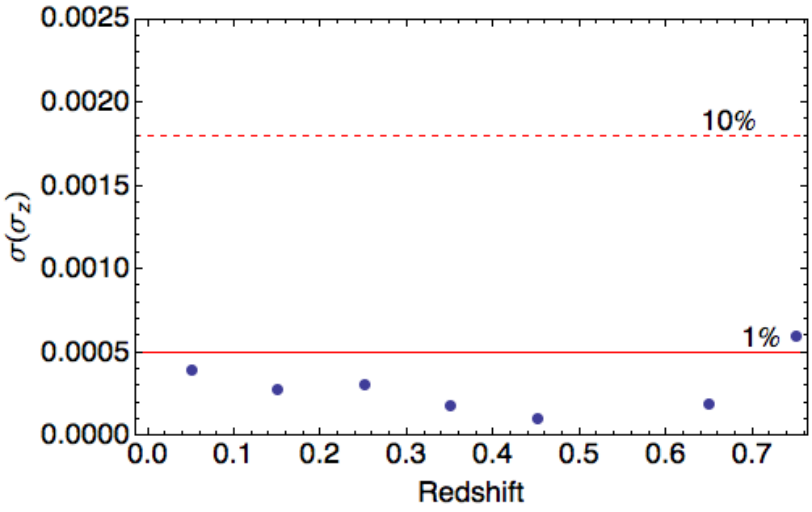

(b) Uncertainty in the photo-z error

Figure 4.11: The uncertainty in ArborZ's photo-z estimates in the SDSS validation test. The points indicate the uncertainties. The solid red line marks the 1% degradation in the dark energy figure of merit. The dashed red line marks the 10% degradation. In the upper plot, the DES science requirement is the solid black line; in the lower plot, the DES science requirement is above the range of the plot.

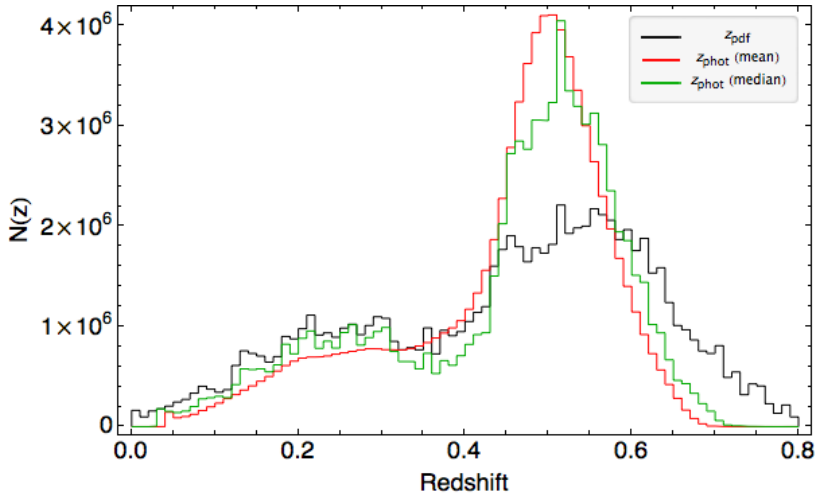

Figure 4.12: The estimated redshift distribution $N(z)$ in the SDSS DR10 imaging catalog.

With the promising results of the validation and unrepresentation tests in hand, we now proceed to produce a photo-z catalog for the SDSS DR10 catalog of Section 4.1.1. We use the same parameters as in the validation test. The redshift distribution $N(z)$ is shown in Figure 4.12. In Figures 4.13a and 4.13b we show, respectively, the bias versus $z_{\rm phot}$ and the error versus $z_{\rm phot}$. These plots are generated by binning $p(z)$ in bins of $z_{\rm phot}$ and then plotting the probability contours (along with means and widths) of the summed $p(z)$s in each bin.

The mean bias is exceptionally low across the entire redshift range, despite the increased uncertainty near $z \approx 0.52$. In the reconstructed redshift distribution, we see a probable overshoot in the $z_{\rm phot}$ distribution, similar to the validation sample. We expect, therefore, that the $z_{\rm pdf}$ distribution is a more accurate representation of the underlying redshift distribution. This photo-z catalog is an important tool for tackling cosmology using data from the largest sky survey ever undertaken.

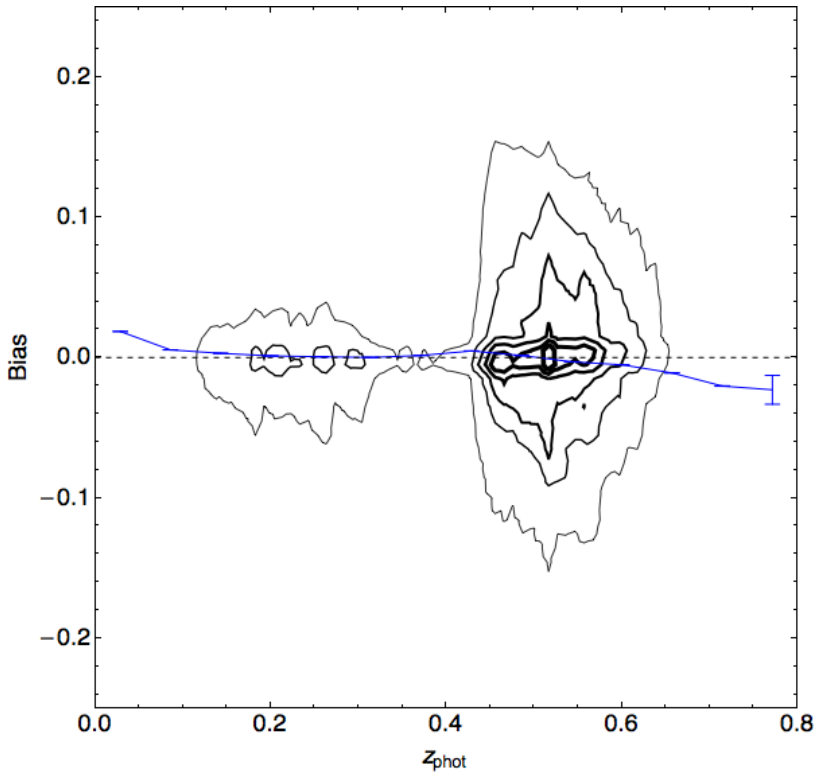

(a) Photo-z bias

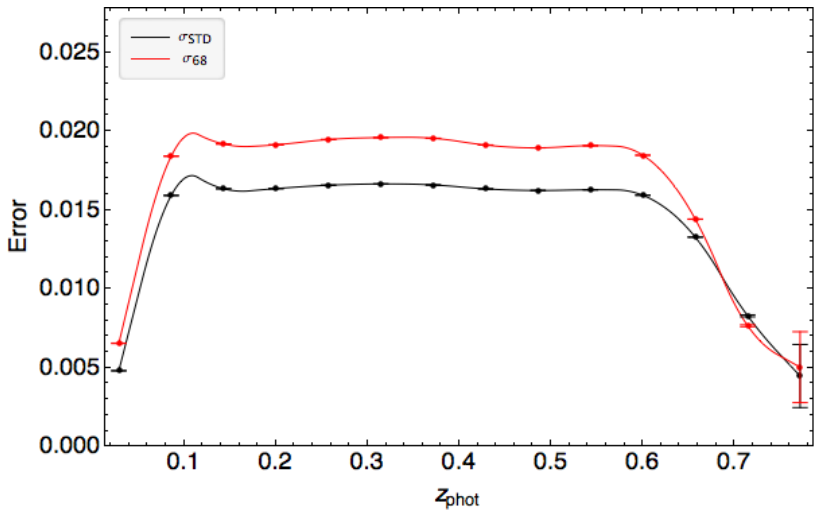

(b) Photo-z error

Figure 4.13: Bias and error plots for the SDSS DR10 catalog. These are generated by summing $p(z)$ distributions in bins of z_{phot}. The resulting contours and means are shown in Figure 4.13a, and the widths are shown in Figure 4.13b.

4.2 The Dark Energy Survey

The Dark Energy Survey (DES; T. Abbott et al. 2005) is a large-scale galaxy survey which will collect optical imaging data in $grizY$ bands for approximately 3×10^8 galaxies spread over $5\,000\,\mathrm{deg}^2$ of the southern sky. These data will be collected over the course of 525 nights over five years, beginning officially on August 31, 2013. The DES's instrument of choice is the Dark Energy Camera (DECam; DePoy et al. 2008). Installed on the 4m Blanco telescope at CTIO and sporting a 570 megapixel camera with a 2.2 deg field-of-view, DECam is the largest, most powerful instrument of its type ever built[2]. Once finished, the DES will be the largest imaging survey ever completed, surpassing even the SDSS.

As part of its commissioning and science verification, a great deal of science data was taken using DECam during the year preceding the start of the DES. This data was processed through the official DES data management (DESDM; Mohr et al. 2012) pipeline, and a coadded catalog—called SV-A1—was constructed. This data is the perfect playground for photometric redshift estimation, and although not much cosmology is expected from the SV-A1 data release, it will help pave the way for many DES science working groups to constrain cosmology over the next five years.

4.2.1 The SV-A1 Imaging Catalog

We begin processing the SV-A1 catalog by selecting all objects in the co-added catalog that were not flagged as bad in i-band and which had SPREAD_MODEL > 0.002 (to exclude stars). Extinction corrections were applied on an object-by-object basis using the D. J. Schlegel, D. P. Finkbeiner, and M. Davis 1998 dust maps with the Schlafly and D. P. Finkbeiner 2011 recalibration. We then apply the photometric calibrations of H. Lin 2013. These are generated on a tile-by-tile basis by first selecting galaxies

[2]See Appendix A for a more detailed hardware description.

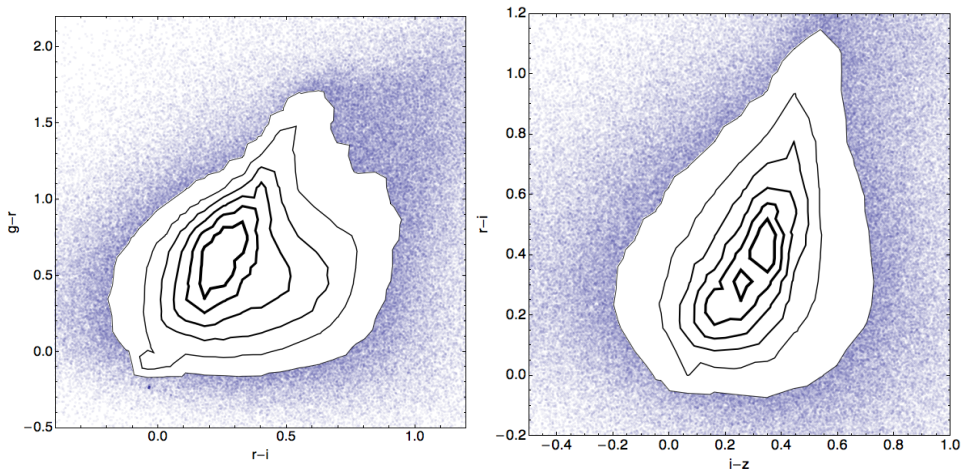

(a) $g - r$ vs $r - i$ (b) $r-$ vs $i - z$

Figure 4.14: Color-color plots showing the distribution of colors in the DES SV-A1 catalog.

with $18 < r < 22$. The r-band "MAG_DETMODEL" zero-point is then shifted to match the median r-band MAG_DETMODEL − MAG_AUTO value in the VVDS Deep 02hr field (Le Fèvre, Vettolani, et al. 2005). Finally, the $gizY$-band zero-points are shifted such that the $g - r$, $r - i$, $z - r$, and $Y - r$ colors match their respective medians in the VVDS Deep 02hr field. If applying this procedure to a tile results in a shift of more than 0.2 magnitudes in any band, the tile is discarded.

After applying these cuts, we keep 17 097 383 galaxies. We show the $g - r$ versus $r - i$ and $r - i$ versus $i - z$ color-color plots in Figure 4.14, and the magnitude distributions are shown in Figure 4.16.

4.2.2 DES Spectroscopic Training Set

Having selected a sample from the SV-A1 catalog, we can follow the same prescription in Section 4.1.1 as we did for SDSS to match redshifts from a variety of spectroscopic surveys to DES imaging. The redshifts sources used are listed below.

- **2dFGRS**: 2 120 redshifts from the 2dF Galaxy Redshift Survey (Colless et al. 2001).

- **ACES**: 3 117 redshifts from the Arizona CDFS Environment Survey (Cooper, Yan, et al. 2012).
- **GOODS**: 854 redshifts from the Great Observatories Origins Deeps Survey (Balestra et al. 2010).
- **K20**: 80 redshifts from the K20 survey (Mignoli et al. 2005).
- **OzDES**: 14 218 redshifts obtained on the Anglo-Australian Telescope as part of the OzDES program, the primary spectroscopic followup program for the Dark Energy Survey (DES).
- **PRIMUS**: 47 997 redshifts from the Prism Multi-Object Survey (Coil et al. 2011).
- **SDSS**: 1 100 redshifts from the Sloan Digital Sky Survey (Dawson et al. 2013).
- **VIPERS**: 6 423 redshifts from the VIMOS Public Extragalactic Redshift Survey (Guzzo et al. 2013).
- **VVDS**: 1 822 redshifts from the VIRMOS-VLT Deep Survey (Le Fèvre, Vettolani, et al. 2005).
- **zCOSMOS**: 3 692 redshifts from from the Cosmological Evolution Survey (COSMOS) field (Lilly, Le Fèvre, Renzini, et al. 2007).

In total, the DES spectroscopic training set contains 81 423 redshifts matched to SV-A1 imaging. The redshift range is well-represented out to $z \approx 1.2$ with a tail extending out to $z \approx 1.5$. This distribution is shown in Figure 4.15, and the magnitude distributions are shown in Figure 4.16.

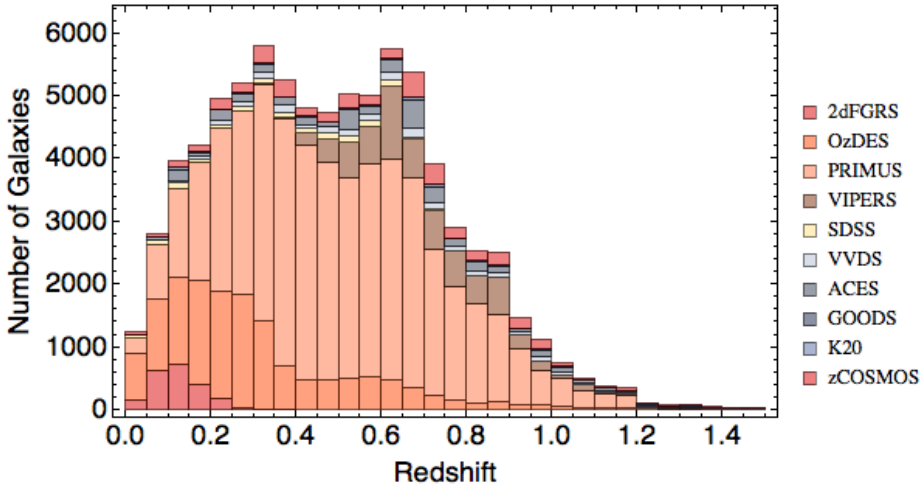

(a) Overall redshift distribution (stacked)

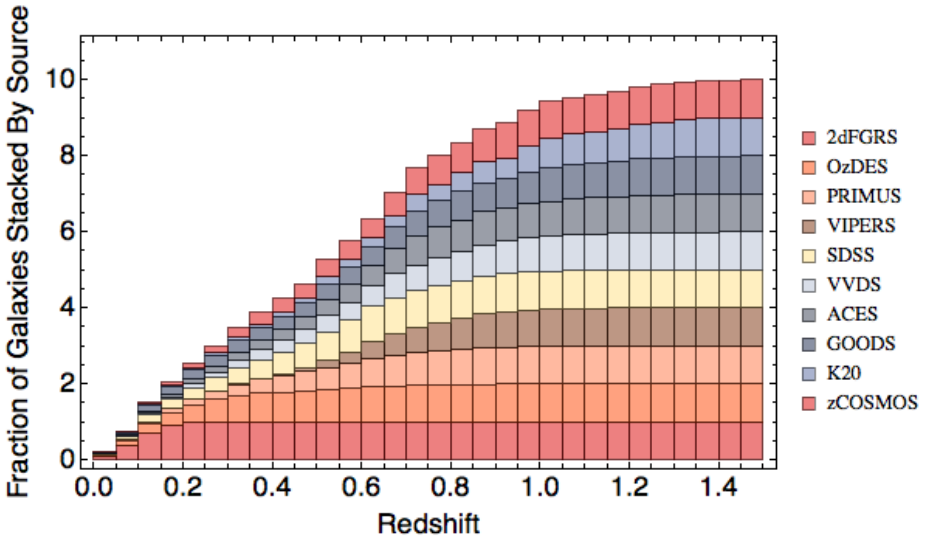

(b) Cumulative distribution functions (stacked)

Figure 4.15: The redshift distribution of the DES spectroscopic training set. In Figure 4.15a is the overall redshift distribution, including contributions from all ten source catalogs. In Figure 4.15b, we stack the cumulative distribution functions (CDFs) of the redshift distributions from each catalog. In both plots, 2dFGRS is shown on the bottom of the stack, and zCOSMOS is shown on top.

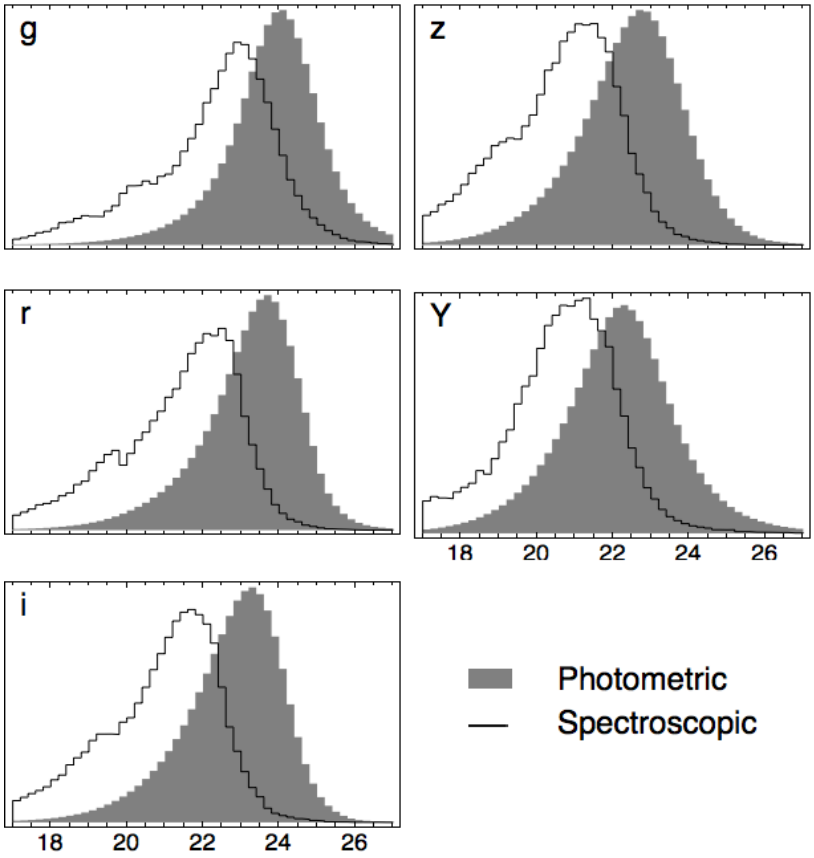

Figure 4.16: The distributions of $m_{\rm auto}$ magnitudes in the DES SV-A1 photometry and the spectroscopic training set. Both distributions are normalized to more easily compare contributions at each magnitude.

4.2.3 DES Photometric Redshift Catalog

In a similar spirit to Section 4.1.4, we first run a validation test on the DES data—training on a subset of the spectroscopic catalog and evaluating on the remainder—to test our training parameters and ensure that the algorithm is behaving appropriately. We therefore train ArborZ on one-third of the spectroscopic catalog, 26 869 galaxies in all, using all five magnitudes as training variables. We use 80 equal-width bins out to redshift 1.2, with a photo-z resolution of 0.02.

Figure 4.17 shows the reconstructed redshift distribution $N(z)$. The $p(z)$ reconstruction is particularly good in comparison to the other z_{phot} measures, which exhibit noticeable bias which manifests as a compressed distribution. The χ^2 goodness-of-fit metric of Equation 3.24 gives 24.9 for z_{phot}, 14.4 for $z_{\text{phot}}^{\text{med}}$, and 1.7 for z_{pdf}.

Figure 4.18a shows the photo-z bias as a function of redshift, and the photo-z error versus redshift is shown in Figure 4.18b. Due to both the poorer data quality in the SV-A1 catalog (photometric calibration, source finding, and deblending are still under heavy development) and the less representative training set, we generally see larger bias and error than in the SDSS catalog. That being said, ArborZ $p(z)$ estimates still work well, and a well-characterized bias can be calibrated out of cosmological measurements.

Now we turn to the final SV-A1 catalog. We use the same parameters as in the validation test. The redshift distribution $N(z)$ is shown in Figure 4.19. In Figures 4.20a and 4.20b we show, respectively, the bias versus z_{phot} and the error versus z_{phot}, generated in the same way as Figure 4.13.

We see good performance and low bias across the entire redshift range, especially considering the relative immaturity of the SV-A1 catalog. In Figure 4.21 we compare the photo-z uncertainties to the DES science requirements. Except for the high redshift end, the uncertainty in the photo-z error is within the limits of the science

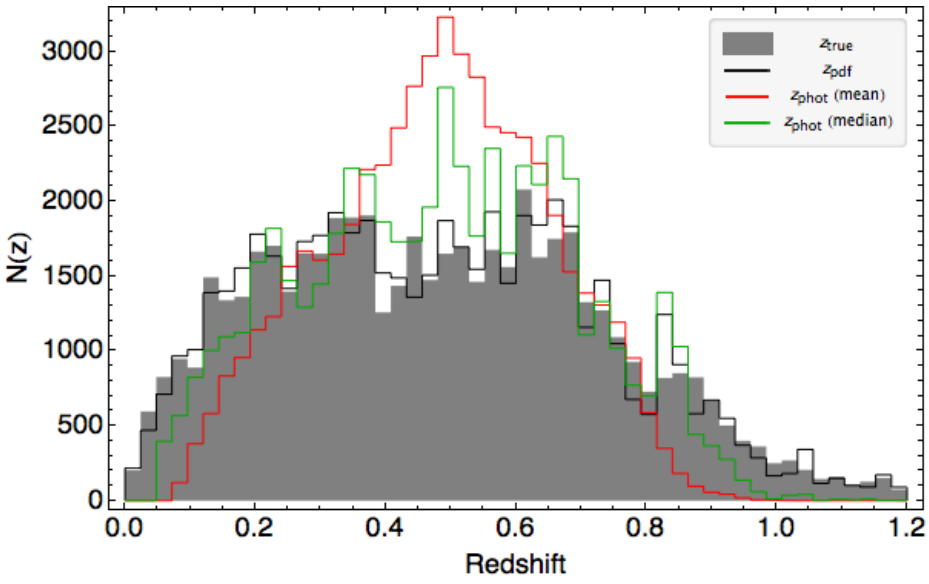

(a) Reconstructed redshift distribution

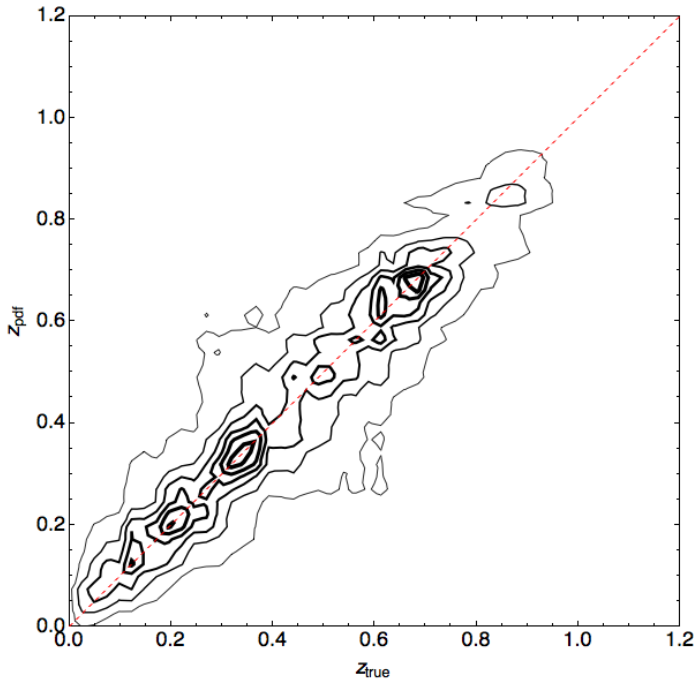

(b) $p(z)$ vs. $z_{\rm true}$ contours

Figure 4.17: The photo-z results in the DES validation sample, calculated by training on a random one-third of galaxies in the overall spectroscopic training set, and evaluating on the remainder. Good agreement between ArborZ $p(z)$ and the true redshift is seen across the entire redshift range.

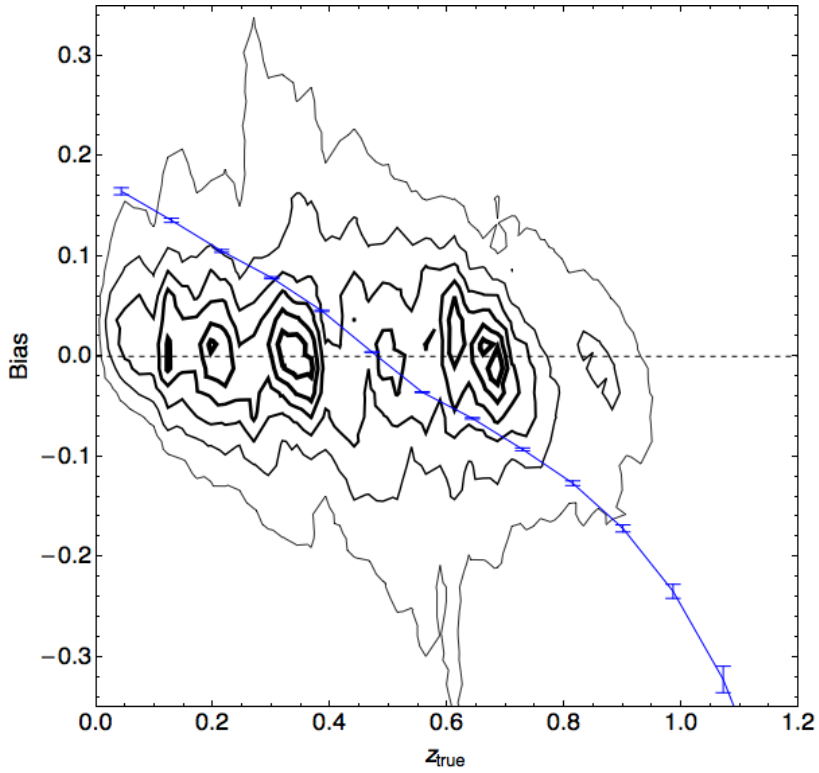

(a) Photo-z bias

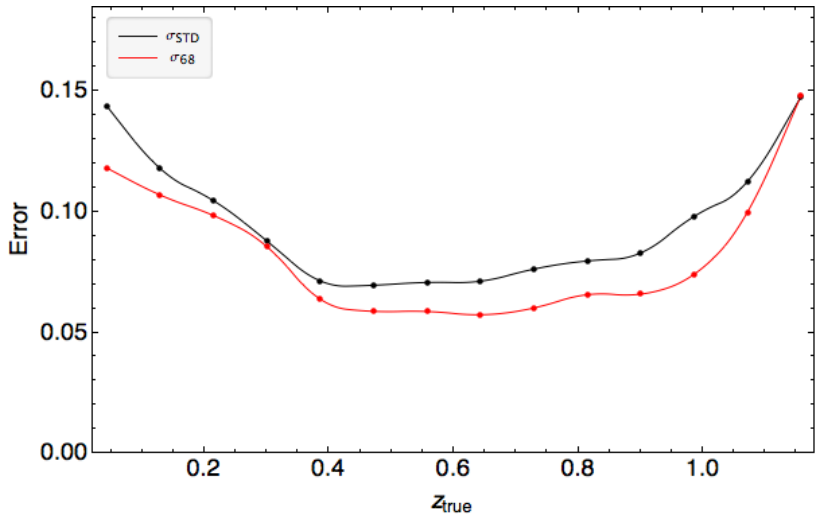

(b) Photo-z error

Figure 4.18: Bias and error plots for the DES validation sample, indicating good photo-z results across the entire redshift range.

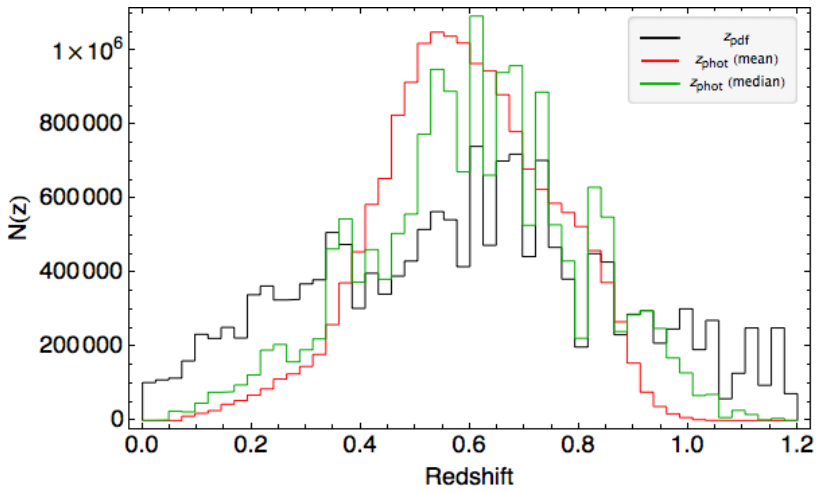

Figure 4.19: The estimated redshift distribution $N(z)$ in the DES SV-A1 imaging catalog.

requirement. Within $0.2 < z < 0.8$ the bias limits have met the requirements, but at low and high redshifts the photo-z bias uncertainty still has room for improvement. That being said, for $z < 1.0$ the photo-z uncertainties are within the 50% degradation limit. This gives us confidence that the ArborZ photo-z estimates are ready to be used for preliminary science. As the DES progresses and as further spectroscopic followup occurs, we expect these uncertainties to shrink. With such a large catalog available, many tests of cosmology can be performed, and the DES can begin to reveal some of the universe's secrets.

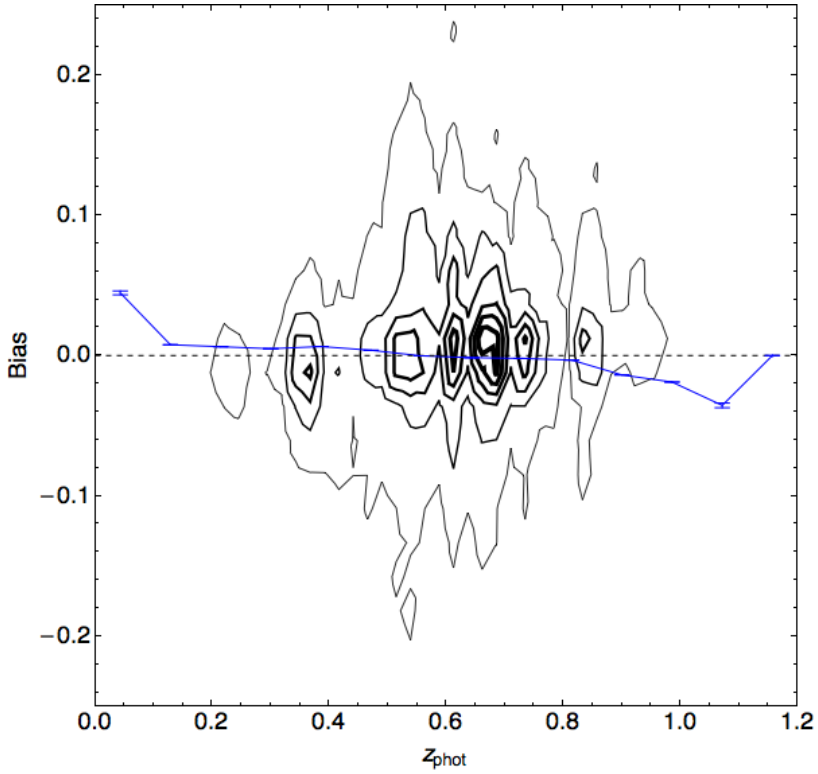

(a) Photo-z bias

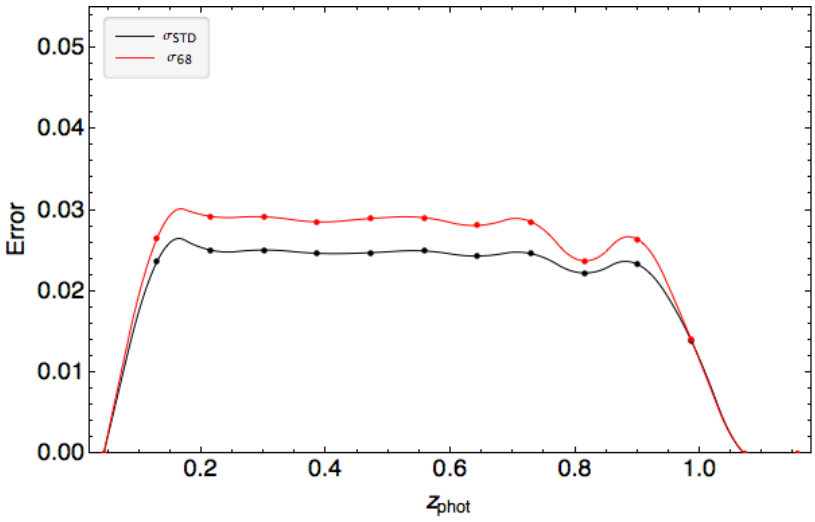

(b) Photo-z error

Figure 4.20: Bias and error plots for the DES SV-A1 imaging catalog.

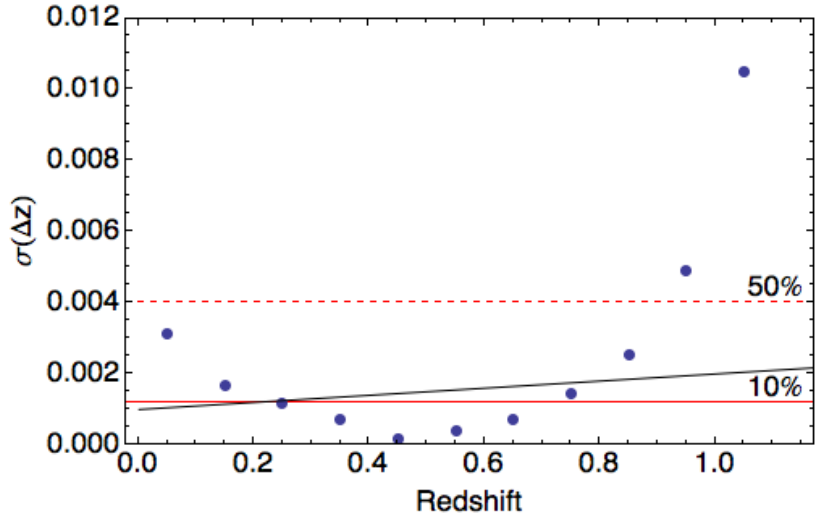

(a) Uncertainty in the photo-z bias

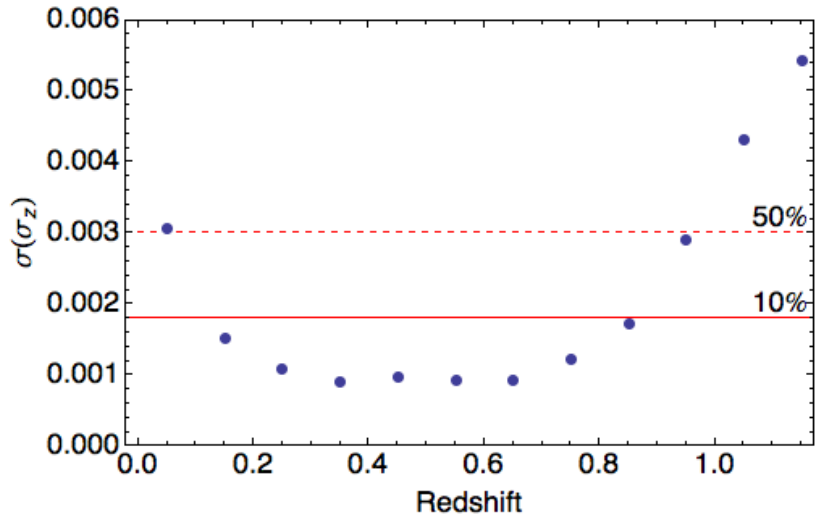

(b) Uncertainty in the photo-z error

Figure 4.21: The uncertainty in ArborZ's photo-z estimates in the DES validation test. The points indicate the uncertainties. The solid red line marks the 1% degradation in the dark energy figure of merit. The dashed red line marks the 10% degradation. In the upper plot, the DES science requirement is the solid black line; in the lower plot, the DES science requirement is above the range of the plot.

Chapter V

Constraining Cluster Redshifts

So far, we have shown that ArborZ provides an excellent methodology for estimating galaxy redshifts, and we have argued that such estimates are necessary for doing precision cosmology. In this chapter, we will address a particular example of ArborZ's application to cluster science, namely, using galaxy $p(z)$s to obtain cluster redshifts.

5.1 Cluster Cosmology

Clusters are an important probe of cosmology, as they are the largest gravitationally bound structures in the universe. Quantum fluctuations seeded the initial density perturbations of the universe. Gravitational instabilities caused these overdense regions to accrue additional dark matter, forming bound structures called dark matter halos. Galaxy clusters form in these halos, and so clusters are tracers of the dark matter distribution of the universe. The mass function $dN(M,z)/dM$ (Jenkins et al. 2001; Press and Schechter 1974; Tinker et al. 2008) of these halos is exponentially sensitive to cosmological parameters (Evrard 1989; Frenk et al. 1990), making them excellent probes of cosmology (Mana et al. 2013). Changes in cosmological parameters affect not only the observed mass function (due to a different distance-redshift relation), but also the predicted model (due to different predictions for the growth of struc-

ture). An example of this is shown in Figure 5.1, where the mass function is used to determine a best-fit cosmology using data from *Chandra* (Vikhlinin, Kravtsov, et al. 2009). The chief challenge associated with cluster cosmology is that cluster masses is not directly observable. Instead, proxy measurements must be used which have a known relationship with the cluster's mass (Song et al. 2012; Vikhlinin, Burenin, et al. 2009). Calibrating these proxies and controlling systematics is itself a difficult task.

Regardless of how cluster mass is estimated, however, one must identify clusters in the first place. How this is done depends strongly on the observed wavelength. In the microwave, CMB photons can be inverse-Compton scattered as they pass through the hot cluster gas, an effect called the *Sunyaev-Zeldovich* (SZ) effect (R. A. Sunyaev and Y. B. Zeldovich 1972; Y. B. Zeldovich and R. A. Sunyaev 1969). The hot cluster gas also emits high-energy photons which can be observed in the X-ray (Gursky et al. 1971; Kellogg et al. 1971). In the optical, there are several techniques which have emerged to address this task of cluster identification. Some of these techniques search for local overdensities in color space (Miller et al. 2005). Others (Hao et al. 2010; Koester 2006; Rykoff et al. 2013) take advantage of the cluster red-sequence, also called the E/S0 ridgeline. The cluster red-sequence describes a low-scatter, linear correlation in color-magnitude space associated with the bright, passively-evolving elliptical galaxies which populate cluster cores (Bower, Lucey, and Ellis 1992; M. D. Gladders and H. K. C. Yee 2000; López-Cruz, Barkhouse, and H. K. Yee 2004; Sandage and Visvanathan 1978). This relationship is remarkably homogeneous across galaxy clusters and has been observed for clusters at redshift $z > 1$ (Hilton et al. 2009; Papovich et al. 2010). To measure the red-sequence, candidate cluster members[1] are plotted in a color-magnitude space, such as $B -$

[1]In practice, this is often done by projecting all galaxies onto the celestial sphere and measuring the color-magnitude relationship for galaxies near the cluster's center. The projection effect is not particularly important, since foreground/background galaxies do not form a coherent red-sequence (M. D. Gladders and H. K. C. Yee 2000).

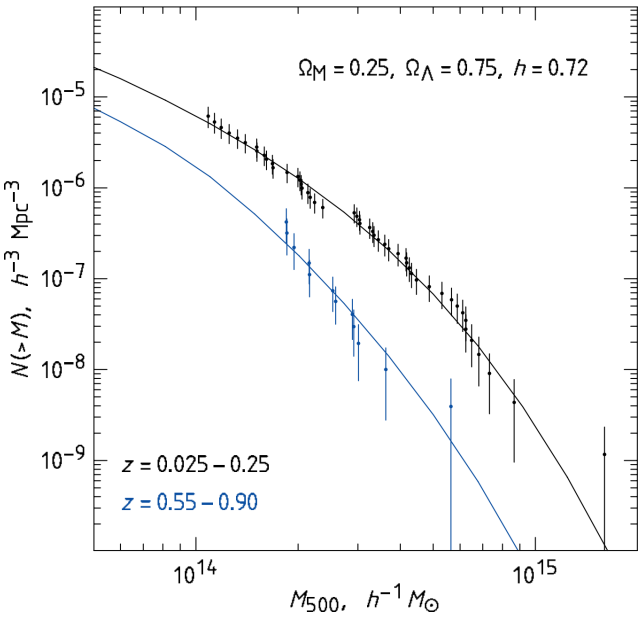

(a) Fiducial model

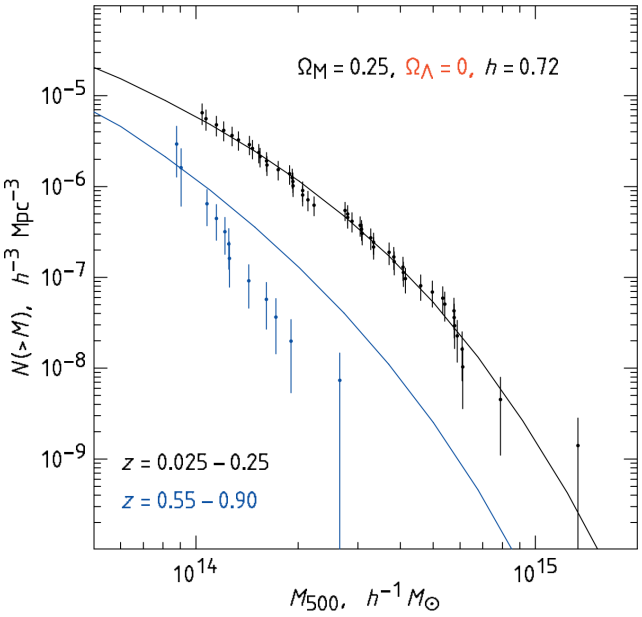

(b) No dark energy

Figure 5.1: Data from *Chandra* is used to constrain cosmological parameters using the cluster mass function. On the top we see good agreement between the observed mass function (points) and the predicted model (lines), indicating that the chosen cosmological parameters (inset at top of plot) are a good fit to reality. On the bottom, the dark energy contribution to the mass-energy of the universe is set to $\Omega_\Lambda = 0$. The resulting disagreement between the observed and predicted mass functions suggests that we do not live in a universe without dark energy. Figures taken from Vikhlinin, Kravtsov, et al. 2009.

R versus R, as shown in Figure 5.2, where the red-sequence becomes manifestly apparent. The problem of measuring the slope of the red sequence is then reduced to a relatively straightforward fitting problem (M. D. Gladders et al. 1998), although more sophisticated weighting schemes can be used (M. D. Gladders and H. K. C. Yee 2000). By measuring the slope of the red-sequence at different redshifts, redshift estimates can be determined (Stott et al. 2009). Thus, the red-sequence can be used not only for optical cluster detection, but also to estimate the photometric redshifts of these clusters.

Cluster photo-zs estimated from the red-sequence typically have a scatter $\sigma_z \lesssim 0.02$ (Koester 2006; Rykoff et al. 2013), but there are other ways to estimate redshift. For example, if a cluster-finding algorithm is capable of producing a list of probable cluster members, then one could average the individual $z_{\rm phot}$ estimates for each cluster member. Since the error in the photo-z for the cluster would decrease with the number $N_{\rm member}$ of members in the cluster, one could expect reasonably accurate redshift estimates even for small clusters with $N_{\rm members} \gtrsim 5$ (H. Lin et al. 2006). Since ArborZ is capable of producing $p(z)$ estimates for each cluster member, we are interested in investigating the possibility of using members' $p(z)$ distributions to construct a single cluster $p_{\rm cluster}(z)$. This would presumably contain more information than a single, best-estimate photo-z for the cluster, analogous to what we have observed for galaxies in the previous chapter. Therefore, we now turn our attention to estimating a cluster's $p_{\rm cluster}(z)$ distribution using the $p(z)$s of its members.

5.2 Cluster $p(z)$

5.2.1 Stacking Methods

There are several ways we can imagine combining ("stacking") cluster members' $p(z)$s. The most obvious method would be to simply sum the member $p(z)$s:

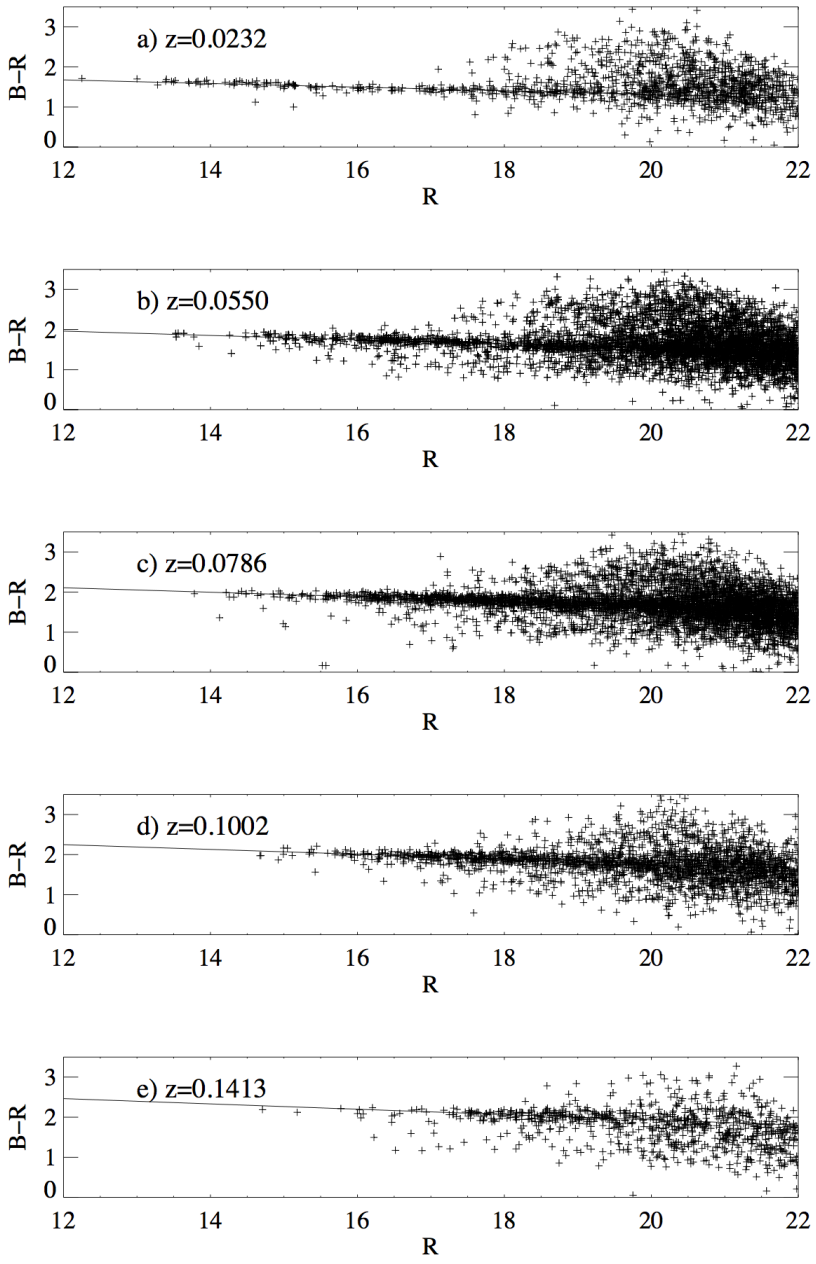

Figure 5.2: Six examples of the cluster red-sequence found in Abell clusters. The color-magnitude relationship works well to find the red-sequence (marked by the lines), since it makes bright, red cluster members obvious. Figure taken from M. D. Gladders et al. 1998.

$$p_{\rm cluster}(z_i) = \sum_{j=1}^{N} p_j(z_i) \tag{5.1}$$

where the sum runs over the N members of the cluster of interest. This is the same as to calculating the redshift distribution $N(z)$ for the cluster members. The best-estimate cluster photo-z could then be defined as the mean or median of this $p_{\rm cluster}(z)$ distribution. This can work moderately well; however, unlike the redshift distributions of entire samples—composed of thousands of $p(z)$s—the stacked cluster $p(z)$ contains relatively few contributions, and so it is susceptible to noise from individual galaxy $p(z)$s. At zeroth order, this acts to flatten the cluster $p(z)$, driving the median redshift toward the center of the redshift range.

Another problem with this stacking method is that it ignores our prior knowledge that all of these galaxies ought to lie at the same redshift. This suggests another approach: the joint redshift distribution. Since all members of a cluster lie at the same redshift, the probability that the cluster's redshift is z_0 is the probability that each member is at $z = z_0$ also. In other words, we consider the product of the member $p(z)$s:

$$p_{\rm cluster}(z_i) = \prod_{j=1}^{N} p_j(z_i) \tag{5.2}$$

If each cluster member has a well-estimated $p(z)$, this works well, and generally results in a single, narrow peak at the cluster's true redshift. We show an example of this stacking method in Figure 5.3 using data from DES mocks (which we will discuss presently).

5.2.2 Failure Modes

Although Equation 5.2 work well as a definition for the cluster $p_{\rm cluster}(z)$, it is not uncommon to find a situation where at least one member's $p(z)$ has no support

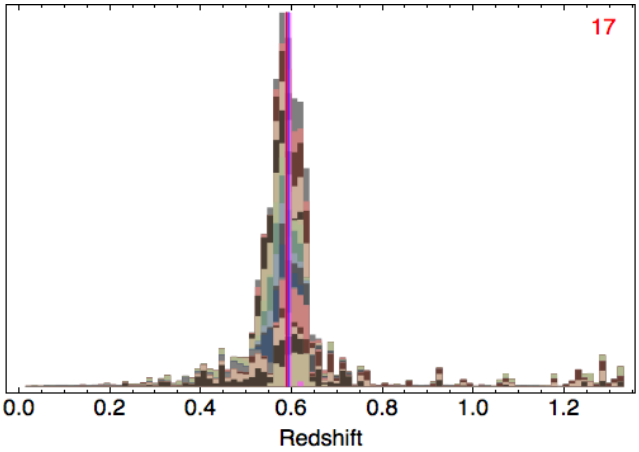

(a)

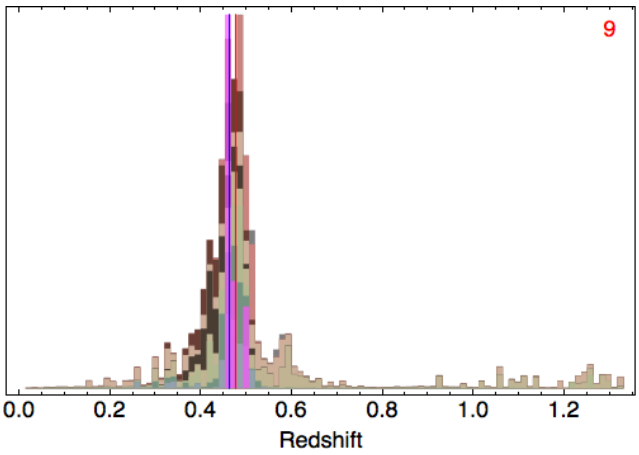

(b)

Figure 5.3: Examples of cluster $p_{\rm cluster}(z)$s constructed using the joint probability distribution of Equation 5.2. Data is taken from the DES mock catalog. The member galaxies' $p(z)$s are shown as a stacked histogram and the cluster $p_{\rm cluster}(z)$ distribution is overlaid in magenta. The true cluster redshift is indicated with a red line, and the best-estimate photo-z—defined as the median of $p_{\rm cluster}(z)$—is indicated with a blue line. The red number inset in the upper-right of each plot is the number of members stacked in the cluster.

near the cluster's true redshift[2]. In these cases, that single "rogue" $p(z)$ will cause the cluster $p_{\rm cluster}(z)$ to have zero area (or a statistically meaningless distribution, in the event that one attempts to renormalize after being dominated by numerical rounding error). When this occurs—which can be detected by checking if the resulting distribution is compatible with a flat, zero-area distribution—one must discard this cluster entirely or find an alternative way to include the rogue member (e.g., apply a weighting scheme or use a different stacking method).

One can imagine several methods for handling this failure mode. We briefly discuss a simple method here which constructs an estimator $\hat{p}_{\rm cluster}(z)$ for the cluster's redshift rather than a true probability distribution. This estimator avoids the failure mode by treating the summed probability as a log-likelihood:

$$\ln \hat{p}_{\rm cluster}(z_i) = \sum_{j=1}^{N} p_j(z_i) \tag{5.3}$$

Because we are using the sum, we avoid the problem of a single galaxy $p(z)$ from destroying the information content of the cluster redshift estimator. The logarithm will also retain the position of the most probable redshift for the stacked $p(z)$; moreover, the logarithm will give higher weight to more probable redshifts, helping to reduce the noise from the sum. The choice of the logarithm is motivated only by the analogy to log-likelihood analysis, and similar functions could be easily be used to produce different estimators. The last remaining problem is that the minimum value of our current estimator in any given bin is now 1. Rearranging the equation and subtracting out the baseline value of 1, our final estimator becomes:

$$\hat{p}_{\rm cluster}(z_i) = \exp\left[\sum_{j=1}^{N} p_j(z_i)\right] - 1 = \left[\prod_{j=1}^{N} e^{p_j(z_i)}\right] - 1 \tag{5.4}$$

[2]Depending on the photo-z algorithm and its implementation, the value $p(z_i)$ of a particular bin may either be exactly or approximately zero. By approximately zero we mean that the value is so small as to be statistically indistinguishable from zero (particularly after accounting for any numerical rounding errors on a computer).

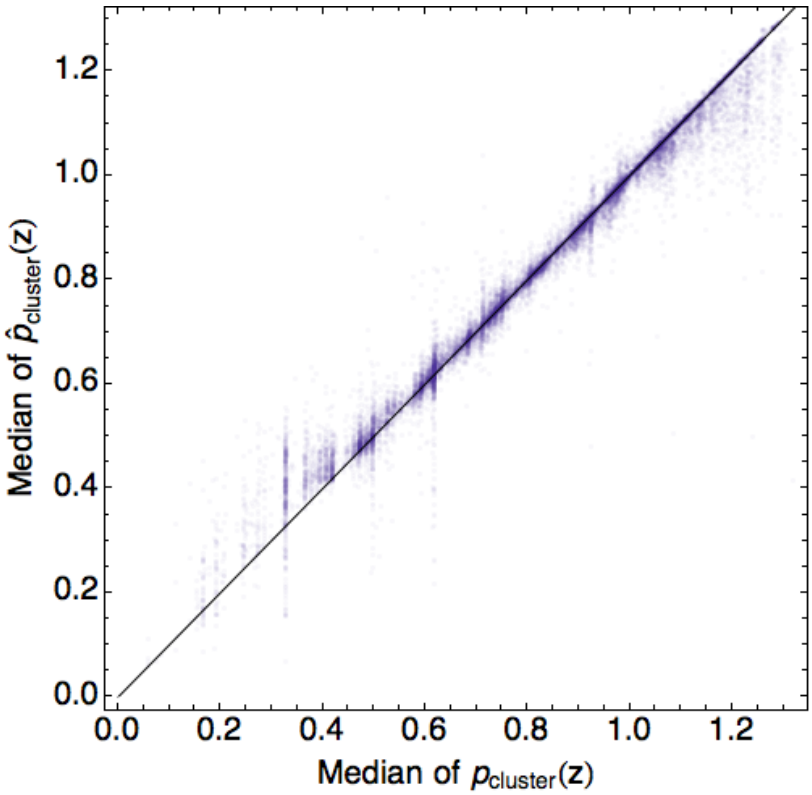

Figure 5.4: Relationship between the best-estimate photo-z obtained from the $p_{\rm cluster}(z)$ distribution and the median of the cluster photo-z estimator $\hat{p}_{\rm cluster}(z)$. Data is taken from the DES mocks.

The median of this distribution then provides an estimate of the cluster's redshift, which is usually very similar the median of Equation 5.2, as we see in Figure 5.4. The relationship is not perfect, but it provides a simple way to produce a meaningful cluster photo-z without introducing more sophisticated weighting schemes.

5.3 Performance in Simulated Catalogs

Before we apply the stacking method of Equation 5.2 on real clusters, we first test it on simulated galaxy catalogs. We turn again to the DES mocks discussed in Section 3.4. We can identify clusters using the simulation's halo catalog, which stores the halo properties which emerged from the N-body simulation. In order to make this test more observationally plausible, we place a series of cuts on the halos and halo

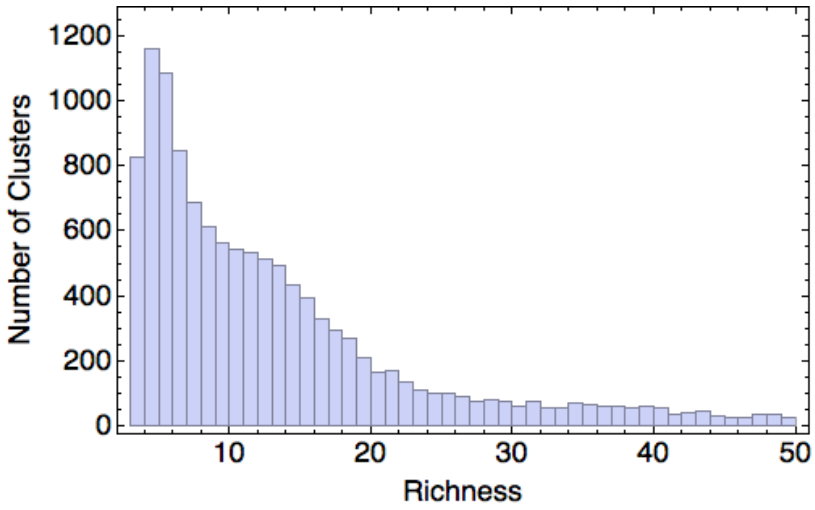

Figure 5.5: Cluster richness in the DES mocks after our member cuts are applied.

members. First, we only keep halos with mass $M > 1.0 \cdot 10^{13} M_{\odot}$. We then select red-sequence member galaxies that can identified using the $g-r$ versus r color-magnitude relationship, discarding members with $i > 23$. After applying these cuts, if a halo has fewer than three members associated with it, we discard it. We are left with 12 507 halos containing a total of 102 697 member galaxies; the richness distribution (number of members in a cluster) is shown in Figure 5.5.

We train ArborZ on the $grizY$ magnitudes of 100 000 randomly chosen galaxies from a tile of the DES mocks, using 100 evenly-spaced bins out to $z = 1.33$, with $\sigma_{\rm res} = 0.02$. We then evaluate ArborZ on the cluster members we identified. We stack each cluster's members' $p(z)$s using Equation 5.2 to produce $p_{\rm cluster}(z)$. The cluster's best-estimate photo-z $z_{\rm phot}^{\rm cluster}$ is defined to be the median of this distribution, and the cluster's error $\sigma_z^{\rm cluster}$ is defined as half of the width of the central 68% of the area of distribution. In Figure 5.6a we plot $p_{\rm cluster}(z)$ as a function of cluster redshift, and in Figure 5.6b we plot the cluster photo-z versus $z_{\rm true}$. We see that $p_{\rm cluster}(z)$s perform similarly to individual galaxy $p(z)$s in that keeping the full distribution preserves a more accurate measure of the cluster's redshift.

In Figure 5.7 we show the reconstructed redshift distribution of the clusters. Be-

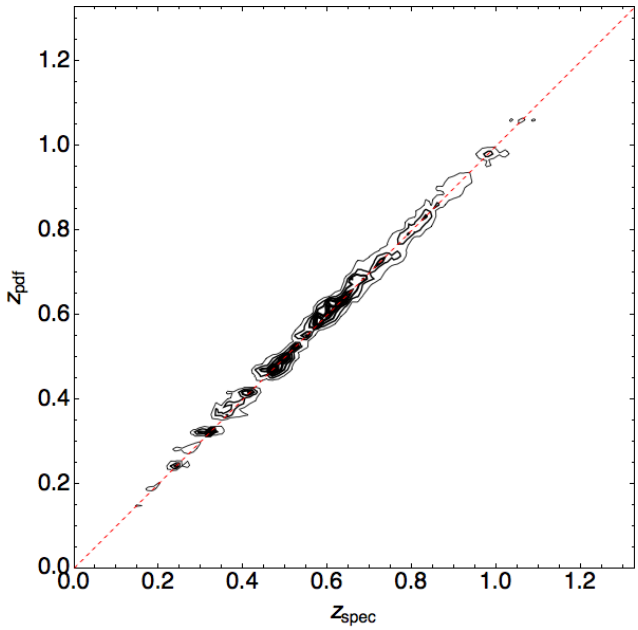

(a) Cluster $p_{\text{cluster}}(z)$ versus redshift

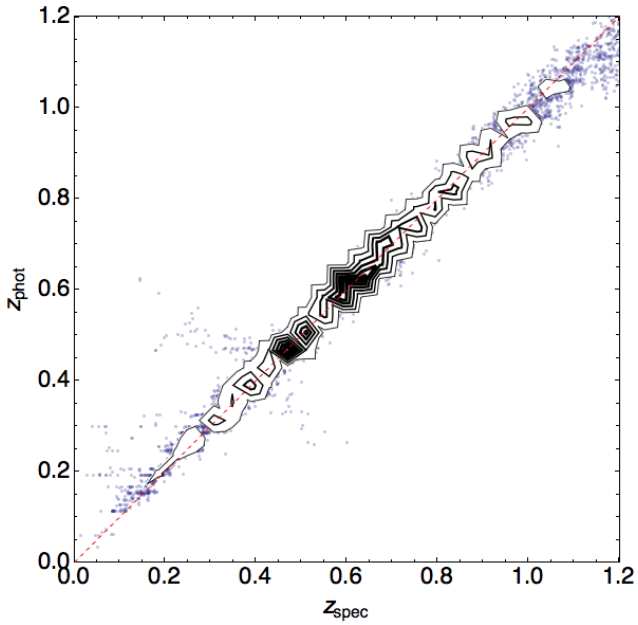

(b) Cluster $z_{\text{phot}}^{\text{cluster}}$ versus redshift

Figure 5.6: Cluster photo-z versus true redshift, using both the full $p_{\text{cluster}}(z)$ distribution as well as its median, $z_{\text{phot}}^{\text{cluster}}$. Using the full probability distribution produces a less biased fit to the line $z_{\text{phot}} = z_{\text{true}}$.

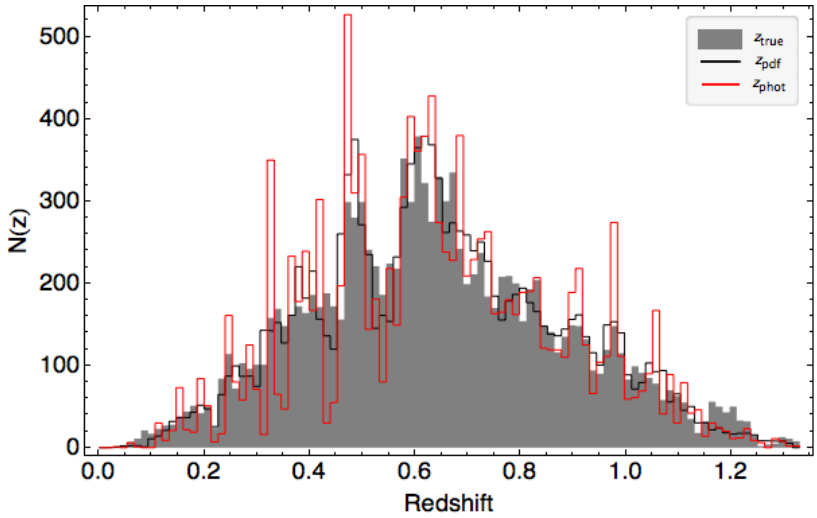

Figure 5.7: Reconstructed redshift distributions $N(z)$ for the mock cluster catalog.

cause the tendency of Equation 5.2 is to drive the cluster's $p_{\rm cluster}(z)$ toward a delta function, we smear the $N(z)$ measurement calculated using the probability distribution in order to smooth out high-frequency noise; in this case, we use a top-hat distribution with width equal to three times the bin width. Without this convolution, the stacked $p_{\rm cluster}(z)$ distribution looks almost identical to the $z_{\rm phot}^{\rm cluster}$ histogram. In Figures 5.8a and 5.8b we show the cluster photo-z bias and error, respectively, as a function of redshift. We see low bias and error across the entire redshift range, lower than we typically see in the photo-zs of individual galaxies (c.f. Section 3.4), which we expect from stacking the statistical signal of multiple galaxies in the cluster. In the range $0.3 < z < 1.1$ we obtain bias and error values which approach the rough constraint $\Delta z \approx \sigma_z \lesssim 0.02$ for using cluster photo-zs to constrain cosmology (D. Huterer, Kim, et al. 2004; Ma, W. Hu, and D. Huterer 2006). With this in mind, we optimistically continue our study of stacking $p(z)$ in real data and compare it to traditional red-sequence redshift estimates.

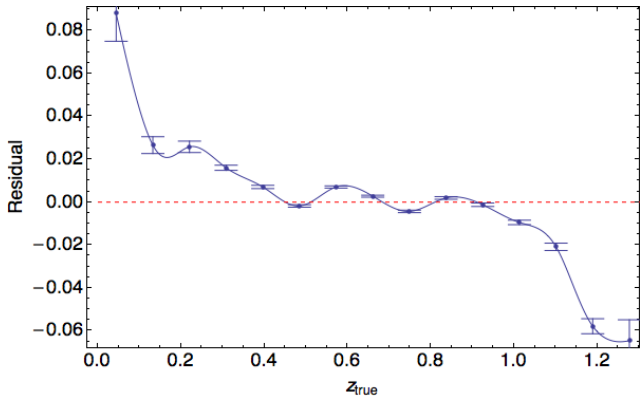

(a) Cluster photo-z bias versus redshift

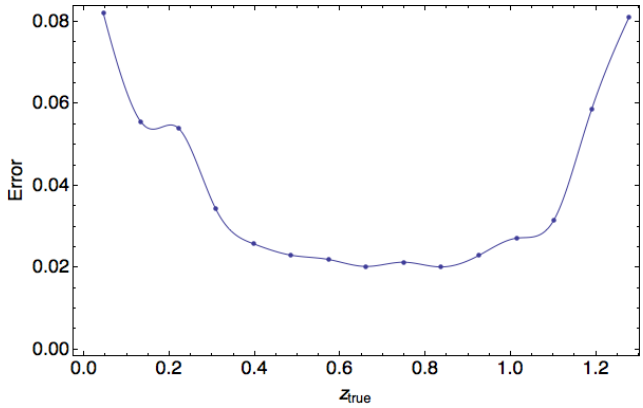

(b) Cluster photo-z error versus redshift

Figure 5.8: Cluster photo-z error and bias in the DES mock galaxy catalogs.

5.4 The DES-XCS Cluster Catalog

As we mentioned earlier, clusters can be identified by their X-ray emissions using space-based[3] X-ray telescopes. The hot gas in clusters emits X-rays, and tracing X-ray contours can trace the underlying baryonic matter distribution (Byram, Chubb, and H. Friedman 1966; Cavaliere, Gursky, and Tucker 1971). The *XMM* Cluster Survey (XCS; Mehrtens et al. 2012) uses archival data in the *XMM-Newton* science archive to search for serendipitous X-ray galaxy clusters. Where possible, redshifts were assigned to each cluster from a compilation (Miller et al. in prep) of previous literature references and follow-up spectroscopic observations; cluster members were identified using the red-sequence.

In order to apply our cluster photo-z technique to the XCS clusters, we first need to assign imaging observables to the cluster members. We therefore match the XCS clusters members to the DES SV-A1 imaging described in Chapter 4. In this way, a total of 6 071 galaxies are matched to one of 127 clusters. The cluster richness distribution is shown in Figure 5.9. Of these clusters, 74 have redshifts associated with them, identified in Miller et al. in prep. This redshift distribution is shown in Figure 5.10. In order to control noise, we apply a magnitude cut $i < 21.5$ on cluster members, and we only keep clusters with at least three cluster members which pass this cut. This reduces the catalog to 1 964 galaxies distributed among 73 clusters.

The next step in our analysis is producing photo-z estimates for each of the DES-XCS cluster members. We apply the same ArborZ formalism discussed in Chapter 3 and which we used to estimate photo-zs for DES in Chapter 4. We use the DES training set in the five $grizY$ magnitudes in 50 equally-spaced bins out to $z = 1.0$ with $\sigma_{\rm res} = 0.02$. We then evaluate on the six thousand galaxies in the DES-XCS cluster catalog.

[3]The Earth's atmosphere absorbs X-rays, requiring space-based missions to obtain precision X-ray measurements.

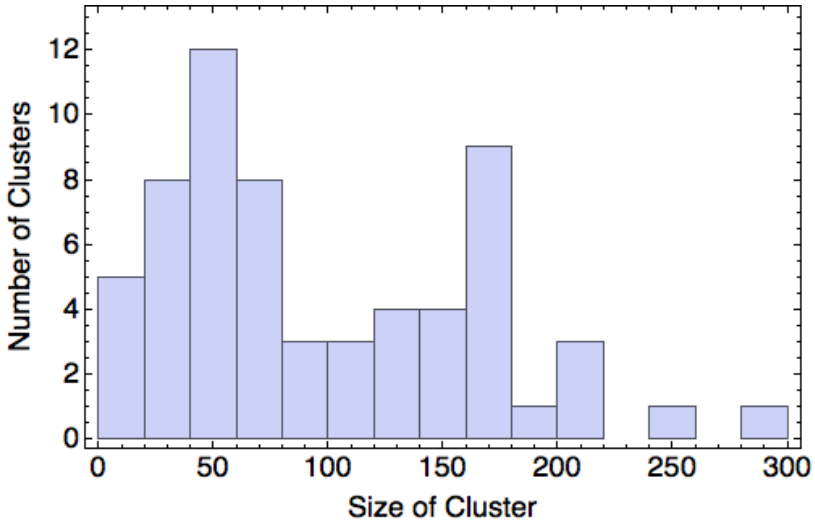

Figure 5.9: The distribution of cluster sizes in the DES-XCS cluster catalog.

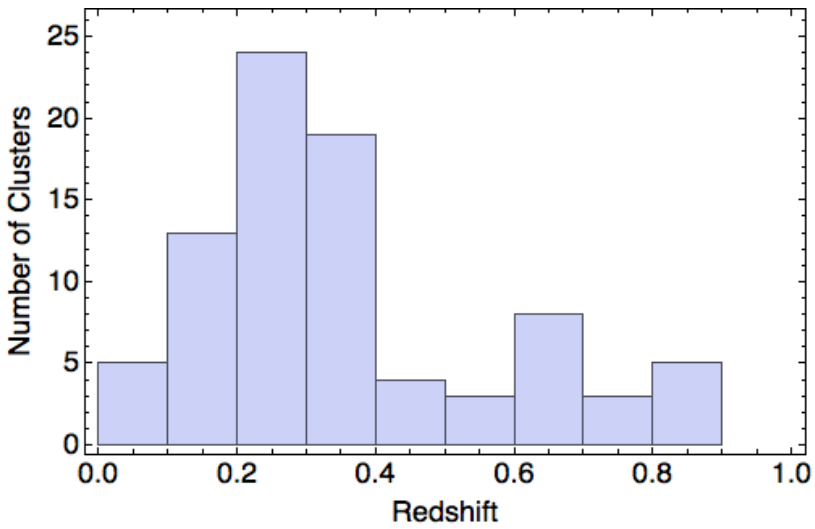

Figure 5.10: The redshift distribution of clusters in the DES-XCS cluster catalog which have known redshifts.

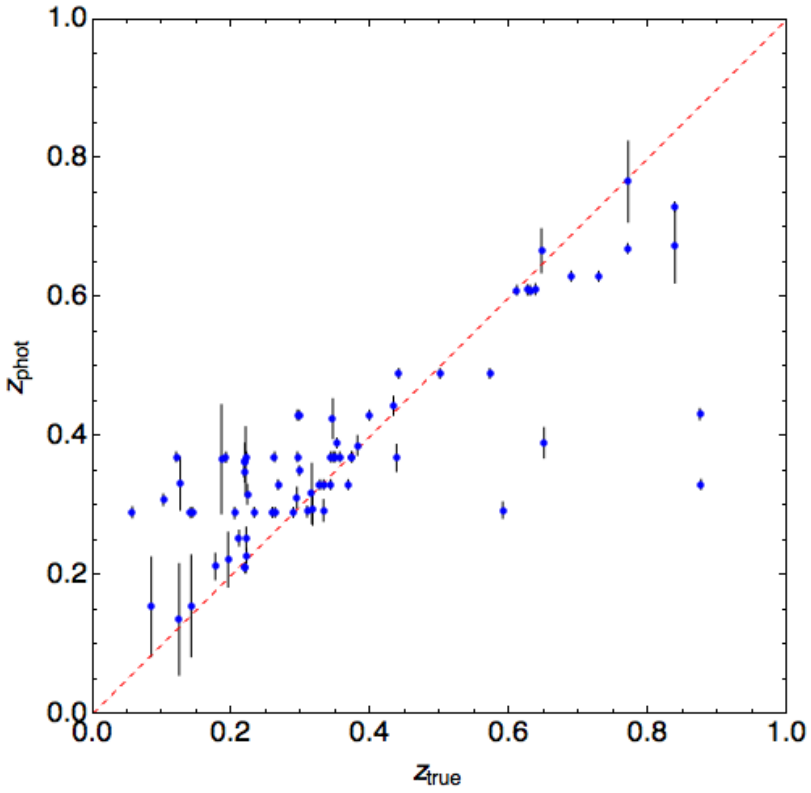

Figure 5.11: The initial DES-XCS photo-z results before applying any further cuts to the cluster members.

We can then use Equation 5.2 to construct cluster $p_{\text{cluster}}(z)$ redshift distributions. However, when we plot the resulting z_{phot} versus z_{true} relation (Figure 5.11) we see an unsatisfactory high-scatter, high-bias result. To combat this, we apply an additional cut to cluster members, demanding that an individual galaxy must pass a p_{peak} threshold. Such a cut will only include galaxies with lower errors (since p_{peak} is correlated with the error σ_z of a photo-z estimate). To determine the optimal cut, we plot p_{peak} versus the RMS of $z_{\text{phot}} - z_{\text{true}}$ of clusters after the cut is applied (keeping only clusters with at least three members). As we see in Figure 5.12, we find a noticeable minimum near $p_{\text{peak}} = 0.75$. This corresponds to keeping 52 clusters in the catalog.

We now return to estimating cluster photo-zs. The new scatter plot is shown in Figure 5.13, where we observe a drastically improved situation. Although we have thrown out about 20 clusters as a result, the remaining clusters photo-zs exhibit a

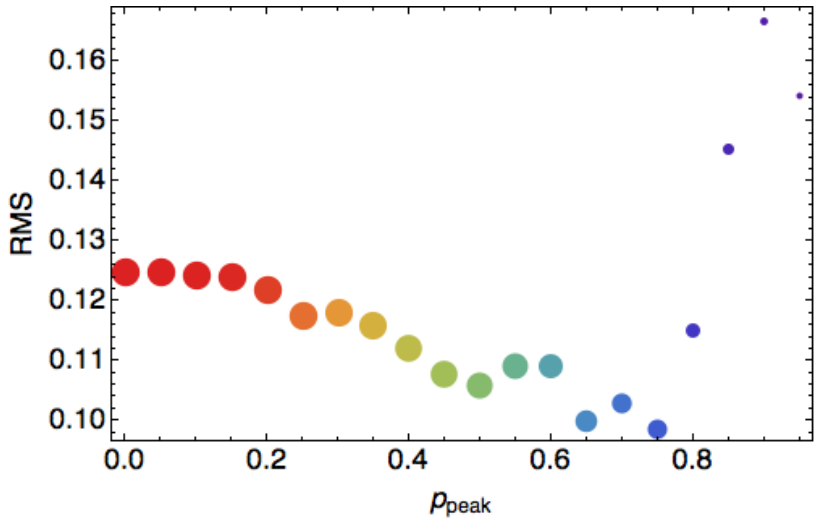

Figure 5.12: Scatter in photo-z residuals $z_{\rm phot} - z_{\rm true}$ as a function of $p_{\rm peak}$. When we apply these cuts, we discard member galaxies, which in turn could remove clusters from the catalog entirely. We show this effect using colors and point sizes. Larger points are associated with a greater number of clusters kept in the catalog. Redder colors are associated with a greater number of member galaxies summed across all clusters.

much tighter relationship with true redshift. Figure 5.14 shows the photo-z distribution for the entire cluster catalog, including clusters for which we do not have spectroscopic redshift estimates.

In Figure 5.15 we show the bias and error plots for this sample. We see that the bias is hovering very close to zero over most of the redshift range. The error is typically near $\sigma_z \approx 0.02$–0.03 for the redshift regions with highest statistics. This makes the stacked cluster photo-z estimates close to the outer limit of current cluster redshift estimation techniques. To compare this performance to other methods on the same data, we turn again to the XCS cluster catalog—the source of our cluster list—which also contains cluster redshifts estimated from the red-sequence.

In Figure 5.16 we show the $z_{\rm phot}$ versus $z_{\rm spec}$ scatter plot using the cluster red-sequence estimates from the XCS cluster catalog. By eye, the red-sequence results look very similar to the stacked $p(z)$ method, except that the red-sequence method has included more high-redshift galaxies which were excluded by the $p_{\rm peak}$ cut from

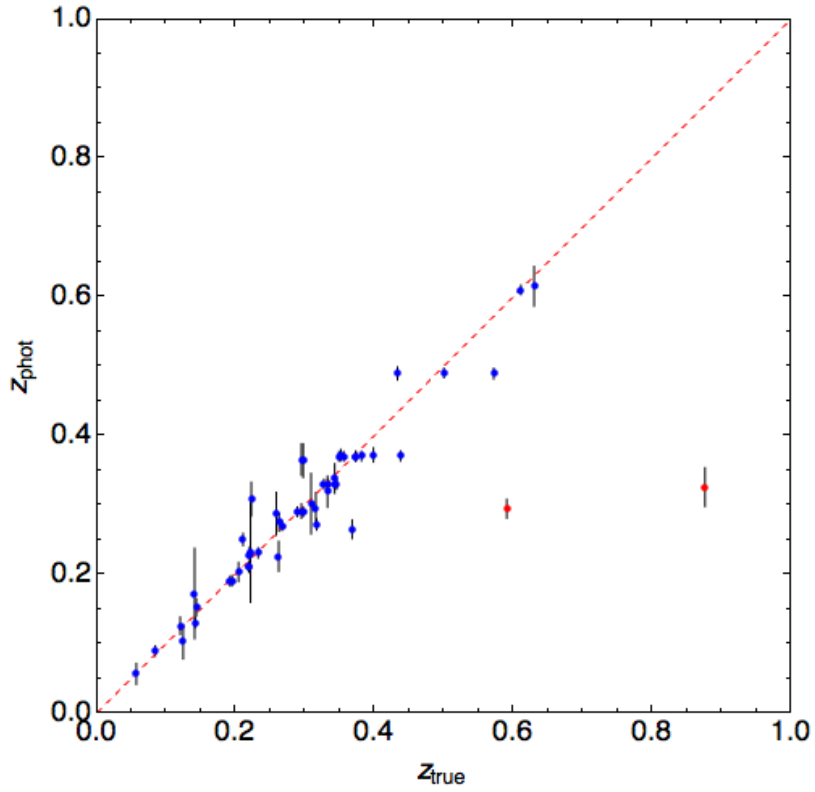

Figure 5.13: The final DES-XCS photo-z results after applying a $p_{\rm peak}$ cut to the cluster members.

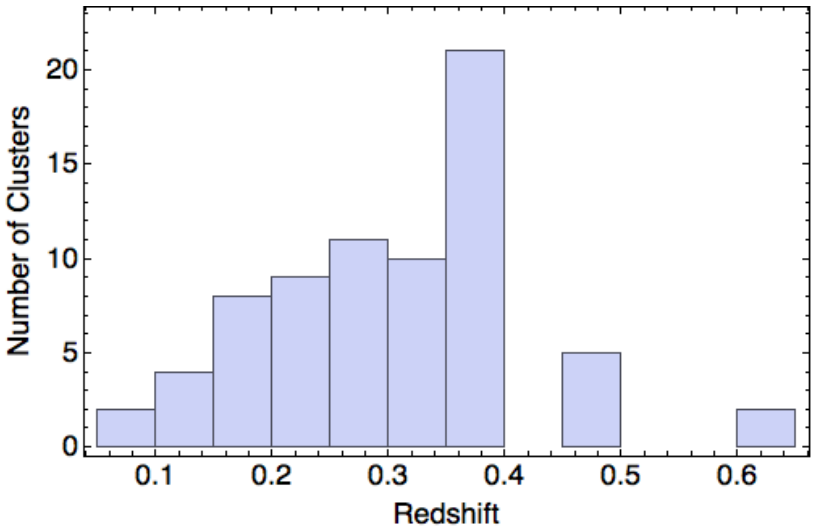

Figure 5.14: The reconstructed redshift distributions $N(z)$ for the entire DES-XCS cluster catalog, including clusters for which we do not have spectroscopic redshifts.

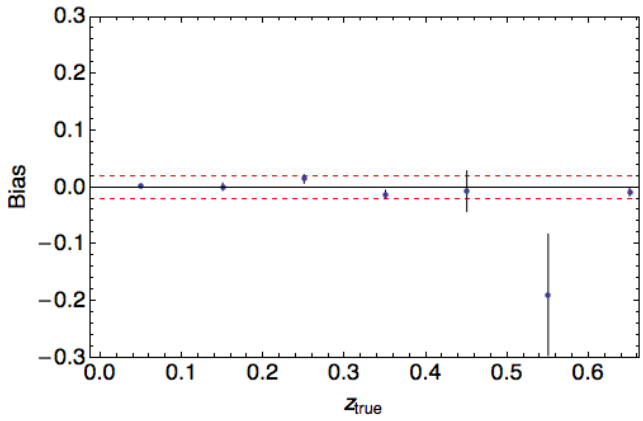

(a) Photo-z bias in the DES-XCS cluster catalog

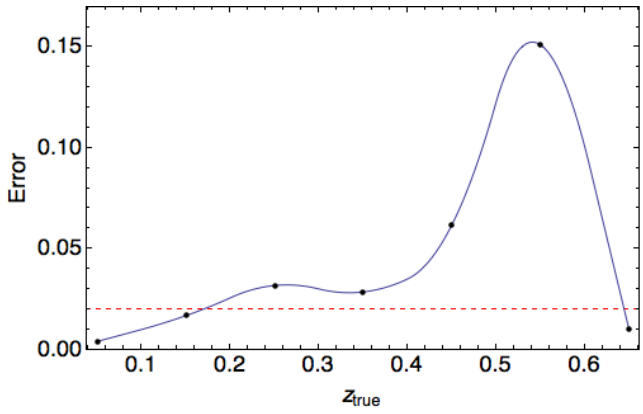

(b) Photo-z error in the DES-XCS cluster catalog

Figure 5.15: The photo-z bias and error plots for the DES-XCS cluster photo-z catalog. The bias and error are both lower than for the DES catalog at large. The dashed red lines indicate the 0.02 level in bias or error, as appropriate, which is typical of cluster redshift estimates (Rykoff et al. 2013).

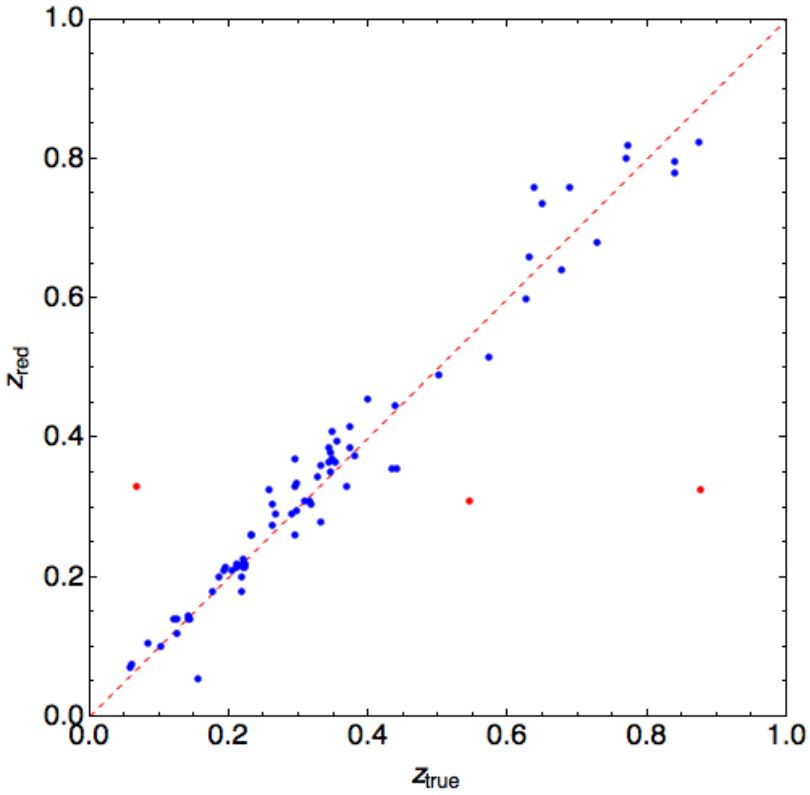

Figure 5.16: Relationship between red-sequence photo-z estimates ($z_{\rm red}$) and true redshift for the DES-XCS catalog.

the stacking method. The bias and error plots for the red-sequence measurements are also qualitatively similar, as we see in Figure 5.17. Although the outliers in the cluster red-sequence method make comparison difficult by eye (except in $0.1 < z < 0.5$), the bias in the red-sequence method is still compatible with that produced by the stacked $p(z)$ method.

To construct a more quantitative comparison, we plot the $z_{\rm phot} - z_{\rm spec}$ residuals in Figure 5.18. Both methods are peaked near zero, though the stacked $p(z)$ method computed using ArborZ seems to show marginally lower scatter. Excluding outliers with $\Delta z > 0.2$ and restricting our redshift range to $z < 0.6$, the RMS scatter in the stacked $p(z)$ method is 0.033 compared to the 0.036 RMS scatter for the red-sequence method. Excluding outliers but without the restricted redshift range, the scatters become 0.033 (ArborZ) and 0.039 (red-sequence). Including outliers and not restricting

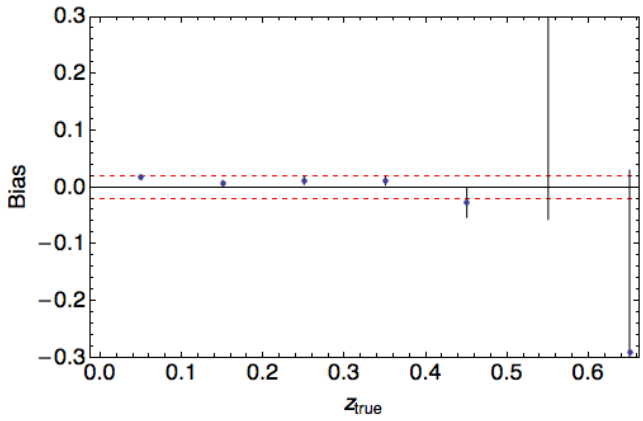

(a) Photo-z bias in the DES-XCS cluster catalog using the red-sequence

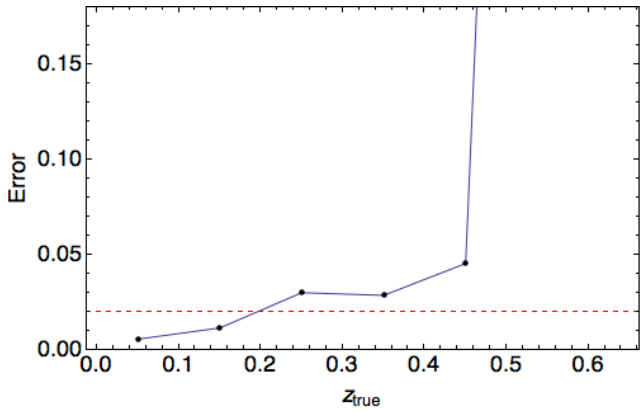

(b) Photo-z error in the DES-XCS cluster catalog using the red-sequence

Figure 5.17: The photo-z bias and error plots for the DES-XCS cluster photo-z catalog, this time calculated from red-sequence estimates. The bias and error are both lower than for the DES catalog at large. The dashed red lines indicate the 0.02 level in bias or error, as appropriate, which approximates the desired performance for competitive cluster redshift estimates.

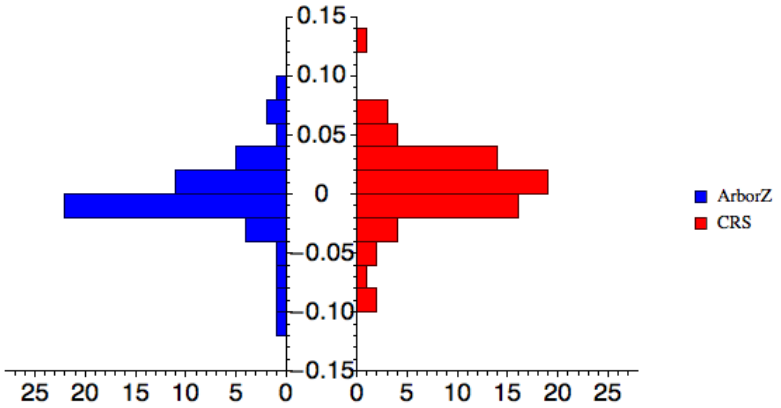

Figure 5.18: Comparison of the photo-z residuals $z_{\rm phot} - z_{\rm spec}$ calculated using the stacked $p(z)$ method (blue) and the traditional red-sequence method (red).

the redshift range yields scatters of 0.092 (ArborZ) and 0.178 (red-sequence) Including outliers and restricting the redshift range gives the largest difference: 0.053 (ArborZ) and 0.178 (red-sequence). To summarize, the photo-zs produced by the red-sequence and by the stacked $p(z)$ method are remarkably comparable. The red-sequence is able to extend out further in redshift without any cuts; on the other hand, ArborZ is able to refine its cluster member selection using $p_{\rm peak}$, which gives it a lower scatter.

Overall, the stacked $p(z)$ method performs well, producing cluster redshift estimates comparable to the commonly-used red-sequence method. ArborZ is capable of refining the quality of its estimates by leveraging the additional information (e.g., $p_{\rm peak}$) produced by full $p(z)$ distributions. This suggests that using ArborZ to produce $p(z)$ estimates, and stacking them as described here, will become a useful tool in cluster cosmology, particularly as large optical surveys, such as DES, continue to observe the universe.

Chapter VI

Conclusion

Cosmology is a fascinating and quickly evolving field, probing the fundamental laws of physics on its largest scales. With the birth of CCDs and the advent of powerful computers, cosmology has truly become a data-driven science, enabling precision cosmology with deep, large-area surveys. In particular, the DES and the SDSS have put tremendous science potential at our fingertips, and given us more data than all previous astrophysical experiments combined. We find ourselves at the golden age of cosmology.

In this thesis, we argued that robust redshift estimates are necessary for achieving many science goals. To this end, we have derived a novel method for estimating photometric redshifts (photo-zs), called ArborZ. We have demonstrated that it performs well, even when the training set is under-represented, as is often the case for many realistic training sets. Moreover, its $p(z)$ distributions make it a particularly powerful method for constraining cosmology.

We applied the ArborZ algorithm to data from the SDSS and the DES surveys, producing photo-z catalogs—complete with $p(z)$ estimates for each galaxy—for hundreds of millions of galaxies. These catalogs open the floodgates to all sorts of cosmological probes. As a particular example, we show that using $p(z)$s can produce good

cluster redshift estimates for joint DES-XCS data.

Despite these optimistic outlooks, photometric redshift estimation still has its limitations. Their errors are several orders of magnitude larger than those for spectroscopic redshifts. Additionally, photo-z estimation algorithms require spectroscopic training sets in order to accurately determine the mapping from photometric observables to redshift. An accurate mapping becomes more and more critical as large sky surveys push our observations into high-redshift realms where few spectroscopic redshift samples exist. This necessitates large-area spectroscopic surveys in order to calibrate photo-z estimates. Experiments such as MS-DESI (Levi et al. 2013) will allow photo-z algorithms to extend their application into high-redshift regimes.

That said, as DES—and future optical imaging surveys—continue to improve and enlarge our galaxy imaging catalogs, so, too, can we expect photo-z techniques to improve. Indeed, DES and other future surveys, such as the LSST, will rely critically on quality photo-z estimates to obtain good science results. With such a tool in hand as the methodology described here, we have bright hopes for the future of cosmology and for our understanding of dark energy in the years to come.

Appendices

Appendix A

The BCAM Alignment System

A.1 Camera Alignment

The Dark Energy Survey (DES; T. Abbott et al. 2005) is a large-scale galaxy survey which will obtain optical imaging for approximately 300 million galaxies during 525 nights of observation spread over five years. The heart of the experiment is the Dark Energy Camera (DECam), which, on account of its 570 megapixel camera, 2.2° diameter field-of-view (3 $\deg^2$ per exposure), and 0.27″ / pixel resolution, is one of the most powerful digital cameras ever built (DePoy et al. 2008). It is installed on the Blanco 4-meter telescope at the Cerro Tololo Inter-American Observatory (CTIO) near La Serena, Chile. Light from the sky reflects off of the Blanco's primary mirror, passes through DECam's five corrector lenses and a broad-band optical filter, and hits one of the seventy-four CCDs on the focal plane.[1] The lenses, filter changer mechanism, shutter, and CCD electronics are supported by a steel structure called the *barrel* (Flaugher et al. 2012). The barrel is mechanically connected to the Blanco's prime focus cage through a hexapod, which can fine-tune DECam's alignment.

[1] Sixty-two CCDs are used for science imaging and are 2048 px × 4096 px. The remaining CCDs are 2048 px × 2048 px, with four begin used for telescope guiding and eight being used for optical focus and alignment.

Aligning such an instrument is a difficult task, and there are countless tolerances which constrain the alignment process. At the detector level, individual CCDs must be flat to within $3\,\mu$m. The entire focal plane must be flat to within $30\,\mu$m (Cease et al. 2008). Particularly difficult to align are the optical correctors, the largest of which weighs 380 lb and yet has an installation tolerance of $\pm 50\,\mu$m in DECam's barrel (Doel, T. Abbott, et al. 2008; Doel, Brooks, et al. 2012). The prime focus cage itself is installed in the telescope's "spider" trusses to within 1.8 mm decenter and 9 mm focus (Muñoz et al. 2012).

In fact, the entire DECam assembly—including corrector lenses, filter changer mechanism, shutter, hexapod, CCD electronics, prime focus cage, spider, and flip rings—weighs approximately 8600 kg (Muñoz et al. 2012). The camera is held by a Serrurier truss 8.5 m from the declination axis of the telescope (as measured along the optical axis), or approximately 10 m from the surface of the primary mirror. As the telescope moves to different pointings on the sky, this massive instrument applies tremendous torques on the telescope system due to the changing gravity vector. These shifting forces inevitably introduce misalignment of the optical and mechanical collimation. Although the Serrurier truss is designed to minimize misalignments of the optical collimation as the telescope moves, it can only do so in a crude, course-grained manner. Small scale misalignment—on the order of a few millimeters in displacement and a hundred arcseconds in rotation—must be corrected using the fine-tuning abilities of the hexapod.

The hexapod can fine-tune the collimation of DECam with respect to the primary mirror in five degrees-of-freedom: displacement (x, y), focus (z), and tip / tilt (θ_x, θ_y). In displacement, the hexapod has a range of ± 32 mm (though the permissible range decreases linearly as focus is increased). It can adjust focus by ± 27 mm. Absolute rotations must be $< 100''$ (Diehl 2012). The hexapod coordinate system, as used here, is a right-hand coordinate system in which $\hat{x}$ is East, $\hat{y}$ is North, and $\hat{z}$ is zenith, all

measured when the telescope is pointing toward zenith. Rotations are defined around the respective $\hat{x}$ and $\hat{y}$ axes in the usual manner for a right-handed coordinate system, and the origin around which rotations are performed is the center of the focal plane.

Although the hexapod is capable of precision alignment, it needs an independent alignment system to provide information about the current misalignment. In essence, a feedback system is necessary, one which can both measure and correct misalignment effects. By measuring misalignments of the telescope between exposures—while the DECam CCDs are being read-out—updated hexapod corrections can be applied in preparation for the next exposure.

There are two independent systems integrated with DECam which can measure misalignments: the BCAM system and the Donut system. The BCAM system reports on the mechanical collimation of the telescope: how DECam is oriented and positioned with respect to the primary mirror. Its counterpart, the Donut system, is designed to estimate the optical collimation of the telescope: artifacts in DECam's images which can be accounted for by relative misalignments of the telescope's optics (Roodman 2012). The author was responsible for researching and developing the BCAM system, which is the focus of this appendix.

A.2 Overview of BCAMs

BCAMs (Brandeis CCD Angle Monitors) are small devices equipped with a CCD and two red diode lasers (Hashemi and Bensinger 2000); see Figure A.1 for a solid model of the BCAMs used on DECam. When a pair of BCAMs are pointed at each other, they can each image the other's lasers. With knowledge of the geometry of a BCAM setup, the locations of the laser spot positions from BCAM pairs can be used to determine static alignments. This has already been done with great success on the ATLAS muon spectrometer at CERN (Amelung et al. 2008). Even without a perfect

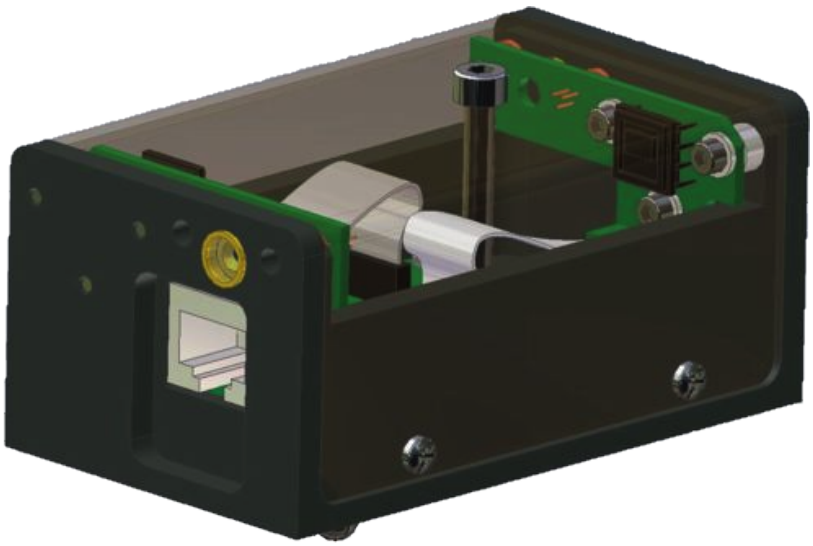

Figure A.1: Solid model of the polar BCAMs used on DECam (top cover not shown). The only difference between this model and the actual BCAMs used on DECam is that the ethernet port is located on the end opposite the laser/CCD. Image taken from Hashemi 2013a.

knowledge of BCAM geometry, relative motions can be inferred by observing how the spot positions shift over time.

The BCAMs installed on DECam were purchased from Brandeis University. They use a TC225P image sensor from Texas Instruments; this CCD has an active area of $3.2\,\text{mm} \times 2.4\,\text{mm}$ with $10\,\mu\text{m}$ square pixels. Since the CCD is positioned 75 mm from the idealized focal plane of the BCAM lens, it has a $43\,\text{mrad} \times 32\,\text{mrad}$ field-of-view. The lens itself is plano-convex with a focal length of 72 mm and a 2 mm aperture. The two lasers are LDP65001E lasers from Lumex with a power output of $< 5\,\text{mW}$ which project a rectangular cone of light with internal angles approximately $40^\circ \times 14^\circ$ (Hashemi 2013a). When a pair of BCAMs are aligned with each other, one can flash its lasers while the other records the image on its CCD; a typical image is shown in Figure A.2. Measuring these spot positions allows the relative positions of the pair to be constrained to within $5\,\mu\text{rad}$.

The BCAMs communicate with the outside world via a device called the Long Wire Data Acquisition (LWDAQ) driver, which was also purchased from Brandeis

Figure A.2: A typical image taken using a BCAM. The BCAM which flashed its lasers in this image did so for 0.6 ms per laser, for a total exposure time of 1.2 ms. The separation between the lasing and imaging BCAMs in this pair is approximately 10 m, as measured along the line-of-sight between them. The image contrast was increased to show the spots more clearly.

University. This is a rather primitive, hub-like device to which all the BCAMs attach. It has one ethernet port which is used to interface with modern computer networks via TCP/IP. The remaining eight ports are used for attaching LWDAQ devices (such as BCAMs), and although they accept standard 8P8C connectors[2], they use a custom protocol and offer a powered connection. The LWDAQ standard is so named because when properly shielded cables are used, it offers reliable signal transmission up to 130 m without a repeater (Hashemi 2013b). In fact, for distance < 13 m, standard, non-shielded CAT-5 cables are sufficient for reliable signal transmission.

An outline of the intended usage of BCAMs to align DECam is as follows. Four pairs of BCAMs have been installed on the Blanco 4 m telescope: four BCAMs mounted directly against the primary mirror (the "lower" BCAMs) and four BCAMs attached to DECam (the "upper" BCAMs). These locations are discussed in detail in

[2]8P8C is the correct name for the connector which is commonly used in ethernet and CAT-5 cables, though it is often called "RJ-45."

Section A.3. At the end of a DECam exposure, the CCDs are read-out, which takes approximately 17 s. Simultaneously, the Observing Tactician (ObsTac) software begins slewing the telescope in preparation for the subsequent exposure. Slewing is expected to complete before the CCDs have been fully read-out. Once the slew is complete, one expects, *a priori*, that the telescope may no longer be in mechanical alignment. So at this point, each BCAM pair flashes its lasers and has its CCDs read out. Shifts in laser spot positions on the BCAM CCDs can then be correlated to the relative mechanical misalignment of DECam and—as we discuss in Section A.8—corrections can be sent to the hexapod to restore alignment and improve image quality. The next exposure begins, and this procedure repeats throughout the night.

The BCAMs supply the hexapod with displacement (x, y) and tip / tilt (θ_x, θ_y) measurements. Although spin θ_z can, in principle, be measured, the hexapod is not capable of motion in this degree-of-freedom, and so we neglect it. The BCAMs can only estimate focus z crudely—to within a few centimeters—by measuring the distance between the two laser spots, but this is too large an error to be of any use in focusing DECam; instead, the Donut system needs to be used for focus.

A.3 Hardware Placement

Four pairs of BCAMs were installed at CTIO in order to align DECam. Although one pair is sufficient to measure all four of the BCAM-sensitive degrees-of-freedom, multiple pairs allow for redundancy as well as the opportunity to reduce errors on the misalignment measurements by a factor of $\sqrt{N_{\text{pairs}}}$. If two pairs are used, then they should optimally be positioned so they are not directly across from each other with respect to the optical axis (i.e., they should not lie on a common radial line through the optical axis); this will ensure that misalignment measurements are, in principle, sensitive to trivial deformations of the primary mirror which would otherwise be

misinterpreted by one pair (or two co-radial pairs). The most natural number of pairs which is not only radially symmetric, but which also ensures redundancy and maximally sensitive measurements, is four.

A.3.1 Upper BCAM positions

One BCAM from each pair needs to be attached to DECam; we call these the "upper" BCAMs, since they are further from the ground when the telescope is at zenith. These BCAMs need to be attached to the barrel of DECam, and not the prime focus cage, since it is the barrel position which houses the camera and which the hexapod can move. The BCAMs also need to have a clear line-of-sight view of the primary mirror so that they can reliable image the BCAMs attached to the primary mirror (the "lower" BCAMs). Additionally, the upper BCAMs need to be attached to a sturdy mounting location. The "ears" of the filter-changer mechanism (FCM) provide such a natural mounting point. These ears are constructed of rigid aluminum, are 0.190 in thick, and welded for additional stiffness. The mounting points are 35.20 in from the optical axis, $\pm 26.5°$ off the $\pm\hat{x}$ axis (East-West), which is the same direction that the FCM extends from DECam (as seen when the telescope is at zenith). See Figure A.3 for an image of the mounting location.

A.3.2 Lower BCAM positions

Since we are interested in misalignments of DECam with respect to the primary mirror, the optimal installation location for the lower BCAMs is directly on the Blanco 4 m primary mirror. The primary mirror is attached to a structure called the mirror cell (see Figure A.4a). It has a large, circular base with twenty-four H-shaped radial supports evenly spaced around its perimeter. Each of these supports has four mounting plates, one in each "corner" of the H, which bolt to the primary mirror. The radial supports themselves serve two purposes beyond merely holding the mirror:

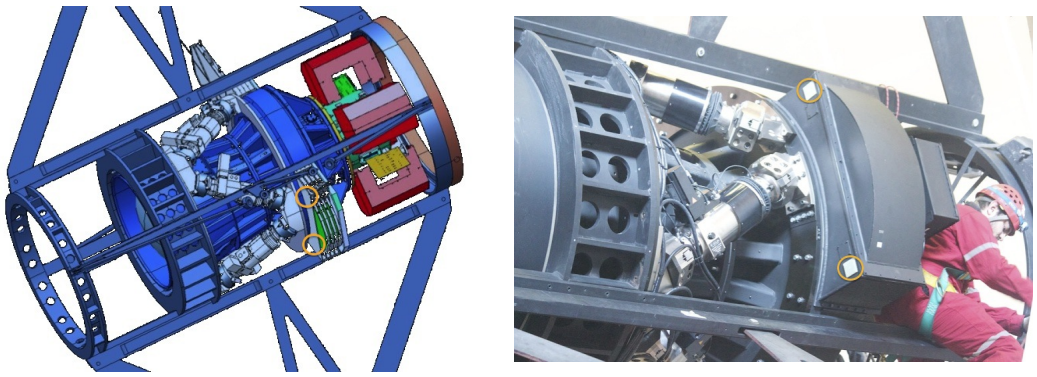

(a) Solid model of DECam

(b) Mounting location on DECam

Figure A.3: The upper BCAMs are mounted on the ears of the filter-changer mechanism. Two of these locations are indicated in orange circles in each image; the other two locations are reflected across the $\hat{y}$ axis.

they have counter-weights attached to distribute the weight of the mirror, and they also attach the primary mirror and mirror cell to the rest of the telescope system.

In the past, these radial supports on the Blanco have failed at a rate of one to two per year (T. M. C. Abbott et al. 2012). They were replaced with twenty-four new supports in 2009, though the old attachment points on the primary mirror were never removed. Thus, there are a total of forty-eight sets of four mounting plates located on the primary mirror, with every other one no longer being used. These mounting plates offer an ideal attachment point for the lower BCAMs (see Figure A.4b).

With these mounting plates available for attaching BCAMs, the author needed to determine the angular position (around the circumference of the mirror) to attach the lower BCAMs. Ideally, we want the lower BCAMs placed at the same angular positions as the upper BCAMs, since this will simplify the alignment process. Note that these angular positions are the same positions which minimize the line-of-sight distances between pairs of upper and lower BCAMs. Therefore, the lower BCAMs were installed at angular positions of 24.0° above and below the East-West axis, as indicated by the orange circles in Figure A.5. In terms of the "spaces" between attachment points seen in Figure A.5, this corresponds to the second space above and below the East-West axis. As we discuss in Section A.4, only two of the four mounting

(a) Mirror cell for primary mirror (b) Mirror cell with primary mirror mounted

Figure A.4: The Blanco 4 m telescope primary mirror cell, with and with the mirror in place. The H-shaped radial radial supports are clearly visible, as are the unused mounting plates on the primary mirror from the old radial supports. Four pairs of these unused mounting plates were used to attach the lower BCAMs.

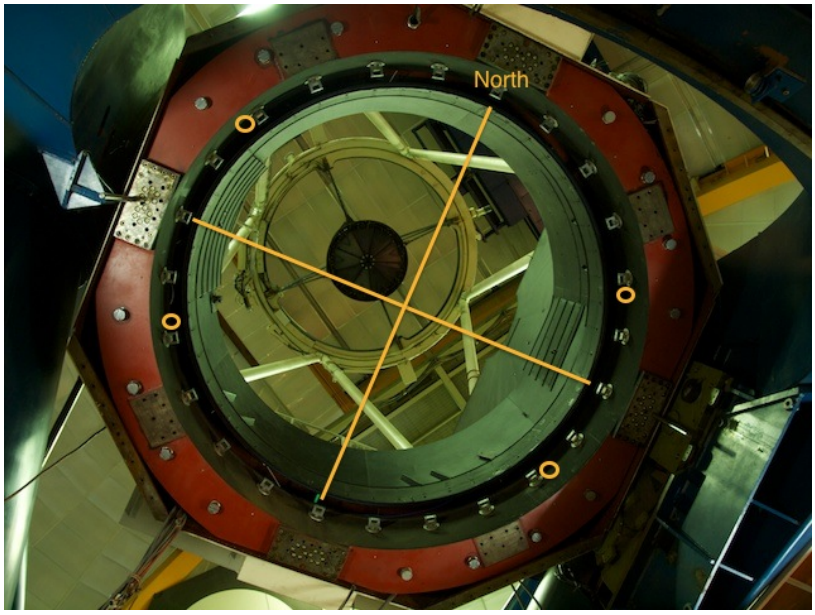

Figure A.5: The Blanco 4 m telescope with its primary mirror removed. This image was taken from underneath the telescope looking toward zenith. The twenty-four attachment points for primary mirror cell are visible just within the red ring (they are small, rectangular metal blocks). The lower BCAMs fit between these attachment points (at the orange circles). Note that DECam is not visible in this image; the old prime focus cage is still in place.

plates at a given angular position on the primary mirror are used for mounting the BCAMs; therefore, we chose the pair of mounting plates furthest from the East-West axis, as this proves to be slightly closer to the optimal mounting point.

A.3.3 LWDAQ position

The LWDAQ driver has several constraints governing the choice of its location. First, it must be within reach of the cables from each BCAM. Second, it requires a network connection in order that the alignment software can communicate with it and the BCAMs. Third, it needs a power supply. Finally, it should be located somewhere which is relatively convenient to access, in case maintenance needs to be performed. The natural location satisfying all of these conditions is the Cassegrain cage below the primary mirror of the telescope.

A.4 Hardware Design

A.4.1 Lower mounting hardware

The lower BCAM mounts were designed in two parts: a large, invar bracket which bolts to the primary mirror and extends up the side of the mirror; and a smaller, kinematic angle mount which holds the BCAM and attaches to the bracket. Together, these two pieces of hardware must secure the lower BCAMs to the primary mirror and allow for some simple angular adjustment in order to align the BCAMs. In the design of these parts, it is critical that the BCAM hardware touch only the mirror and not the telescope; otherwise, the mirror would be mechanically coupled to the telescope in a way that would not be adequately understood.

We first discuss the kinematic angle mounts. These mounts must hold the BCAMs firmly in place and reliably retain their alignment in the event that a BCAM is

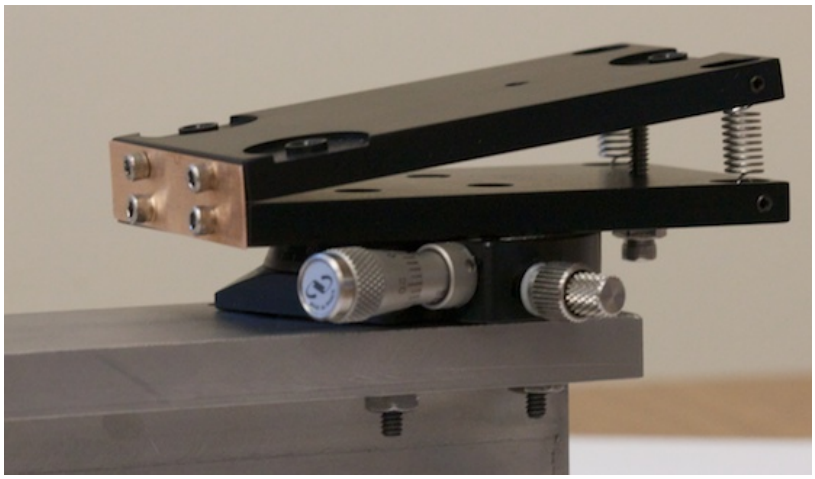

Figure A.6: The kinematic angle mounts which provide angular adjustment for the lower BCAMs. The two plates make an angle of 6.4 deg at the hinge. The three kinematic mounting points are visible on the upper plate. The RS-40 stage is visible between the hinged kinematic mount and the invar bracket.

removed and reattached. We can accommodate reliable and repeatable alignment by use of a kinematic mount. This mount is machined from aluminum, and has three cone-shaped groves in it. Quarter-inch steel ball bearings sit in these grooves and provide precision contact points for the BCAMs. The BCAMs screw into the kinematic mount with a M4 $\times$ 50 mm screw.

This kinematic plate is then hinged to a second plate, also machined from aluminum. In order to prevent reflection during observing, both plates of the hinged kinematic mount are anodized black (type II, class 2). See Figure A.6 for a detailed image of this mount. The hinge between the two plates provides angular adjustment in the radial direction (with respect to the primary mirror). The hinge itself has three parts. First, there is a phosphor-bronze sheet spring which is screwed into both plates, acting as a traditional hinge. Second, there are two, steel expansion springs (Lee Spring, P/N: LE 022C 01 S) which help to close the hinge and guide the plates. These springs were chosen so that their spring constant provided more than twice the force needed to keep the hinge closed even with the BCAM attached and held upside down. Finally, there is an adjustable limit screw which sets the angular adjustment.

This screw is rotated and locked into place so that the two plates form an angle of 6.4°, which is the radial alignment angle predicted by our telescope model.

The tangential angular degree-of-freedom is accommodated by an RS-40 rotation stage from Newport, which provides 360° of course adjustment and 10° of fine adjustment. The two hinged kinematic mount plates screw into the rotation stage, which in turn is bolted to the large, invar bracket (described below). In principle, our telescope model suggests that no tangential adjustment is necessary in order to align the BCAMs; however, we decided to keep this degree-of-freedom available in case it proved necessary to allow a wider range of alignments.

Now we discuss the large, invar brackets. These brackets need to bolt to the primary mirror and extend up the side of the mirror (without contacting the telescope), sticking out a few inches above the mirror in order to attach the BCAMs. The dominant constraints to consider when designing the brackets are flexure due to gravity and deformation due to thermal expansion. Since the BCAMs are precision alignment devices, even small, micron-scale motions of the brackets will dominate the BCAMs statistical error, unnecessarily reducing the accuracy of their misalignment measurements. In fact, using simple trigonometry we can see that a 10 μrad deflection in the bracket will introduce a 100 μm deflection of the BCAMs' lasers as seen by the upper BCAMs.

Several preliminary measures were taken to reduce these effects. First, each bracket was milled from a solid block of invar. This removes shear forces which would otherwise be present at any mechanical junctions. Second, the bracket bolts to two of the mounting plates on the primary mirror, which minimizes both shear on the mirror as well as flexure due to gravity. Finally, gussets were milled into the brackets to distribute strain and minimize flexure.

The exact shape, material, and dimensions of the brackets were fixed using Euler-Bernoulli beam theory to model the cantilever section of the bracket. This equation

relates a simple beam's deflection to its load:

$$\frac{d^2}{dx^2}\left(EI\frac{d^2w}{dx^2}\right) = q \tag{A.1}$$

where $w(x)$ is the deflection of the beam at a distance x from its fixed end, q is the load (force per unit length), E is the elastic modulus of the material, and I is the second moment of area of the beam's cross-section. For the lower BCAM mounting bracket, we consider a constant E (homogeneous material) and constant I (fixed cross-section). We model the load as a uniform load due to the weight of the beam itself (mass M and length L), plus a distributed load from position x_a to x_b due to the added mass of the BCAM and mount which attach to the bracket (mass m and length ℓ). Mathematically, this is:

$$q(x) = \frac{Mg}{L} + \frac{mg}{\ell}\left\langle x - x_a\right\rangle^0\left\langle x_b - x\right\rangle^0 \tag{A.2}$$

where $\langle x-a\rangle^n$ are the singularity functions and g is the acceleration due to gravity. The properties related to the small BCAM mounts (m, ℓ, $\Delta x = x_b - x_a$) are fixed by our previous discussion, and we attach the BCAM mounts to the very end of the bracket (so that $x_b = L$). We chose invar as the material due to its extraordinarily small coefficient of thermal expansion: $\approx$ 1–2 μm/ m/ °C (Shackelford 2001). This choice of material sets the value of E and the density (which gives us M once we choose a cross-section). The cross-section selected is an I-beam, as it proved to be the stiffest cross-section we examined. We chose the precise dimensions of the bracket cross-section by numerically minimizing the maximum deflection of the bracket. As used here, "maximum" deflection is the maximum possible deflection induced by gravity, which occurs when the brackets are oriented horizontally.

The final, machined brackets, shown in Figure A.7, each weigh approximately 12 lb. The maximum flexure predicted by the model for deflections in the radial direction of the primary mirror (i.e., along the "strong" axis of the I-beam) is 1.8 μm,

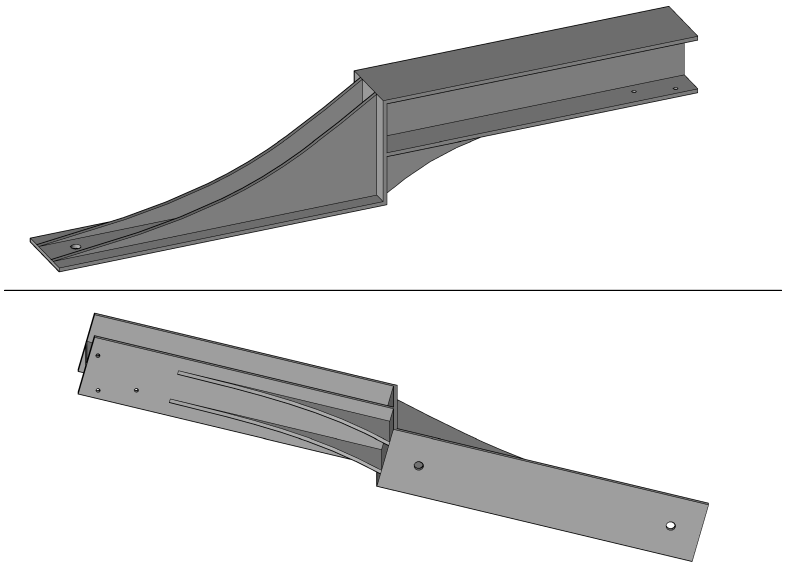

Figure A.7: Solid model of the large, invar mounting brackets used to attach the lower BCAMs and their kinematic angle mounts to the primary mirror. The side which bolts to both the primary mirror and the kinematic angle mount is facing out of the page in the bottom view.

corresponding to an angle of 6.3 μrad and an induced deflection of the BCAM lasers of 63 μm (as seen by the upper BCAMs). The maximum flexure along the tangential ("weak") direction is 3.4 μm, corresponding to an induced deflection of the BCAM lasers of 119 μm. Of course, the telescope will never point horizontally, so the maximum deflection one expects during normal observing will be smaller than the numbers quoted here. In practice, one expects the errors introduced by flexure to be of the same order as the statistical accuracy of the BCAMs.

A fully assembled and aligned lower BCAM is shown in Figure A.8a.

A.4.2 Upper mounting hardware

Compared to the lower mounting hardware, the upper BCAM hardware is smaller and lighter, since no large extension bracket is necessary. However, since a BCAM's field-of-view is an elliptical cone (rather than a circular cone), we can maximize the

(a) Lower BCAM (b) Upper BCAM

Figure A.8: Fully assembled and aligned BCAMs on the Blanco 4 m telescope. The lower BCAM is attached directly to the primary mirror, and extrudes minimally over the mirror through a gap between the mirror surface and the baffle on the telescope wall. The upper BCAM is attached to the ears of the FCM.

overlap of the field-of-views of paired BCAMs by making sure that, in addition to simply "looking toward" one another, they are also aligned in the same way around their line-of-sight. Since the lower BCAMs cannot rotate around their line-of-sight, we build this last rotational degree-of-freedom into the upper mounts. This was achieved with a M-RN-50 ball-and-socket stage from Newport which allows for a full $360 minimizes$ rotation along the line-of-sight and $\pm 25 minimizes$ of tip / tilt in any direction. It also features a locking clutch to prevent further motion after being aligned.

Like the lower BCAMs, the upper BCAMs are attached via a kinematic mount. This kinematic mount is part of a right-angle mount which is bolted to the ball-and-socket stage. The right-angle mount was machined from aluminum in three pieces, all of which were anodized in the same way as the lower BCAM mounting hardware.

The ball-and-socket stage, now holding the right-angle mount and attached BCAM, is mounted onto a thin piece of G10 using eight mounting bridles. The G10 was necessary to ensure electrical isolation between the BCAM system and the rest of DECam. This G10 plate was screwed onto the FCM from the inside (although the screw doesn't pass all the way through, so electrical isolation still holds). It was also glued in place

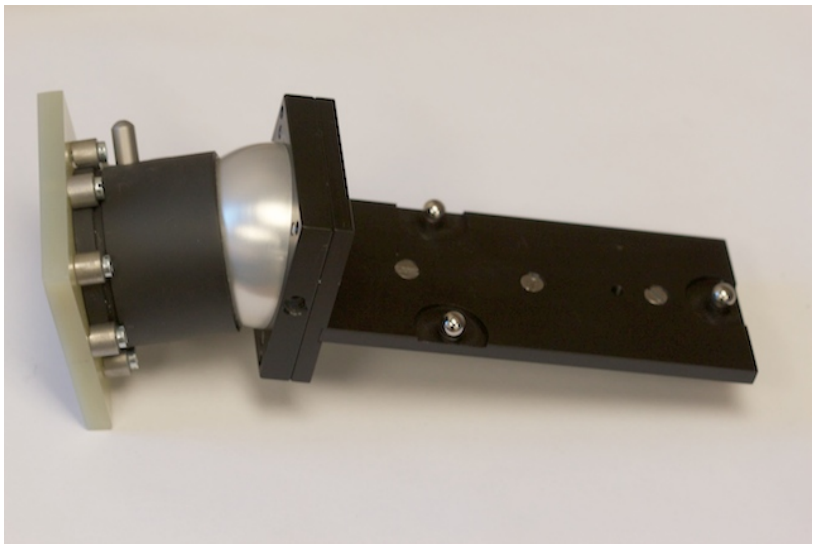

Figure A.9: The right-angle mount, ball-and-socket stage, and G10 plate which constitute the upper BCAM mounting hardware. The three ball bearings can be seen sitting in place on the kinematic mount. The right-angle mount is reinforced with a spine (not visible in this image) on the rear of the kinematic mount.

using DP-190 structural epoxy as an additional precaution against mount failure.

The assembled upper mount (sans BCAM) can be seen in Figure A.9, and a completely assembled and aligned upper BCAM is shown in Figure A.8b.

A.4.3 Cabling

Each of the eight BCAMs requires a cable to connect it to the LWDAQ driver located in Cassegrain cage. In the case of the lower BCAMs, which are mounted to the primary mirror, four 20 m cables were run from the LWDAQ driver, around the perimeter of the primary mirror, and up to each of the four BCAMs. The upper BCAMs were slightly more complicated: four 50 m cables were run from the Cassegrain cage; up the cable wrap, which is attached to one of the legs of the Serrurier truss; through the flip ring (which allows DECam to rotate 180 deg and reflect light back into a detector in the Cassegrain cage, thereby allowing for a larger camera aperture); and finally down the spider to the upper BCAMs.

The cables used for both the upper and lower BCAMs are longer than the non-

shielded limit of 13 m; therefore, the author custom-made CAT-5e cables. The cable was obtained from Four Star Wire & Cable, Inc., P/N F20682ST-20. This cable has four twisted pairs of 24 AWG solid, copper conductors surrounded by a foil shield. Standard CAT-5 modular plugs, strain reliefs, and boots were attached to the cables[3]. These cables were tested before shipment to CTIO and no noticeable signal loss was observed.

A.5 Alignment

As mentioned previously, our telescope models were a great aid in determining course alignment for the BCAMs. In order to align a pair of BCAMs, we continuously flash one of the BCAMs and keep its alignment fixed. We then read off the second BCAM's CCD while adjusting its alignment. This works well because the BCAM lasers act like point sources, so adjusting the flashing BCAM only minimally changes its image as seen by the second BCAM (in the limit of large distances).

Our alignment process for BCAMs on the Blanco 4 m telescope was performed in three parts. First, we roughly aligned the lower BCAMs before installation using the information from our telescope model. The purpose of this was to ensure that the lower BCAMs' lasers would be visible to the upper BCAMs. Then we installed the lower BCAMs, followed by the upper BCAMs. While flashing the lower BCAMs, we moved the upper BCAMs around until alignment was achieved. We then locked the clutch on the upper BCAMs and epoxied the ball-and-socket joint to prevent any long-term shifts in alignment. With the upper BCAMs now aligned, we finally returned to the lower BCAMs and fine-tuned their alignment while flashing the upper BCAMs.

[3]Digi-Key part numbers A9115-ND, A9130-ND, and H11478-ND, respectively

A.6 Software Design

The minimal requirements for BCAM software are two-fold: it must be responsible for communicating with the BCAMs via the LWDAQ driver, and it must interface with the DES software architecture known as SISPI (Survey Image System Process Integration; see Honscheid et al. 2008 for details). SISPI will be responsible for determining when the BCAMs should be read out, it will store all acquired BCAM data in the telemetry database, and it will (via the active optics system) feed back BCAM measurements to the hexapod. However, since the BCAMs are flashed between DECam exposures, the BCAM software also must acquire and analyze BCAM data quickly, lest precious survey time is unnecessarily lost.

Although Brandeis University provides software for communicating with the BCAMs via the LWDAQ driver, the author developed in-house software to solve this problem. There were several factors influencing this decision. First, SISPI uses Python for high-level interfacing with the various software subsystems, whereas the provided LWDAQ software is written in Pascal. By writing hardware drivers in Python, we can ensure smoother software integration. Second, Pascal is itself an obsolete language; code maintenance can become an avalanching nightmare if the software is not written in a ubiquitously supported modern language. Finally, and perhaps most importantly, the provided LWDAQ software is meant to solve the general problem of using BCAMs for alignment, reading off the entire BCAM CCD after every exposure and processing pairs of BCAMs serially. By writing custom software, we can drastically increase the computational performance of the BCAM system.

Let us discuss this last issue in more detail. At 10 m separation, the BCAM spots are small and close together, and the small motions induced by misalignment span only a few pixels on the CCD. This means that we can read out small portions of the CCD without losing any information. On some systems, this speed gain could be

negligible; however, the LWDAQ driver has a relatively slow clock speed, and reading off eight pairs of BCAMs would take nearly eight seconds. Therefore, the author's software determines the optimal CCD patch to read out, reducing readout time by a factor of more than four. This heuristic is described below as part of the `calibrate()` function.

Having decided that these advantages were worth the effort of developing custom in-house software, the author implemented a `BCAM` class in Python which is responsible for creating a network connection with the LWDAQ driver and managing communication with the BCAMs. It functionality is divided across several member functions:

- `start()`. This function takes the IP address of the LWDAQ driver (supplied by SISPI) and initializes the network connection.

- `expose()`. This function clears the BCAM CCDs and flashes their lasers for given exposure interval.

- `readout()`. This function reads out the BCAM CCDs (using the heuristic determined by `calibrate()`). It then analyzes the signal / noise properties of the resulting image and uses them search for laser spots. When it finds a spot, it returns the position of the spot's intensity-weighted centroid, which has a resolution better than $0.5\,\mu$m (5% of a pixel width) (Hashemi 2013a). The function ultimately returns the two (x, y) spot positions from each of the eight BCAMs (32 floating-point numbers total), reporting $(-1, -1)$ for a spot position if it failed to find that spot. All valid spot positions are quoted in units of microns relative to the origin (upper-left corner) of the CCD.

- `acquire()`. This is essentially a convenience function which calls `expose()` and `readout()` in succession.

- `calibrate()`. This function calibrates the BCAM system, and is designed to be called once per night, preferably with the telescope at the same position. The

BCAMs on DECam are calibrated once a night during dome flats when the telescope is pointing at the flat-field screen. The `calibrate()` function works by disabling any heuristics and then calling `acquire()`. It records the current spot positions, which are used to tare the BCAMs system later in the night during calls to `analyze()`. If a calibration of this type were not performed, both long- and short-term misalignments (say, due to nightly or seasonal temperature fluctuations) would contaminate the BCAM signal. The `calibrate()` function also uses these spot positions to determine a rectangular box on each BCAM's CCD where the spots are likely to be. This heuristic is then used (until the next call to `calibrate()` to drastically increase the readout speed.

- `analyze()`. This function takes the spot positions found during `acquire()`, subtracts out the calibration determined by `calibrate()`, and applies the formalism discussed in Section A.7.1 to transform spot positions into hexapod corrections. The return value is a Python dictionary with keys: `dx`, `dy`, `aX`, `aY`, `aZ`. The respective dictionary values are misalignment corrections in units of microns (for displacements `dx`, `dy`) or arcseconds (for rotations `aX`, `aY`, `aZ`). Thus, a `dx` value of 105.4 means that the BCAMs believe that the hexapod should increase its x-position by 105.4 μm. Note that although this dictionary returns `aZ`, it is always set to zero, since we do not attempt to measure this degree-of-freedom.

A.7 Measuring Alignment

Although reading spot positions off of each BCAM CCD is a straightforward task, we still need to map these readings to useful information about misalignments. This transformation is essentially determined by the geometry of the telescope and the positions of the BCAMs. We discuss the model we use, how we calibrate its parameters,

and how we test the calibration in the following sections.

A.7.1 The Alignment Model

The primary assumption in the BCAM alignment model is that the system behaves approximately linearly; that is, shifts in the BCAM spot positions are proportional to hexapod motions. *A priori*, we expect this assumption to hold, since the misalignments we are measuring are extremely small. This assumptions greatly reduces the complexity of the misalignment problem.

To understand the reasoning behind the alignment model we use, it is worth considering what types of motion the hexapod induces in the BCAM spot positions. If the hexapod translates in x or y, the upper and lower BCAMs will measure a symmetric (equal magnitude) spot shift. A hexapod rotation in θ_x or θ_y, however, will induce asymmetric motion. This is because the center of rotation for the hexapod is the center of the focal plane, and the lever arm from the focal plane is much larger for the lower BCAMs than for the upper BCAMs. Thus, one expects the upper BCAMs to measure a larger spot shift than the lower BCAMs when a rotation is performed.

This is the key to being able to separate the degeneracy between translation and rotation. If all hexapod motions effect linear shifts of the BCAMs' spots, then we can separate translation from rotation with a single pair of BCAMs. By subtracting the spot shift seen by the lower BCAMs from the spot shift seen by the upper BCAMs, we are left with asymmetric motion which we can attribute to rotation. Similarly, we can then determine what the symmetric motion due to translation is. Note that in the formalism which follows, we average the positions of the two spots on the CCD in order to get a more precise estimate of the spot position.

Since our fundamental BCAM observable is the shift in spot position relative to the calibration zero-point, it will become useful to talk about "CCD coordinates,"

which are spot positions measured in microns on the BCAM CCDs relative to the current calibration zero-point:

$$\begin{pmatrix} x \\ y \end{pmatrix}_{\text{CCD}} = \begin{pmatrix} x \\ y \end{pmatrix}_{\text{spot}} - \begin{pmatrix} x \\ y \end{pmatrix}_{\text{calibration}} \tag{A.3}$$

and so forth for all BCAMs. Of course, each BCAM is rotated with respect to the East-West axis; that is, a hexapod motion toward East (hexapod $\hat{x}$) will be seen as a motion in both x and y on the BCAM CCDs. Thus, our next step is to rotate the CCD coordinates so that align with the hexapod coordinate system. We call these rotated coordinates "BCAM coordinates." The angle θ which rotates the coordinate system is different for each BCAM. Thus, the conversion from CCD coordinates to BCAM coordinates is:

$$\begin{aligned} \text{Lower:} \quad & \begin{pmatrix} x \\ y \end{pmatrix}_{\text{BCAM}} = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}_{\text{CCD}} \\ \text{Upper:} \quad & \begin{pmatrix} x \\ y \end{pmatrix}_{\text{BCAM}} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}_{\text{CCD}} \end{aligned} \tag{A.4}$$

where the reflection matrix converts the upper BCAM coordinate system into a right-handed coordinate system. The additional reflection for the upper BCAMs is due to the geometry of the BCAM CCD.

Now we need to relate these BCAM coordinates to "hexapod coordinates," the physical coordinates system of the hexapod. To do this, we must first realize that a translation in $-\hat{x}$ aliases a rotation θ_y around $+\hat{y}$, and a translation in $+\hat{y}$ aliases a rotation θ_x around $-\hat{x}$. Proving this is simply of matter of visualization the telescope system and imagining what pairs of BCAMs observe as the hexapod is moved. With this in mind, we can use our assumption about linearity to write:

$$\begin{pmatrix} x \\ y \end{pmatrix}_{\text{BCAM}} = k_t \begin{pmatrix} x \\ y \end{pmatrix}_{\text{hexapod}} + k_r \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} \theta_x \\ \theta_y \end{pmatrix}_{\text{hexapod}} \tag{A.5}$$

Here, k_t is the constant of proportionality for translations; it is the same for all BCAMs. k_r is the constant of proportionality for rotations; it has one value for the lower BCAMs and a second value for the upper BCAMs, due to the nature of the asymmetric spot motion caused by rotations. x, y, θ_x, θ_y are the four hexapod coordinates. Multiplying these hexapod coordinates by -1 yields the values returned by `analyze()`.

Since we have four unknowns—the four hexapod coordinates—we can solve the system exactly by choosing any pair of BCAMs: one lower and one upper. Note that it is necessary to have one lower and one upper BCAM; otherwise, the degeneracy between translations and rotations cannot be broken.

A.7.2 Determining Model Parameters

There are a total of 11 parameters in this alignment model: k_t, the two values of k_r, and the eight possible rotation angles θ. To determine these parameters, we performed a hexapod test after aligning the BCAMs. With the telescope at a fixed location, we recorded the movement of the BCAM spots while moving the hexapod in $+\hat{x}$. This is sufficient to determine k_t and all eight θ values. We then repeated this test, but moved in θ_x instead of $\hat{x}$. This determined the two values of k_r. We continued the test for the remaining degrees-of-freedom to ensure the consistency of our model. We also moved in combinations of the degrees-of-freedom to guarantee that the assumptions about linearity still held. We saw no noticeable deviations from the model throughout these commissioning tests.

We found that $k_t^{-1} = 132\,\mu\text{m}/\,\mu\text{m}$, meaning that a hexapod translation of $132\,\mu$m induces a $1\,\mu$m shift of the spot positions on the BCAMs' CCDs. Since our spot analysis algorithm can determine the centroid of the spot positions to better that $0.5\,\mu$m on the CCDs, our expected statistical error for measuring translation is $< 66\,\mu$m. Similarly, we found that $k_{r,\text{upper}}^{-1} = 2.54\,\text{arcsec}/\,\mu\text{m}$ and $k_{r,\text{lower}}^{-1} = 32.4\,\text{arcsec}/\,\mu\text{m}$, val-

idating our argument that the upper BCAMs are more sensitive to rotations. This also caps our statistical error on measuring rotations to $< 1.25''$. Both of these are in agreement with the quoted BCAM resolution of $5\,\mu$rad.

Of course, we can cross-check the sanity of these measurements by comparing against our detailed telescope model. Since the upper BCAMs are aligned by hand, it is difficult to accurately estimate the θ value *a priori*. However, the three most important parameters—k_t, $k_{r,\text{lower}}$, and $k_{r,\text{upper}}$—-can be predicted. The largest error is for $k_{r,\text{lower}}$, which we measured to be 0.66% larger than predicted, indicating that the model parameters are behaving as expected.

In fact, since these model parameters encode information about the telescope geometry, we can go a step further and attempt to extract geometric information from them. For example, we know that the center of rotation for the hexapod is the DECam focal plane. Therefore, we can numerically find the distance to the focal plane which produces model predictions in agreement with our measured parameters. We estimate the focal plane to be 25.23 in away from the upper BCAMs, in the $+\hat{z}$ direction. The DECam engineers tell us (Stefanik 2012) that the actual specification is 25.45 in, an error of only 0.87%. Again, this gives us confidence that our model is working as expected.

A.7.3 Commissioning Tests

After checking our model parameters, we performed a hexapod test to ensure that the BCAMs were properly inferring misalignments. This was done with the telescope stationary, pointing toward the flat-field screen. The hexapod was stepped slowly in each of its degrees-of-freedom. Each degree-of-freedom was exercised independently, and when it wasn't being changed it was set to zero. Figure A.10a shows the position of each hexapod degree-of-freedom. The transformation from BCAM spots to misalignments was performed using the model parameters determined in Section A.7.2.

The results of each of the four BCAM pairs were averaged together (to improve statistical precision). This final misalignment estimate should be equal to the current hexapod position, up to a constant offset (due to the fact that BCAMs measure relative alignment), so in Figure A.10b we plot the difference and subtract the constant factor, which should be consistent with zero. We also calculate the standard deviations in each of the degrees-of-freedom, and find them to be 27.4 μm, 16.5 μm, 0.71$''$, and 0.8$''$ for each of x, y, θ_x, and θ_y, respectively. After accounting for the improved statistical precision gained by combining multiple BCAM readings, these errors are consistent with the expected errors discussed earlier.

After satisfying ourselves that the BCAM model was working correctly, we used this additional data to better constrain the BCAM model parameters. Re-analyzing this data with the updated model improved the errors by approximately 15%.

A second test was done while the telescope was tracking at the sidereal rate over the course of an hour. Not only was tracking performed (and images taken), the hexapod settings were purposely changed between many of the exposures. We applied our same BCAM model to the data, averaging over the four pairs to improve precision. Figure A.11 shows the mean-subtracted residual between BCAM misalignment measurements and the hexapod settings, which should be consistent with small flexures due to tracking. This is exactly what is observed, with errors of 27 μm, 12 μm, 0.6$''$, and 0.7$''$ on x, y, θ_x, and θ_y, respectively, around their best-fit lines.

A.8 Current Applications and Results

The BCAMs are part of the DECam active optics system (AOS). They are currently flashed between every exposure and the misalignment measurements and fed into the AOS, which uses BCAM and Donut data to produce a hexapod correction. This is done using a pre-generated look-up table (LUT) which contains the expected correc-

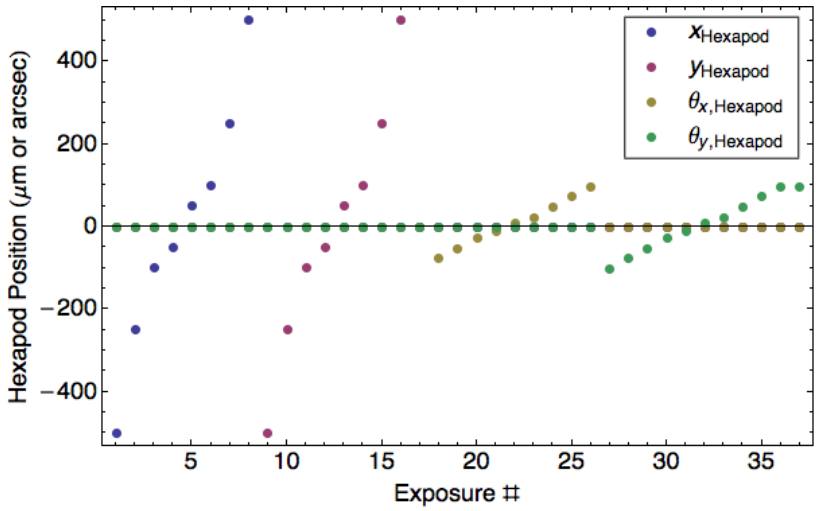

(a) Hexapod Position

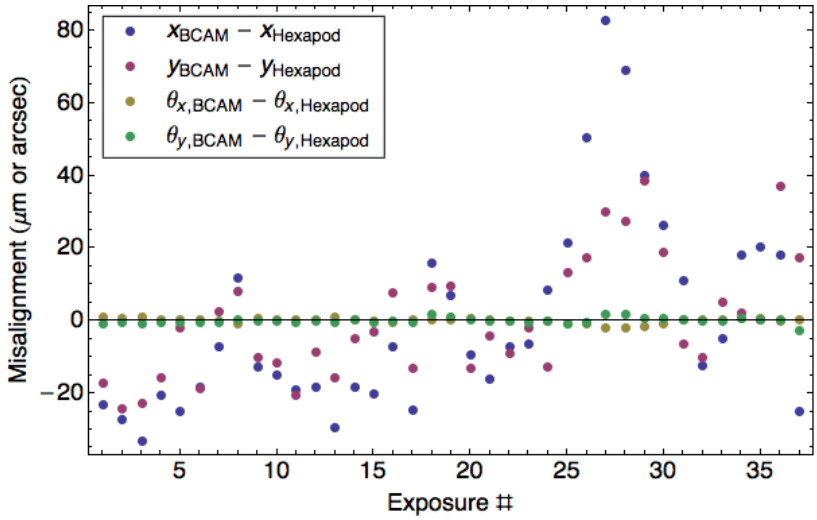

(b) BCAM − Hexapod Residual

Figure A.10: Results of the commissioning test for the BCAMs. The upper figure shows the current hexapod position at each exposure. Each degree-of-freedom was exercised independently, and when a given degree-of-freedom wasn't being tested, it was set to zero. The lower figure shows the BCAM − hexapod residual, which—if the BCAMs are working properly—should be consistent with zero (within errors).

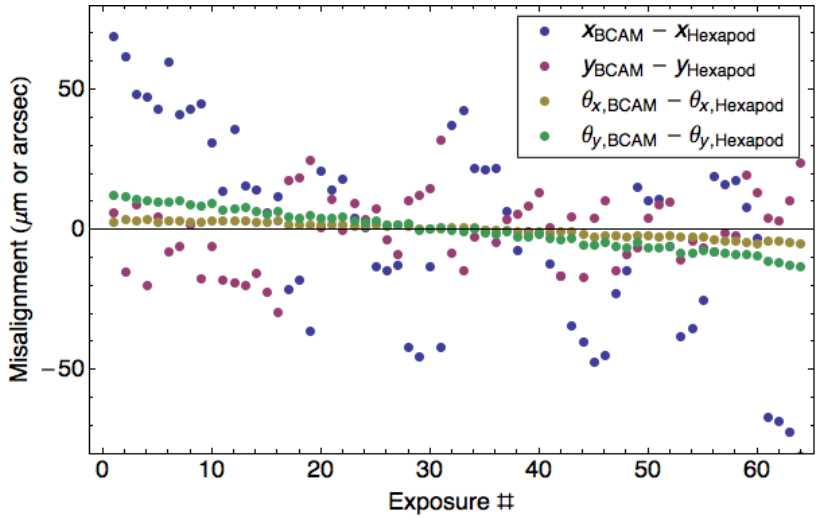

Figure A.11: Results of the BCAM-hexapod test performed while tracking with the telescope at a fixed (RA, Dec) pointing. Shown is the residual between BCAM-inferred misalignments and the actual hexapod settings. Since BCAM measurements are relative, the mean was subtracted to make the residual easier to see. The slopes in the plot are consistent with tracking.

tions for the current telescope pointing. The AOS then uses the current misalignment readings to tweak the LUT values. This allows for crude AOS adjustments even in the absence of BCAM or Donut data, or when data from either system is too noisy to be useful. Moreover, having a LUT contributes greatly to a better understanding of flexure in the DECam / telescope system.

The BCAMs generated their LUT by during an altitude-azimuth scan, which took data at 25 different azimuth angles for each of five different altitudes. See Figure A.12 for plots of the results. We find that displaying these results as a vector field offers the most intuitive insight into the measured misalignments. In words, DECam as a whole sags away from zenith, but the bottom (mirror side) tips toward zenith. This LUT indicates that the maximum absolute misalignments in x, y, θ_x, and θ_y are 2860 μm, 3232 μm, 157″, and 149″, respectively. These are large deviations in optical terms, illustrating the importance of a precision alignment system.

In addition to produce a LUT for the AOS, additional benefits are reaped from multiple altitude-azimuth scans. For example, by taking several scans in succession,

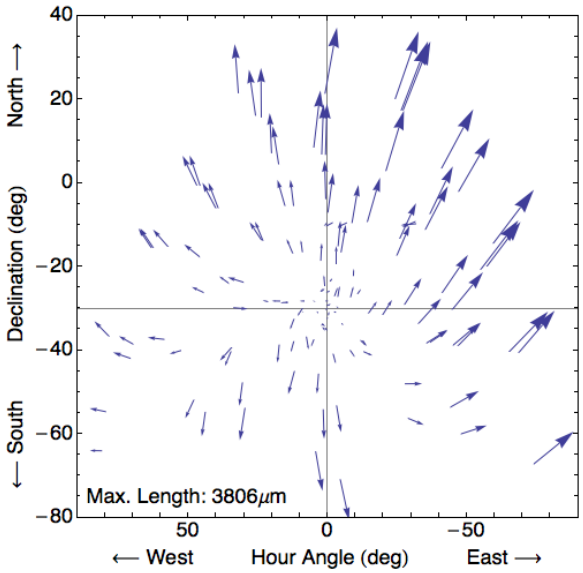

(a) (x, y) Look-up Table

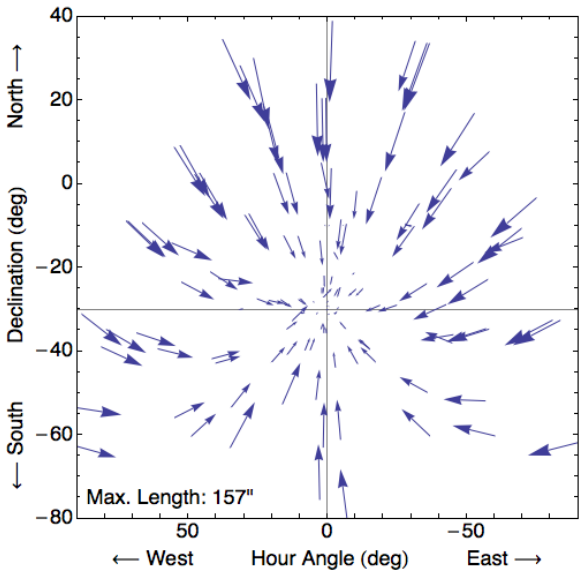

(b) (θ_x, θ_y) Look-up Table

Figure A.12: The BCAM look-up table (LUT). These tables were generated by moving the hexapod in circles in azimuth, and repeating at several altitudes. The results are shown as a vector field zeroed at zenith. The centers of the arrows correspond to the telescope pointing, the length of the arrows are proportional to the magnitude of the misalignment, and the direction of the arrows is the BCAM measurement of the current misalignment. An arrow pointing to the North-East indicates that the hexapod is currently displaced—or angled, depending on the plot—too far toward the North-East (as measured if the telescope were at zenith). For scale, the magnitude of the largest arrow is quoted in the corner of each plot. The gray guide lines indicate zenith.

the effects of mechanical hysteresis in the telescope system can be measured. This was done during DES science verification (SV) to help the telescope engineers understand certain problems with the telescope system. Hysteresis was observed on the order of 100 μm–200 μm in translation and $\approx 10''$ in rotation. Additionally, by repeating altitude-azimuth several times over the course of an observing season, one can observe any long-term shifts in the alignment. Qualitatively, misalignments behave the same from month to month, but quantitively one observes offsets of order 400 μm–500 μm in translation and $\approx 10''$ in rotation. Although some of this shift is due to hysteresis, this may indicate that there are long-term shifts of DECam with respect to the primary mirror.

The BCAMs also have potential for improving image quality. The Donut system measures optical alignment directly, and so, in principle, their misalignment measurements are more strongly correlated with image quality than the BCAMs'. However, the Donut estimates are also much noisier than BCAM data. Thus, it is possible that even though mechanical and optical collimation differ, the higher signal-to-noise measurements produced by the BCAMs may be sufficient to improve image quality.

To study this effect, image quality data from over four months of DES data was collated. During this period, the AOS used only Donut data to adjust the hexapod. In Figure A.13 we show the BCAM misalignments in each degree-of-freedom plotted against the PSF flux radius from DES images. There are clear correlations between BCAM data and image quality. In fact, since the BCAMs were not used to determine the hexapod position, this data necessarily includes noise from the Donut system. If BCAM measurements were incorporated into the AOS, then we would expect the correlation to strengthen and the noise decrease.

The only thing necessary for adapting the BCAMs for use in the AOS is that, since the BCAMs measure relative misalignments, a reference point must be chosen. The best choice, of course, can be read directly off of A.13, and is approximately 680 μm,

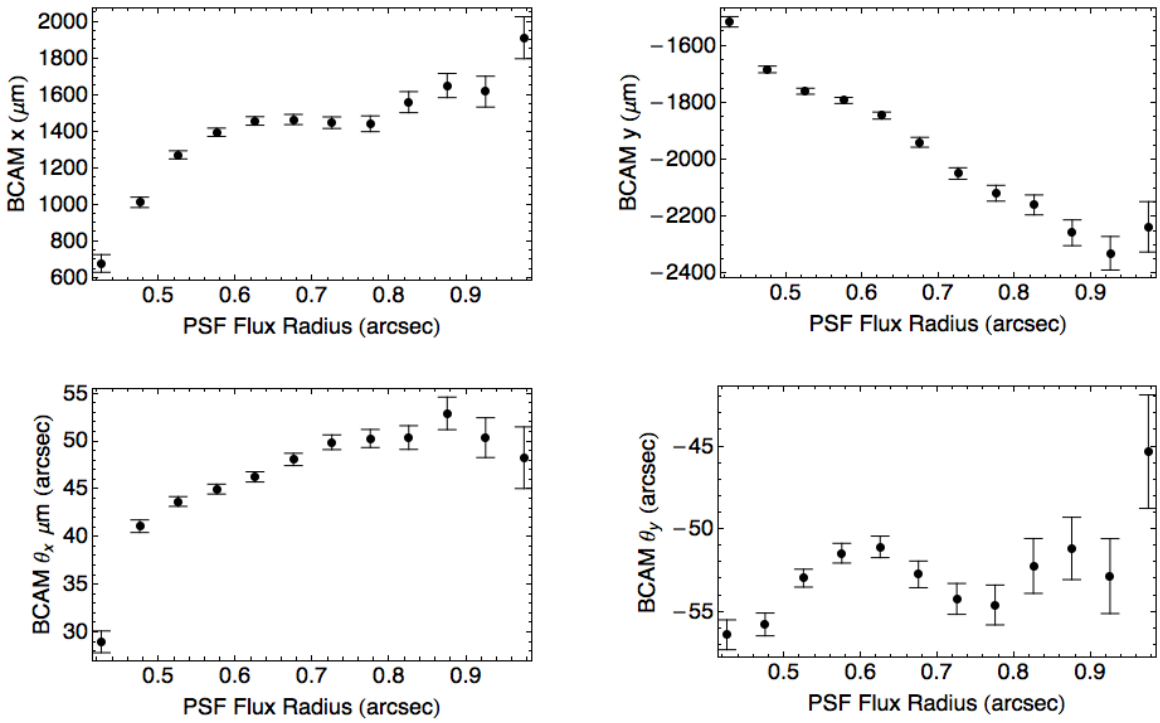

Figure A.13: Correlation of image quality with BCAM data for each of the BCAM degrees-of-freedom. Error bars are standard errors on the mean. Clear trends are seen, implying that the BCAMs could potentially be used for active hexapod corrections, despite measuring mechanical—not optical—collimation.

$-1520\,\mu$m, $29''$, and $56''$ in each of x, y, θ_x, and θ_y, respectively. So to produce the BCAMs' estimate of the optimal optical alignment (relative to the current hexapod position), one simply subtracts these values from the BCAM readings.

A.9 Conclusions

The BCAMs are precision alignment devices which use lasers to determine relative motion. Four BCAMs have been installed on DECam and are paired with four respective BCAMs attached directly to the Blanco 4 m primary mirror. By measuring the shifts of the laser spot positions on each BCAM's CCD, the mechanical misalignment of DECam with respect to the primary mirror can be estimated. Their measured accuracy is $< 30\,\mu$m in translation and $< 1''$ in rotation. These low noise measurements provide intuitive insight into the mechanical deformation of the telescope system.

Although optical and mechanical collimation are different, the BCAMs readings are clearly correlated with image quality, suggested that improved image quality can be obtained by further integration of the BCAMs into the AOS / hexapod feedback system.

Appendix B

Adaptive Catalog Matching

B.1 The Matching Problem

One of the major impediments to producing quality photometric redshift (photo-z) catalogs is the lack of sufficient spectroscopic catalogs. Empirical photo-z estimation algorithms require a training set of imaging observables paired with precision (spectroscopic) redshifts. However, photometric and spectroscopic data are obtained separately on different instruments, and often on different telescopes and at different sites. Moreover, spectroscopic and imaging surveys have different science goals and methodologies. Spectroscopic surveys usually focus on relatively small patches of sky with specific targets designated for spectroscopy; imaging surveys, on the other hand, generally cover a larger patch of sky and do not need to pre-select targets, since everything in the instrument's field-of-view is imaged. To further complicate the discussion, galaxy positions are subject to astrometry errors, and so the same galaxy will be measured to have a slightly different position on the sky when measured multiple times or by different instruments. This makes the problem of combining, or *matching*, spectroscopic and imaging catalogs into a single training set a challenging problem.

An obvious and simple way to solve this problem is to project both catalogs

onto the two-dimensional sky. Each spectroscopic entry is then *matched*—assigned to a imaging counterpart—to the nearest object in the imaging catalog. As a sanity check, a maximum matching distance (called the "cutoff" radius) is usually enforced. For its simplicity, this method works quite well, and is the most commonly used matching method is astronomy (Blanton, D. J. Schlegel, et al. 2005; Dahlen et al. 2010; Schneider et al. 2010).

This "nearest-neighbor" matching is simple, and as such overlooks several subtleties. For example, the deeper the imaging catalog is—or, conversely, the shallower the spectroscopic catalog—the more likely it is that you will produce an incorrect match. In addition, different surveys have different astrometry errors, so although a match may be correct if the astrometry is good in both the photometric and spectroscopic surveys, bad astrometry in one survey may lead to incorrect matches (or an overly conservative cutoff radius, artificially reducing the number of matches). To solve these shortcomings, additional complexity needs to be introduced into the matching algorithm. To the author's knowledge, the only alternative is a likelihood ratio approach (Sutherland and Saunders 1992; Wang and Rowan-Robinson 2009; Wolstencroft et al. 1986). This method is much improved over the crude, nearest-neighbor approach, though it does require estimating several *a priori* distributions.

As an alternative, we introduce a new method for matching. It has the advantage of being simple, with an elegant analytic result. Of particular interest is its ability to estimate the purity of the matched catalog. This allows one to study the effects of impure training sets—ones sprinkled with incorrect redshifts—on the quality of photo-z estimates.

B.2 Derivation

To motivate the derivation, consider again the nearest-neighbor method. This method works well most of the time because modern astrometry errors are small compared to the average nearest-neighbor distance between two galaxies on the sky. Of course, error distributions have tails, and sometimes a galaxy's true position must fall further than average from its true position. Additionally, once a clean sample is selected in the imaging catalog, there may no longer exist a true match to a given galaxy. These problems introduce rogue matches which require a cutoff radius to curtail. This is not a problem *per se*, but it begs the fundamental question: what cutoff radius should be chosen so as to maximize the number of correct matches found while simultaneously minimizing the number of incorrect matches?

To answer this, consider the problem of matching a particular spectroscopic entry to exactly zero or one entries in an imaging catalog. We assume there is no systematic offset in the position measurements of the two catalogs[1]. If the spectroscopic and imaging surveys have normally-distributed statistical errors σ_s and σ_p, respectively, then we can simplify the math by treating the astrometry in the spectroscopic catalog as perfect and instead considering the astrometry in the imaging catalog to have an error $\sigma = \left(\sigma_s^2 + \sigma_p^2\right)^{1/2}$. The second assumption that we make is to ignore clustering in the imaging catalog; that is, we treat the spatial distribution of galaxies in the imaging catalog as a Poisson distribution with expectation value λ. This is a simplifying assumption: we know that galaxy-galaxy spatial correlations exist; however, these correlations are small on the arcsecond scales that we are considering.

Now consider a single entry in the spectroscopic catalog. Let f be the probability that the imaging catalog contains an entry describing the same physical object; we

[1] Although the formalism could be extended to include a systematic offset, it is typically easier to simply use several bright stars to subtract out any systematic differences between the photometric and spectroscopic catalog before matching.

do not, *a prioiri*, know f, and the value of f is not affected by our ability or inability to actually find the match in the imaging catalog. By hypothesis, astrometry errors are Gaussian, so the probability $p_m(r)$ of the photometric match falling within a distance r—as measured on the sky (in, for example, arcseconds)—is simply f times the integrated two-dimensional Gaussian PDF:

$$P_m(r) = f\left(\frac{1}{\sqrt{2\pi\sigma}}\right)^2 \int_{r'=0}^{r}\int_{\theta=0}^{2\pi} e^{-(r')^2/(2\sigma^2)} r'\, dr'\, d\theta \tag{B.1}$$

$$= f\left(1 - e^{-r^2/(2\sigma^2)}\right) \tag{B.2}$$

Meanwhile, "background" galaxies may randomly fall within this radius r; these are galaxies which are not associated with the spectroscopic object, but rather other entries in the imaging catalog which introduce matching noise in the catalog. If these are approximately Poisson distributed, then on average we will find $\lambda = \pi r^2 \rho$ within a radius r (where ρ is the average areal density of galaxies in the imaging catalog) and the probability of finding exactly k background galaxies within radius r is:

$$P_{\rm Pois}(k \mid \lambda(r) = \pi r^2 \rho) = \frac{\lambda^k e^{-\lambda}}{k!} \tag{B.3}$$

Therefore, the probability of finding at least one background galaxy in radius r is one minus the probability of not finding any galaxies in the area:

$$P_b(r) = P_{\rm Pois}(k > 0 \mid \lambda(r)) \tag{B.4}$$

$$= 1 - P_{\rm Pois}(k = 0 \mid \lambda(r)) \tag{B.5}$$

$$= 1 - e^{-\pi r^2 \rho} \tag{B.6}$$

The probability, therefore, that the distance $r_{\rm NN}$ to the nearest neighbor of the spectroscopic galaxy—whether or not it is a correct match at radius r_m or a background galaxies at radius r_b—falls within a radius r is:

$$P(r) = \Pr(r_{\rm NN} < r) = \Pr(r_b < r \cup r_m < r) \tag{B.7}$$
$$= \Pr(r_b < r) + \Pr(r_m < r) - \Pr(r_b < r \cap r_m < r) \tag{B.8}$$
$$= P_b(r) + P_m(r) - P_b(r)P_m(r) \tag{B.9}$$
$$= \left(1 - e^{-\pi r^2 \rho}\right) + f\left(1 - e^{-r^2/(2\sigma^2)}\right) - f\left(1 - e^{-\pi r^2 \rho}\right)\left(1 - e^{-r^2/(2\sigma^2)}\right) \tag{B.10}$$
$$= f\left(1 - e^{-r^2/(2\gamma^2)}\right) + (1 - f)\left(1 - e^{-\pi r^2 \rho}\right) \tag{B.11}$$

where the reduced density γ is defined by:

$$\gamma^2 = \frac{\sigma^2}{1 + 2\pi\sigma^2\rho} \tag{B.12}$$

With Equation B.11 in hand, we begin to derive all sorts of interesting quantities. Particularly, we can calculate the *purity* of the matched sample; that is, the fraction of matches which are correct. To do this, we first calculate the probability $P_{b,m}(r)$ that we encounter a background galaxy before the true match. First we write down the PDF of $P_m(r)$:

$$p_m(r) = \frac{d}{dr}P_m(r) = -\frac{f}{\sigma} r e^{-r^2/(2\sigma^2)} \tag{B.13}$$

Then we have[2]:

[2]If the first line of this derivation isn't clear, consider two distributions A and B with PDFs $p_A(r)$ and $p_B(r)$, respectively. The probability of drawing number from distribution B that is less than r is:

$$\int_{-\infty}^{r} p_B(r')dr' \tag{B.14}$$

The probability of drawing r from A and a number less than r from B is:

$$p_A(r)dr \int_{-\infty}^{r} p_B(r')dr' \tag{B.15}$$

Therefore, the probability of drawing a number less than r from A but a yet smaller number from B is:

$$\int_{r'=-\infty}^{r} \left[p_A(r')dr' \int_{r''=-\infty}^{r'} p_B(r'')dr'' \right] \tag{B.16}$$

Recognizing the inner integral as the probability distribution $P_B(r)$ completes the proof.

$$P_{b,m}(r) = \int_0^r p_m(r')P_b(r')dr' \tag{B.17}$$

$$= \int_0^r \left[p_m(r')\left(1 - e^{-\pi r'^2\rho}\right)\right] dr' \tag{B.18}$$

$$= P_m(r) - f\frac{\gamma^2}{\sigma^2}\left(1 - e^{-r^2/(2\gamma^2)}\right) \tag{B.19}$$

This, in turn, allows us to calculate the probability $P_{m,b}(r)$ that, within radius r, we encounter the true match (a correct match) before encountering a spurious background galaxy[3].

$$P_{m,b}(r) = P_m(r) - P_{b,m}(r) = f\frac{\gamma^2}{\sigma^2}\left(1 - e^{-r^2/(2\gamma^2)}\right) \tag{B.20}$$

This equation is important because we can use it to find a natural cutoff radius for matching. Let N_S be the number of possible spectroscopic matches[4]. Once $\frac{d}{dr}(N_S P_{m,b})$ drops below one we are only contaminating our matched catalog by increasing the matching radius. Thus, the natural cutoff radius r_0 is the solution to:

$$N_S \frac{dP_{m,b}}{dr}\bigg|_{r=r_0} = 1 \tag{B.21}$$

This occurs when:

$$r_0 e^{-r_0^2/(2\gamma^2)} = \frac{\sigma^2}{fN_S} \tag{B.22}$$

Unfortunately, this is a transcendental equation and must be solved numerically. Note that this equation will generally have two solutions; we choose the larger possible value

[3]One might naïvely think that we could just calculate $P_{m,b}(r)$ in a similar fashion, *mutatis mutandis*, to how we calculated $P_{b,m}(r)$. However, although there can only be one true match, it is quite possible that we encounter many background galaxies in a radius r. The integrals in the calculation of $P_{m,b}(r)$ would become far more difficult (if even tractable), and so we elect the roundabout, but much easier, calculation shown here.

[4]This could be calculated, for example, by assigning each spectroscopic entry its nearest imaging entry and counting the number of unique imaging entries assigned.

of r_0. Using r_0 as a cutoff radius will produce the largest matched catalog possible which does not suffer from adding too many false positives.

We continue our derivation by defining the purity $q(r)$ as the fraction of times we find a correct match out of all matches assigned (whether correct or not).

$$q(r) = \frac{P_{m,b}(r)}{P(r)} = \frac{\gamma^2/\sigma^2}{1 + \frac{1-f}{f}\left(\frac{1-e^{-\pi r^2 \rho}}{1-e^{-r^2/(2\gamma^2)}}\right)} \tag{B.23}$$

This lets us quantify the success of our matching algorithm. Moreover, it also hands us another method for determining a cutoff radius: instead of asking what the optimal catalog size should be, we can also impose a purity limit. Inserting this purity limit into Equation B.23 and solving for r, we can determine the cutoff radius needed to satisfy our purity constraint. We also note that the purity is bounded from above and below by values that we can calculate:

$$q_{\max} = \lim_{r\to 0} q(r) = \frac{f}{f + 2\pi\rho\sigma^2} \tag{B.24}$$

$$q_{\min} = \lim_{r\to \infty} q(r) = \frac{f}{1 + 2\pi\rho\sigma^2} \tag{B.25}$$

Intuitively, one can only obtain perfect $q = 1$ purity—even at $r = 0$—only if there are no background galaxies or there is no uncertainty in the astrometry. Similarly, if we disregard the cutoff radius and always match to the nearest neighbor, our chances of finding the correct match decrease quickly as astrometry degrades or as our photometric catalog includes fainter galaxies (effectively increasing ρ).

Equations B.11, B.20, and B.23 are the primary results of this discussion. The first allows us to predict the number of matches we'll find by assigning nearest-neighbor matches out to cutoff radius r. The second informs us of how many of these matches are correct and gives a natural cutoff radius. The last is used to calculate the purity derived from the cutoff radius, and can be used to impose a minimum purity on the

matched catalog.

B.3 Performance

Our model contains three free parameters which need to be determined in order to apply it to the matching problem: f, σ, and ρ. With so few parameters, most fitting algorithms should perform well out-of-the-box. For the sake of completeness, however, we note that there is an additional constraint that can be imposed on ρ. If one considers our derivation of $P_b(r)$, we realize that it is simply the cumulative distribution function of nearest-neighbor distances in the photometric catalog. Thus, we can estimate ρ in a straightforward manner by simply fitting a one-parameter function to the photometric catalog's nearest-neighbor CDF. Of course, at small separations matches are dominated by astrometry errors, and so as long as the astrometry is good ($\lesssim 0.5''$) one can typically ignore ρ altogether (setting it to zero). However, some surveys—particularly older ones or ones which experienced optics problems—have particularly poor astrometry, and so the full formalism must be retained.

As an example of the model's fit to real data, we consider a recent spectroscopic survey, the Arizona CDFS Environment Survey (ACES; Cooper, Yan, et al. 2012), which we match to the photometry of the Dark Energy Survey (DES; T. Abbott et al. 2005). ACES contains approximately 5 000 secure, unique redshifts out to $z \approx 1.4$ and $R_{\mathrm{AB}} = 24.1$ in the $30' \times 30'$ extended Chandra Deep Field–South (CDFS) region.

To match the catalog, we first need to fit $P(r)$, the total number of nearest-neighbor matches, to the data, fixing our three parameters. We find $f = 93.7\%$, $\sigma = 0.18''$, and $\rho = 6.4345 \times 10^{-3}\,\mathrm{arcsec}^{-2}$. The results are shown in Figure B.1. Our model fits the data wonderfully. Already we can determine the number of matches which would be found by a naïve matching algorithm. With our fit in hand we can calculate the purity, shown in Figure B.2. The maximum purity we can expect is

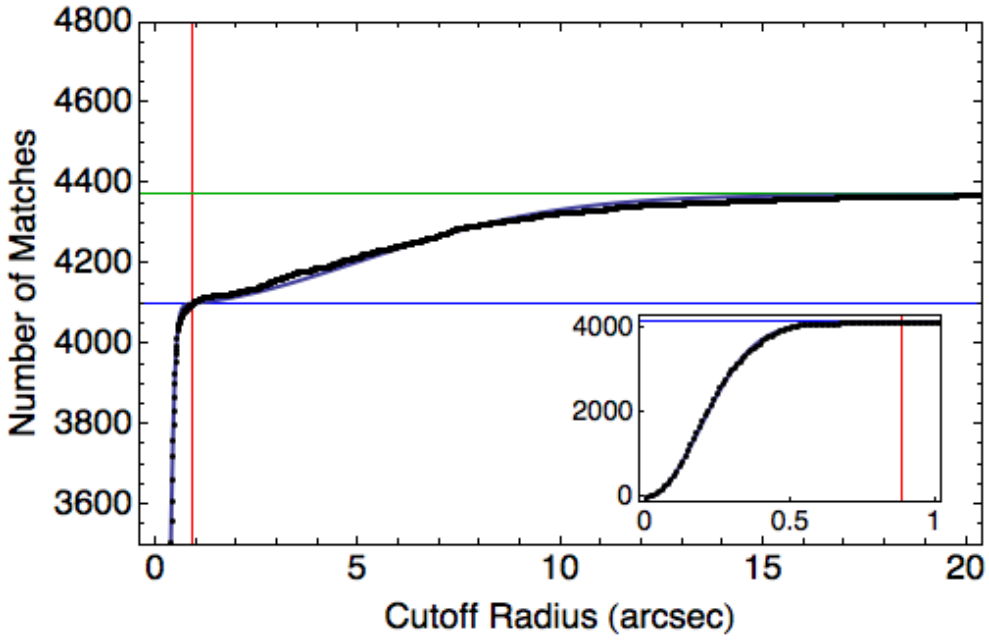

Figure B.1: The results of fitting $P(r)$—the number of nearest-neighbor matches within cutoff radius r—to ACES and DES. We plot $N_S P(r)$ to cast the numbers into absolute terms. The small, inset plot is a close-up of the fit for small r. Across the entire range we see excellent agreement with the model. The red line marks the natural cutoff radius r_0. The blue line corresponds to the size of the catalog at this radius. The green line is the maximum possible number N_S of matches. Holding all other model parameters fixed, changing f will shift the position of the blue line, changing σ shifts the red line (since astrometry errors dominate at small separation), and shifting ρ changes how quickly the fit asymptotically approaches the green line (since background galaxies dominate at larger separation). The position of the green line is determined by the overlap between the two catalogs.

99.85%, and the minimum (if we disregard a cutoff radius) is 93.60%. The optimal catalog size is obtained with $r_0 = 0.89''$, corresponding to 4 101 matches with purity 99.76%. Detailed statistics are given in Table B.1.

Purity (%)	Radius (arcsec)	Matches	Added	Correct
99.85	0	0	—	—
99.82	0.55	4052	4052	4044
99.80	0.68	4094	42	4086
99.76	0.89	4101	7	4091
99.6	1.41	4108	7	4091
99.4	1.89	4116	8	4091
99.2	2.28	4124	8	4091
99	2.63	4133	9	4091
98	4.06	4175	42	4091
97	5.36	4218	43	4091
96	6.74	4262	44	4091
95	8.46	4307	45	4091
94	11.53	4352	45	4091
93.60	∞	4371	19	4091

Table B.1: Detailed statistics for matching ACES to DES. The first and last lines of the table correspond to maximum purity and minimum purity, respectively. The red entries correspond to the optimal cutoff radius r_0. Several intermediate values are shown as well for comparison. The column labeled "Added" is the number of additional matches found compared to the previous line. "Correct" is the number of correct matches found at that cutoff radius (the number of matches found times the purity).

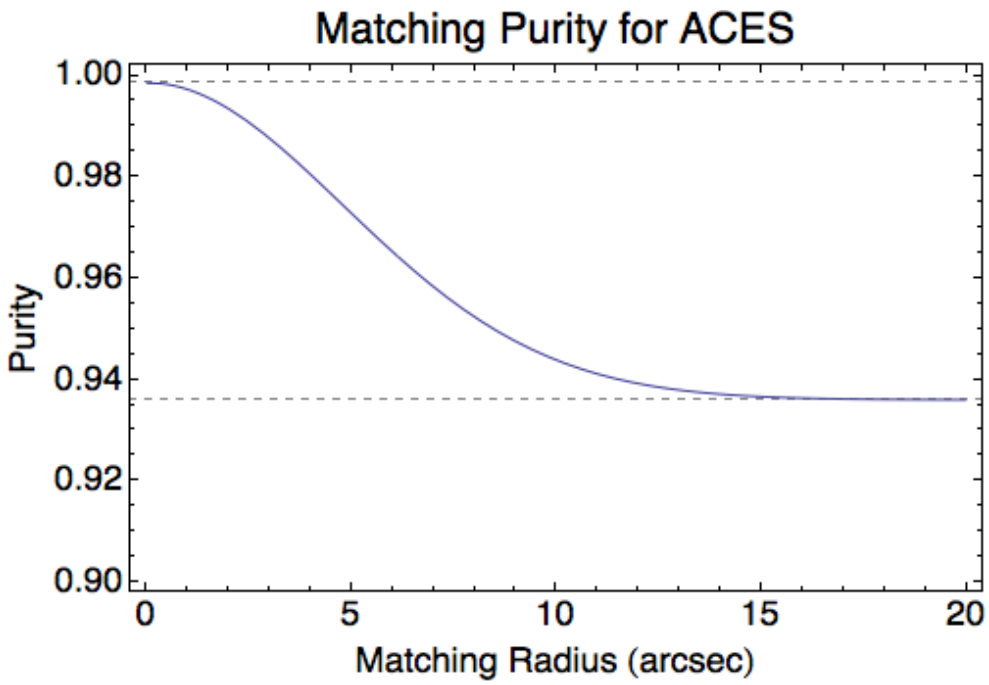

Figure B.2: The purity $q(r)$ resulting from matching ACES galaxies to the nearest DES neighbors with cutoff radius r. The maximum purity (at radius $r = 0$) is 99.85%. The minimum purity is 93.60%. Detailed matching statistics are shown in Table B.1.

Bibliography

Abbott, T. M. C. et al. (Sept. 2012). "Dark Energy Camera installation at CTIO - overview". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Ed. by L. M. Stepp, R. Gilmozzi, and H. J. Hall. Vol. 8444. Ground-based and Airborne Telescopes IV. Amsterdam, Netherlands. DOI: 10.1117/12.925692 (see p. 140)

Abbott, T. et al. (Oct. 2005). "The Dark Energy Survey". In: *pre-print*. arXiv:astro-ph/0510346 (see pp. 31, 40, 97, 133, 171)

Ade, P. A. R. et al. (2013). "Planck 2013 results. XVI. Cosmological parameters". In: *pre-print*. arXiv:1303.5076v1 (see pp. 7, 12)

Ahn, C. P. et al. (July 2013). "The Tenth Data Release of the Sloan Digital Sky Survey: First Spectroscopic Data from the SDSS-III Apache Point Observatory Galactic Evolution Experiment". In: *pre-print*. arXiv:1307.7735 (see p. 79)

Allen, P. D. et al. (Sept. 2006). "The Millennium Galaxy Catalogue: bulge-disc decomposition of 10095 nearby galaxies". In: *Monthly Notices of the Royal Astronomical Society* 371. 2006MNRAS.371....2A, 2–18. DOI: 10.1111/j.1365-2966.2006.10586.x (see p. 83)

Alpher, R. A., H. Bethe, and G. Gamow (1948). "The Origin of Chemical Elements". In: *Physical Review* 73, 803–804. DOI: 10.1103/PhysRev.73.803 (see p. 16)

Amara, A. and A. Réfrégier (Nov. 2007). "Optimal surveys for weak-lensing tomography". In: *Monthly Notices of the Royal Astronomical Society* 381. SAO/NASA:

2007MNRAS.381.1018A, 1018–1026. DOI: 10.1111/j.1365-2966.2007.12271.x (see pp. 41–43)

Amelung, C. et al. (Mar. 2008). *The Optical Alignment System of the ATLAS Muon Spectrometer Endcaps*. ATLAS Note, ATL-MUON-PUB-2008-003 (see p. 135)

Annis, J. et al. (Oct. 2010). *The Dark Energy Survey Science Requirements Document*. Tech. rep. Internal DES document (docdb #20). The Dark Energy Survey Collaboration (see p. 41)

Arnouts, S. et al. (Dec. 1999). "Measuring and modelling the redshift evolution of clustering: the Hubble Deep Field North". In: *Monthly Notices of the Royal Astronomical Society* 310. SAO/NASA: 1999MNRAS.310..540A, 540–556. DOI: 10.1046/j.1365-8711.1999.02978.x (see p. 44)

Balestra, I. et al. (Mar. 2010). "The Great Observatories Origins Deep Survey. VLT/VIMOS spectroscopy in the GOODS-south field: Part II". In: *Astronomy & Astrophysics* 512. SAO/NASA: 2010A&A...512A..12B. DOI: 10.1051/0004-6361/200913626 (see p. 99)

Ball, N. M., R. J. Brunner, A. D. Myers, N. E. Strand, S. L. Alberts, and D. Tcheng (Aug. 2008). "Robust Machine Learning Applied to Astronomical Data Sets. III. Probabilistic Photometric Redshifts for Galaxies and Quasars in the SDSS and GALEX". In: *The Astrophysical Journal* 683. SAO/NASA: 2008ApJ...683...12B, 12–21. DOI: 10.1086/589646 (see p. 45)

Ball, N. M., R. J. Brunner, A. D. Myers, N. E. Strand, S. L. Alberts, D. Tcheng, and X. Llorà (July 2007). "Robust Machine Learning Applied to Astronomical Data Sets. II. Quantifying Photometric Redshifts for Quasars Using Instance-based Learning". In: *The Astrophysical Journal* 663. SAO/NASA: 2007ApJ...663..774B, 774–780. DOI: 10.1086/518362 (see p. 45)

Bardeen, J. M., P. J. Steinhardt, and M. S. Turner (1983). "Spontaneous creation of almost scale-free density perturbations in an inflationary universe". In: *Physical Review D* 28, 679–693. DOI: 10.1103/PhysRevD.28.679 (see p. 11)

Baum, W. A. (Feb. 1957). "Photoelectric determinations of redshifts beyond 0.2 c". In: *The Astronomical Journal* 62. SAO/NASA: 1957AJ.....62....6B, 6–7. DOI: 10.1086/107433 (see p. 40)

— (1962). "Photoelectric Magnitudes and Red-Shifts". In: *Proceedings from IAU Symposium*. Ed. by G. C. McVittie. Vol. 15. Problems of Extra-Galactic Research. SAO/NASA: 1962IAUS...15..390B. International Astronomical Union Symposium. Macmillan Press, New York (see p. 40)

Bender, R. et al. (Oct. 2001). "The FORS Deep Field: Photometric Data and Photometric Redshifts". In: *Deep Fields*. Ed. by S. Cristiani, A. Renzini, and R. E. Williams. SAO/NASA: 2001defi.conf...96B. ESO Astrophysics Symposia. Springer-Verlag, 96. DOI: 10.1007/10854354_18 (see p. 44)

Benítez, N. (June 2000). "Bayesian Photometric Redshift Estimation". In: *The Astrophysical Journal* 536. SAO/NASA: 2000ApJ...536..571B, 571–583. DOI: 10.1086/308947 (see p. 44)

Blanton, M. R., H. Lin, et al. (Apr. 2003). "An Efficient Targeting Strategy for Multiobject Spectrograph Surveys: the Sloan Digital Sky Survey "Tiling" Algorithm". In: *The Astronomical Journal* 125. SAO/NASA: 2003AJ....125.2276B, 2276–2286. DOI: 10.1086/344761 (see p. 79)

Blanton, M. R., D. J. Schlegel, et al. (June 2005). "New York University Value-Added Galaxy Catalog: A Galaxy Catalog Based on New Public Surveys". In: *The Astronomical Journal* 129. SAO/NASA: 2005AJ....129.2562B, 2562–2578. DOI: 10.1086/429803 (see p. 165)

Bolzonella, M., J.-M. Miralles, and R. Pelló (Nov. 2000). "Photometric redshifts based on standard SED fitting procedures". In: *Astronomy & Astrophysics* 363. SAO/NASA: 2000A&A...363..476B, 476–492 (see p. 44)

Bond, J. R. and G. Efstathiou (Oct. 1984). "Cosmic background radiation anisotropies in universes dominated by nonbaryonic dark matter". In: *Astrophysical Journal Letters* 285, 45–48. DOI: 10.1086/184362 (see p. 16)

— (1987). "The statistics of cosmic background radiation fluctuations". In: *Monthly Notices of the Royal Astronomical Society* 226. SAO/NASA: 1987MNRAS.226..655B, 655–687 (see p. 11)

Bond, J. R., L. Kofman, and D. Pogosyan (Apr. 1996). "How filaments of galaxies are woven into the cosmic web". In: *Nature* 380. SAO/NASA: 1996Natur.380..603B, 603–606. DOI: 10.1038/380603a0 (see p. 19)

Bower, R. G., J. R. Lucey, and R. S. Ellis (Feb. 1992). "Precision Photometry of Early Type Galaxies in the Coma and Virgo Clusters - a Test of the Universality of the Colour / Magnitude Relation - Part Two - Analysis". In: *Monthly Notices of the Royal Astronomical Society* 254. SAO/NASA: 1992MNRAS.254..601B, 601 (see p. 109)

Brammer, G. B., P. G. van Dokkum, and P. Coppi (Oct. 2008). "EAZY: A Fast, Public Photometric Redshift Code". In: *The Astrophysical Journal* 686. SAO/NASA: 2008ApJ...686.1503B, 1503–1513. DOI: 10.1086/591786 (see p. 44)

Breiman, L. et al. (1984). *Classification and regression trees*. Wadsworth statistics/probability series. Wadsworth International Group (see pp. 55, 60)

Bromm, V. and R. B. Larson (Sept. 2004). "The First Stars". In: *Annual Review of Astronomy and Astrophysics* 42, 79–118. DOI: 10.1146/annurev.astro.42.053102.134034 (see p. 19)

Bromm, V. and A. Loeb (May 2006). "High-Redshift Gamma-Ray Bursts from Population III Progenitors". In: *The Astrophysical Journal* 642, 382–388. DOI: 10.1086/500799 (see p. 19)

Byram, E. T., T. A. Chubb, and H. Friedman (Apr. 1966). "Cosmic X-ray Sources, Galactic and Extragalactic". In: *Science* 152. SAO/NASA: 1966Sci...152...66B, 66–71. DOI: 10.1126/science.152.3718.66 (see p. 121)

Cannon, R. et al. (Oct. 2006). "The 2dF-SDSS LRG and QSO (2SLAQ) Luminous Red Galaxy Survey". In: *Monthly Notices of the Royal Astronomical Society* 372. SAO/NASA: 2006MNRAS.372..425C, 425–442. DOI: 10.1111/j.1365-2966.2006.10875.x (see p. 82)

Cappellaro, E. et al. (1997). "The rate of supernovae from the combined sample of five searches". In: *Astronomy & Astrophysics* 322, 431–441 (see p. 14)

Carliles, S. et al. (Sept. 2008). "Photometric Redshift Estimation on SDSS Data Using Random Forests". In: *Astronomical Data Analysis Software and Systems XVII*. Ed. by R. W. Argyle, P. S. Bunclark, and J. R. Lewis. Vol. 394. Astronomical Society of the Pacific Conference Series. SAO/NASA: 2008ASPC..394..521C. Astronomical Society of the Pacific, 521 (see p. 45)

Carroll, S. M. (Aug. 2001). "Dark Energy and the Preposterous Universe". In: *preprint*. arXiv:astro-ph/0107571 (see p. 20)

Cavaliere, A. G., H. Gursky, and W. H. Tucker (June 1971). "Extragalactic X-ray Sources and Associations of Galaxies". In: *Nature* 231. SAO/NASA: 1971Natur.231..437C, 437–438. DOI: 10.1038/231437a0 (see p. 121)

Cease, H. et al. (July 2008). "The Dark Energy Survey CCD imager design". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Ed. by I. S. McLean and M. M. Casali. Vol. 7014. Ground-based and Airborne Instrumentation for Astronomy II. Marseille, France. DOI: 10.1117/12.788200 (see p. 134)

Clerke, A. M. (1908). *A Popular History of Astronomy*. 4th. London: Adam and Charles Black (see p. 4)

Coe, D. et al. (Aug. 2006). "Galaxies in the Hubble Ultra Deep Field. I. Detection, Multiband Photometry, Photometric Redshifts, and Morphology". In: *The Astronomical Journal* 132. SAO/NASA: 2006AJ....132..926C, 926–959. DOI: 10.1086/505530 (see p. 44)

Coil, A. L. et al. (Nov. 2011). "The PRIsm MUlti-object Survey (PRIMUS). I. Survey Overview and Characteristics". In: *The Astrophysical Journal* 741. SAO/NASA: 2011ApJ...741....8C. DOI: 10.1088/0004-637X/741/1/8 (see pp. 83, 99)

Colless, M. et al. (Dec. 2001). "The 2dF Galaxy Redshift Survey: spectra and redshifts". In: *Monthly Notices of the Royal Astronomical Society* 328. SAO/NASA: 2001MNRAS.328.1039C, 1039–1063. DOI: 10.1046/j.1365-8711.2001.04902.x (see pp. 82, 98)

Collister, A. A. and O. Lahav (Apr. 2004). "ANNz: Estimating Photometric Redshifts Using Artificial Neural Networks". In: *The Publications of the Astronomical Society of the Pacific* 116. SAO/NASA: 2004PASP..116..345C, 345–351. DOI: 10.1086/383254 (see p. 45)

Connolly, A. J., I. Csabai, et al. (Dec. 1995). "Slicing Through Multicolor Space: Galaxy Redshifts from Broadband Photometry". In: *The Astronomical Journal* 110. SAO/NASA: 1995AJ....110.2655C, 2655. DOI: 10.1086/117720 (see p. 45)

Connolly, A. J., A. S. Szalay, et al. (Sept. 1997). "The Evolution of the Global Star Formation History as Measured from the Hubble Deep Field". In: *The Astrophysical Journal Letters* 486. SAO/NASA: 1997ApJ...486L..11C, L11–L14. DOI: 10.1086/310829 (see p. 45)

Cooper, M. C., J. A. Aird, et al. (Mar. 2011). "The DEEP3 Galaxy Redshift Survey: Keck/DEIMOS Spectroscopy in the GOODS-N Field". In: *The Astrophysi-*

cal Journal Supplement Series 193. SAO/NASA: 2011ApJS..193...14C. DOI: 10.1088/0067-0049/193/1/14 (see p. 83)

Cooper, M. C., R. Yan, et al. (Sept. 2012). "The Arizona CDFS Environment Survey (ACES): A Magellan/IMACS Spectroscopic Survey of the Chandra Deep Field-South". In: *Monthly Notices of the Royal Astronomical Society* 425. SAO/NASA: 2012MNRAS.425.2116C, 2116–2127. DOI: 10.1111/j.1365-2966.2012.21524.x (see pp. 99, 171)

Copernicus, N. (1543). *De revolutionibus orbium coelestium*. Nuremberg apud Johannes Petreius (see p. 3)

Crawford, A. B., D. C. Hogg, and L. E. Hunt (July 1961). "Project Echo: A Horn-Reflector Antenna for Space Communication". In: *Bell System Technical Journal* 40, 1095–1116 (see p. 10)

Csabai, I. et al. (Feb. 2003). "The Application of Photometric Redshifts to the SDSS Early Data Release". In: *The Astronomical Journal* 125. SAO/NASA: 2003AJ....125..580C, 580–592. DOI: 10.1086/345883 (see p. 45)

Cunha, C. E. et al. (July 2009). "Estimating the redshift distribution of photometric galaxy samples - II. Applications and tests of a new method". In: *Monthly Notices of the Royal Astronomical Society* 396. SAO/NASA: 2009MNRAS.396.2379C, 2379–2398. DOI: 10.1111/j.1365-2966.2009.14908.x (see pp. 40, 45, 46, 78)

Dahlen, T. et al. (Nov. 2010). "A Detailed Study of Photometric Redshifts for GOODS-South Galaxies". In: *The Astrophysical Journal* 724. SAO/NASA: 2010ApJ...724..425D, 425–447. DOI: 10.1088/0004-637X/724/1/425 (see p. 165)

Davis, M. et al. (1992). "The end of cold dark matter?" In: *Nature* 356, 489–494. DOI: 10.1038/356489a0 (see pp. 10, 31)

Davis, M. et al. (Feb. 2003). "Science Objectives and Early Results of the DEEP2 Redshift Survey". In: *Proceedings of the SPIE*. Ed. by P. Guhathakurta. Vol. 4834. Dis-

coveries and Research Prospects from 6- to 10-Meter-Class Telescopes II. Waikoloa, Hawai'i, 161–172. DOI: 10.1117/12.457897 (see p. 83)

Dawson, K. S. et al. (Jan. 2013). "The Baryon Oscillation Spectroscopic Survey of SDSS-III". In: *The Astronomical Journal* 145, 10–50. DOI: 10.1088/0004-6256/145/1/10 (see pp. 39, 79, 82, 99)

DePoy, D. L. et al. (July 2008). "The Dark Energy Camera (DECam)". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Ed. by I. S. McLean and M. M. Casali. Vol. 7014. Ground-based and Airborne Instrumentation for Astronomy II. Marseille, France. DOI: 10.1117/12.789466 (see pp. 97, 133)

Descartes, R. (1677). *Le Monde*. Trans. Michael Sean Mahoney, 1978. New York: Abaris Books (see p. 3)

Dicke, R. H. et al. (July 1965). "Cosmic Black-Body Radiation". In: *Astrophysical Journal Letters* 142, 414–419. DOI: 10.1086/148306 (see p. 10)

Diehl, H. T. (2012). private communication (see p. 134)

Doel, P., T. Abbott, et al. (July 2008). "Design and status of the optical corrector for the DES survey instrument". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Ed. by I. S. McLean and M. M. Casali. Vol. 7014. Ground-based and Airborne Instrumentation for Astronomy II. Marseille, France. DOI: 10.1117/12.790049 (see p. 134)

Doel, P., D. Brooks, et al. (Sept. 2012). "Assembly, alignment, and testing of the DECam wide field corrector optics". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Ed. by I. S. McLean, S. K. Ramsay, and H. Takami. Vol. 8446. Ground-based and Airborne Instrumentation for Astronomy IV. Amsterdam, Netherlands. DOI: 10.1117/12.926590 (see p. 134)

Dreyer, J. L. E. (1953). *A History of Astronomy From Thales to Kepler*. 2nd. Dover Publications, Inc. (see p. 3)

Drinkwater, M. J. et al. (Jan. 2010). "The WiggleZ Dark Energy Survey: survey design and first data release". In: *Monthly Notices of the Royal Astronomical Society* 401. SAO/NASA: 2010MNRAS.401.1429D, 1429–1452. DOI: 10.1111/j.1365-2966.2009.15754.x (see p. 83)

Driver, S. P. et al. (May 2011). "Galaxy and Mass Assembly (GAMA): survey diagnostics and core data release". In: *Monthly Notices of the Royal Astronomical Society* 413. SAO/NASA: 2011MNRAS.413..971D, 971–995. DOI: 10.1111/j.1365-2966.2010.18188.x (see p. 83)

Drucker, H., D. Wu, and V. N. Vapnik (Sept. 1999). "Support vector machines for spam categorization". In: *IEEE Transactions on Neural Networks* 10, 1048–1054. DOI: 10.1109/72.788645 (see p. 48)

Eddington, A. S. (1930). "On the Instability of Einstein's Spherical World". In: *Monthly Notices of the Royal Astronomical Society* 90, 668–678 (see p. 6)

Einstein, A. (Nov. 1915). "Die Feldgleichungen der Gravitation". In: *Sitzungsberichte der Preussischen Akademie der Wissenschaften zu Berlin*, 844–847 (see p. 5)

— (1916). "Die Grundlage der allgemeinen Relativitätstheorie". In: *Annalen der Physik* 49, 769–822 (see pp. 4, 5)

— (1917). "Kosmologische Betrachtungen zur allgemeinen Relativitätstheorie". In: *Sitzungsberichte der Königlich Preussische Akademie der Wissenschaften* 6, 142–152. DOI: 10.1002/3527608958.ch10 (see p. 5)

Eisenstein, D. J., J. Annis, et al. (Nov. 2001). "Spectroscopic Target Selection for the Sloan Digital Sky Survey: The Luminous Red Galaxy Sample". In: *The Astronomical Journal* 122. SAO/NASA: 2001AJ....122.2267E, 2267–2280. DOI: 10.1086/323717 (see p. 82)

Eisenstein, D. J., I. Zehavi, et al. (Nov. 2005). "Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies". In: *Astrophysical Journal* 633.2, 560–574. DOI: 10.1086/466512 (see pp. 17, 19)

Eisenstein, D. J. and W. Hu (Apr. 1998). "Baryonic Features in the Matter Transfer Function". In: *The Astrophysical Journal* 496.2, 605–614. DOI: 10.1086/305424 (see p. 16)

Eisenstein, D. J., H.-J. Seo, and M. White (Aug. 2007). "On the Robustness of the Acoustic Scale in the Low-Redshift Clustering of Matter". In: *The Astrophysical Journal* 664. SAO/NASA: 2007ApJ...664..660E, 660–674. DOI: 10.1086/518755 (see p. 18)

Erickson, B. M. S. (2013). "Cyberinfrastructure for Cosmology and Line-of-Sight Projection in Optical Galaxy Clusters". PhD thesis. Ann Arbor, MI: University of Michigan (see p. 3)

Evrard, A. E. (June 1989). "Biased cold dark matter theory - Trouble from rich clusters?" In: *The Astrophysical Journal* 341. SAO/NASA: 1989ApJ...341L..71E, L71–L74. DOI: 10.1086/185460 (see p. 108)

Feldmann, R. et al. (Oct. 2006). "The Zurich Extragalactic Bayesian Redshift Analyzer and its first application: COSMOS". In: *Monthly Notices of the Royal Astronomical Society* 372. SAO/NASA: 2006MNRAS.372..565F, 565–577. DOI: 10.1111/j.1365-2966.2006.10930.x (see p. 44)

Firth, A. E., O. Lahav, and R. S. Somerville (Mar. 2003). "Estimating photometric redshifts with artificial neural networks". In: *Monthly Notices of the Royal Astronomical Society* 339. SAO/NASA: 2003MNRAS.339.1195F, 1195–1202. DOI: 10.1046/j.1365-8711.2003.06271.x (see p. 45)

Fixsen, D. J. et al. (1996). "The Cosmic Microwave Background Spectrum from the Full COBE FIRAS Data Set". In: *Astrophysical Journal* 473, 576–587. DOI: 10.1086/178173 (see p. 11)

Flaugher, B. L. et al. (Sept. 2012). "Status of the Dark Energy Survey Camera (DECam) Project". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Ed. by I. S. McLean, S. K. Ramsay, and H. Takami. Vol. 8446.

Ground-based and Airborne Instrumentation for Astronomy IV. Amsterdam, Netherlands. DOI: 10.1117/12.926216 (see p. 133)

Freeman, P. E. et al. (Oct. 2009). "Photometric redshift estimation using spectral connectivity analysis". In: *Monthly Notices of the Royal Astronomical Society* 398. SAO/NASA: 2009MNRAS.398.2012F, 2012–2021. DOI: 10.1111/j.1365-2966.2009.15236.x (see p. 45)

Frenk, C. S. et al. (Mar. 1990). "Galaxy clusters and the amplitude of primordial fluctuations". In: *The Astrophysical Journal* 351. SAO/NASA: 1990ApJ...351...10F, 10–21. DOI: 10.1086/168439 (see p. 108)

Freund, Y. (Sept. 1995). "Boosting a Weak Learning Algorithm by Majority". In: *Information and Computation* 121, 256–285. DOI: 10.1006/inco.1995.1136 (see p. 47)

— (June 2001). "An Adaptive Version of the Boost by Majority Algorithm". In: *Machine Learning* 43, 293–318. DOI: 10.1023/A:1010852229904 (see p. 57)

— (May 2009). "A more robust boosting algorithm". In: *pre-print*. arXiv:0905.2138 (see p. 57)

Freund, Y. and R. E. Schapire (Aug. 1997). "A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting". In: *Journal of Computer and System Sciences* 55, 119–139. DOI: 10.1006/jcss.1997.1504 (see pp. 56, 57)

— (Sept. 1999). "A Short Introduction to Boosting". In: *Journal of Japanese Society for Artificial Intelligence* 14. Trans. Naoki Abe, 771–780 (see p. 56)

Friedman, J. H. (1999). "Greedy Function Approximation: A Gradient Boosting Machine". In: *The Annals of Statistics* 29, 1189–1232. DOI: 10.1214/aos/1013203451 (see p. 57)

— (Feb. 2002). "Stochastic gradient boosting". In: *Computational Statistics & Data Analysis* 38, 367–378. DOI: 10.1016/S0167-9473(01)00065-2 (see pp. 52, 57)

Friedman, J., T. Hastie, and R. Tibshirani (Aug. 1998). "Additive Logistic Regression: a Statistical View of Boosting". In: *The Annals of Statistics* 28, 2000–2044. DOI: 10.1.1.51.9525 (see p. 57)

— (2000). "Additive logistic regression: a statistical view of boosting". In: *The Annals of Statistics* 28, 337–407. DOI: 10.1214/aos/1016218223 (see p. 52)

Friedmann, A. (1922). "Über die Krümmung des Raumes". In: *Zeitschrift für Physik* 10. Translated in Friedmann 1999a, 377–386. DOI: 10.1007/BF01332580 (see pp. 5, 186)

— (1924). "Über die Möglichkeit einer Welt mit konstanter negativer Krümmung des Raumes". In: *Zeitschrift für Physik* 21. Translated in Friedmann 1999b, 326–332. DOI: 10.1007/BF01328280 (see pp. 5, 186)

— (1999a). "On the curvature of space". In: *General Relativity and Gravitation* 31. Translation of Friedmann 1922, 1991–2000. DOI: 10.1023/A:1026751225741 (see p. 186)

— (1999b). "On the Possibility of a World with Constant Negative Curvature of Space". In: *General Relativity and Gravitation* 31. Translation of Friedmann 1924, 2001–2008. DOI: 10.1023/A:1026755309811 (see p. 186)

Gamow, G. (1970). *My World Line: An Informal Autobiography*. New York: Viking Press (see p. 5)

Gerdes, D. W. et al. (June 2010). "ArborZ: Photometric Redshifts Using Boosted Decision Trees". In: *The Astrophysical Journal* 715, 823–832. DOI: 10.1088/0004-637X/715/2/823 (see pp. 45, 64, 65)

Gladders, M. D. et al. (July 1998). "The Slope of the Cluster Elliptical Red Sequence: A Probe of Cluster Evolution". In: *The Astrophysical Journal* 501. SAO/NASA: 1998ApJ...501..571G, 571. DOI: 10.1086/305858 (see pp. 111, 112)

Gladders, M. D. and H. K. C. Yee (Oct. 2000). "A New Method For Galaxy Cluster Detection. I. The Algorithm". In: *The Astronomical Journal* 120. SAO/NASA: 2000AJ....120.2148G, 2148–2162. DOI: 10.1086/301557 (see pp. 109, 111)

Gursky, H. et al. (Aug. 1971). "A Strong X-Ray Source in the Coma Cluster Observed by UHURU". In: *The Astrophysical Journal* 167. SAO/NASA: 1971ApJ...167L..81G, L81–L84. DOI: 10.1086/180765 (see p. 109)

Guth, A. H. (1981). "Inflationary universe: A possible solution to the horizon and flatness problems". In: *Physical Review D* 23, 347–356. DOI: 10.1103/PhysRevD.23.347 (see p. 12)

Guzzo, L. et al. (Mar. 2013). "The VIMOS Public Extragalactic Redshift Survey (VIPERS). An unprecedented view of galaxies and large-scale structure at 0.5¡z¡1.2". In: *pre-print*. arXiv:1303.2623 (see pp. 83, 99)

Hansen, B. M. S. et al. (2002). "The White Dwarf Cooling Sequence of the Globular Cluster Messier 4". In: *Astrophysical Journal Letters* 574, L155–L158. DOI: 10.1086/342528 (see p. 13)

Hao, J. et al. (Dec. 2010). "A GMBCG Galaxy Cluster Catalog of 55,424 Rich Clusters from SDSS DR7". In: *The Astrophysical Journal Supplement Series* 191. SAO/NASA: 2010ApJS..191..254H, 254–274. DOI: 10.1088/0067-0049/191/2/254 (see p. 109)

Harper, D. R. (2013). "Search for Physics beyond the Standard Model in Events with a Z Boson and Large Missing Transverse Momentum with the ATLAS Detector". PhD thesis. Ann Arbor, MI: University of Michigan (see p. 10)

Hashemi, K. S. (2013a). *BCAM User Manual*. URL: http://alignment.hep.brandeis.edu/Devices/BCAM/User_Manual.html (visited on 06/21/2013) (see pp. 136, 151)

Hashemi, K. S. (2013b). *Long-Wire Data Acquisition Specification*. URL: http://alignment.hep.brandeis.edu/Electronics/LWDAQ/LWDAQ.html (visited on 06/21/2013) (see p. 137)

Hashemi, K. S. and J. R. Bensinger (Dec. 2000). *The BCAM Camera*. ATLAS Note, ATL-MUON-2000-024 (see p. 135)

Hastie, T., R. Tibshirani, and J. Friedman (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. 2nd. Springer Series in Statistics. Springer-Verlag (see pp. 47, 57)

Heath, T. (1913). *Aristarchus of Samos*. Oxford: Oxford University Press (see p. 3)

Herschel, W. (Jan. 1800). "Experiments on the Refrangibility of the Invisible Rays of the Sun". In: *Philosophical Transactions of the Royal Society of London* 90, 284–292. DOI: 10.1098/rstl.1800.0015 (see p. 4)

Hetherington, N. S. (2006). *Planetary Motions*. Ed. by B. Baigrie. Westport, Connecticut: Greenwood Press (see p. 2)

Hilton, M. et al. (May 2009). "The XMM Cluster Survey: Galaxy Morphologies and the Color-Magnitude Relation in XMMXCS J2215.9 - 1738 at z = 1.46". In: *The Astrophysical Journal* 697. SAO/NASA: 2009ApJ...697..436H, 436–451. DOI: 10.1088/0004-637X/697/1/436 (see p. 109)

Hoecker, A. et al. (2007). "TMVA - Toolkit for Multivariate Data Analysis". In: arXiv:physics/0703039 (see p. 51)

Hoekstra, H. and B. Jain (May 2008). "Weak Gravitational Lensing and its Cosmological Applications". In: *pre*. arXiv:0805.0139 (see p. 33)

Hogg, D. W. (Dec. 2000). "Distance Measures in Cosmology". In: *pre-print*. arXiv:astro-ph/9905116 (see p. 33)

Holden, E. S. (1881). *Sir William Herschel*. New York: Charles Scribner's Sons (see p. 4)

Honscheid, K. et al. (July 2008). "The read-out and control system of the DES camera (SISPI)". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Ed. by A. Bridger and N. M. Radziwill. Vol. 7019. Advanced Software and Control for Astronomy II. DOI: 10.1117/12.787926 (see p. 150)

Hotelling, H. (1933). "Analysis of a complex of statistical variables into principal components". In: *Journal of Educational Psychology* 24, 417–441. DOI: 10.1037/h0071325 (see p. 62)

Howe, N. R., T. M. Rath, and R. Manmatha (Aug. 2005). "Boosted Decision Trees for Word Recognition in Handwritten Document Retrieval". In: *Proceedings of the 28th International ACM SIGIR Conference on Research and Development in Information Retrieval*. New York: ACM, 377–383. DOI: 10.1.1.152.1551 (see p. 48)

Hubble, E. P. (Dec. 1926). "Extra-galactic Nebulae". In: *Astrophysical Journal* 64, 321–369. DOI: 10.1086/143018 (see p. 6)

— (Jan. 1929). "A relation between distance and radial velocity among extra-galactic nebulae". In: *Proceedings of the National Academy of Sciences* 15, 168–173. DOI: 10.1073/pnas.15.3.168 (see pp. 6, 7)

Huterer, D., A. Kim, et al. (Nov. 2004). "Redshift Accuracy Requirements for Future Supernova and Number Count Surveys". In: *The Astrophysical Journal* 615. SAO/NASA: 2004ApJ...615..595H, 595–602. DOI: 10.1086/424726 (see pp. 41, 119)

Huterer, D., M. Takada, et al. (Feb. 2006). "Systematic errors in future weak-lensing surveys: requirements and prospects for self-calibration". In: *Monthly Notices of the Royal Astronomical Society* 366. SAO/NASA: 2006MNRAS.366..101H, 101–114. DOI: 10.1111/j.1365-2966.2005.09782.x (see p. 41)

Huterer, D. (Sept. 2010). "Class notes from Physics 525". University of Michigan (see p. 33)

Huterer, D. and M. S. Turner (Oct. 1999). "Prospects for probing the dark energy via supernova distance measurements". In: *Physical Review D* 60. SAO/NASA: 1999PhRvD..60h1301H, 081301. DOI: 10.1103/PhysRevD.60.081301 (see p. 14)

Ilbert, O. et al. (Oct. 2006). "Accurate photometric redshifts for the CFHT legacy survey calibrated using the VIMOS VLT deep survey". In: *Astronomy & Astrophysics* 457. SAO/NASA: 2006A&A...457..841I, 841–856. DOI: 10.1051/0004-6361:20065138 (see p. 44)

Ivezić, Ž. et al. (June 2008). "LSST: from Science Drivers to Reference Design and Anticipated Data Products". In: *pre-print*. arXiv:0805.2366 (see p. 40)

Jenkins, A. et al. (Feb. 2001). "The mass function of dark matter haloes". In: *Monthly Notices of the Royal Astronomical Society* 321. SAO/NASA: 2001MNRAS.321..372J, 372–384. DOI: 10.1046/j.1365-8711.2001.04029.x (see p. 108)

Kearns, M. and L. Valiant (Jan. 1994). "Cryptographic limitations on learning Boolean formulae and finite automata". In: *Journal of the ACM* 41, 67–95. DOI: 10.1145/174644.174647 (see p. 56)

Kellogg, E. et al. (Apr. 1971). "X-Ray Observations of the Virgo Cluster, NGC 5128, and 3c 273 from the UHURU Satellite". In: *The Astrophysical Journal* 165. SAO/NASA: 1971ApJ...165L..49K, L49–L54. DOI: 10.1086/180714 (see p. 109)

Kind, M. C. and R. J. Brunner (Mar. 2013). "TPZ: Photometric redshift PDFs and ancillary information by using prediction trees and random forests". In: *pre-print*. arXiv:1303.7269 (see pp. 45, 46)

Kochanek, C. S. et al. (May 2012). "AGES: The AGN and Galaxy Evolution Survey". In: *The Astrophysical Journal Supplement Series* 200. SAO/NASA: 2012ApJS..200....8K. DOI: 10.1088/0067-0049/200/1/8 (see p. 82)

Koester, B. P. (2006). "MaxBCG: Systematic discovery, characterization and calibration of galaxy clusters from large optical surveys". SAO/NASA: 2006PhDT.....17K. Ph.D. Thesis. Ann Arbor, MI: University of Michigan (see pp. 109, 111)

Kolb, E. et al. (June 2006). *Report of the Dark Energy Task Force*. Tech. rep. http://www.nsf.gov/mps/ast/detf.jsp. Dark Energy Task Force (see p. 31)

Koo, D. C.-Y. (1981). "Multi-Color Analysis of Galaxy Evolution and Cosmology". SAO/NASA: 1981PhDT.....66K. Ph.D. Thesis. Berkeley, CA 94720: University of California, Berkeley (see p. 40)

— (Mar. 1985). "Optical multicolors - A poor person's Z machine for galaxies". In: *The Astronomical Journal* 90. SAO/NASA: 1985AJ.....90..418K, 418–440. DOI: 10.1086/113748 (see p. 40)

Kowal, C. T. (1968). "Absolute magnitudes of supernovae". In: *Astronomical Journal* 73, 1021–1024. DOI: 10.1086/110763 (see p. 14)

Kuhn, T. S. (1957). *The Copernican Revolution*. Cambridge, MA: Harvard University Press (see p. 3)

Le Fèvre, O., D. Crampton, et al. (Dec. 1995). "The Canada-France Redshift Survey. II. Spectroscopic Program: Data for the 0000-00 and 1000+25 Fields". In: *The Astrophysical Journal* 455. SAO/NASA: 1995ApJ...455...60L, 60–74. DOI: 10.1086/176556 (see p. 82)

Le Fèvre, O., G. Vettolani, et al. (Sept. 2005). "The VIMOS VLT deep survey. First epoch VVDS-deep survey: 11 564 spectra with $17.5 \leq \mathrm{IAB} \leq 24$, and the redshift distribution over $0 \leq z\ le5$". In: *Astronomy & Astrophysics* 439. SAO/NASA: 2005A&A...439..845L, 845–862. DOI: 10.1051/0004-6361:20041960 (see pp. 83, 98, 99)

Lemaître, G. (Apr. 1927). "Un Univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques". In: *Annales de la Societe Scientifique de Bruxelles* A47, 49–59 (see p. 6)

— (1931a). "A homogeneous universe of constant mass and increasing radius accounting for the radial velocity of extra-galactic nebulae". In: *Monthly Notices of*

the Royal Astronomical Society 91. Translation of Lemaître 1927, 483–490 (see p. 6)

Lemaître, G. (May 1931b). "The Beginning of the World from the Point of View of Quantum Theory". In: *Nature* 127, 706. DOI: 10.1038/127706b0 (see p. 6)

Levi, M. et al. (Aug. 2013). "The DESI Experiment, a whitepaper for Snowmass 2013". In: *pre-print*. arXiv:1308.0847 (see pp. 39, 131)

Lilly, S. J., V. Le Brun, et al. (Oct. 2009). "The zCOSMOS 10k-Bright Spectroscopic Sample". In: *The Astrophysical Journal Supplement Series* 184. SAO/NASA: 2009ApJS..184..218L, 218–229. DOI: 10.1088/0067-0049/184/2/218 (see p. 83)

Lilly, S. J., O. Le Fèvre, D. Crampton, et al. (Dec. 1995). "The Canada-France Redshift Survey. I. Introduction to the Survey, Photometric Catalogs, and Surface Brightness Selection Effects". In: *The Astrophysical Journal* 455. 1995ApJ...455...50L, 50–59. DOI: 10.1086/176555 (see p. 82)

Lilly, S. J., O. Le Fèvre, A. Renzini, et al. (Sept. 2007). "zCOSMOS: A Large VLT/VIMOS Redshift Survey Covering 0 ¡ z ¡ 3 in the COSMOS Field". In: *The Astrophysical Journal Supplement Series* 172. SAO/NASA: 2007ApJS..172...70L, 70–85. DOI: 10.1086/516589 (see pp. 83, 99)

Lima, M., C. E. Cunha, et al. (Oct. 2008). "Estimating the redshift distribution of photometric galaxy samples". In: *Monthly Notices of the Royal Astronomical Society* 390. SAO/NASA: 2008MNRAS.390..118L, 118–130. DOI: 10.1111/j.1365-2966.2008.13510.x (see pp. 40, 45, 46, 78)

Lima, M. and W. Hu (Dec. 2007). "Photometric redshift requirements for self-calibration of cluster dark energy studies". In: *Physical Review D* 76. SAO/NASA: 2007PhRvD..76l3013L, 123013. DOI: 10.1103/PhysRevD.76.123013 (see p. 41)

Lin, H. (June 2013). *Photometric "re-calibration" procedure*. https://cdcvs.fnal.gov/redmine/projects/des-photoz/wiki/Y1C1_Y1C2_SVA1_data. Internal DES document (see p. 97)

Lin, H. et al. (Dec. 2006). "Improving Galaxy Cluster Photometric Redshifts". In: *American Astronomical Society Meeting Abstracts*. Vol. 38. Bulletin of the American Astronomical Society. SAO/NASA: 2006AAS...20921503L (see p. 111)

Linde, A. D. (Feb. 1982). "A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems". In: *Physics Letters B* 108. SAO/NASA: 1982PhLB..108..389L, 389–393. DOI: 10.1016/0370-2693(82)91219-9 (see p. 12)

Linder, E. V. (Mar. 2003). "Exploring the Expansion History of the Universe". In: *Physical Review Letters* 90. SAO/NASA: 2003PhRvL..90i1301L, 091301. DOI: 10.1103/PhysRevLett.90.091301 (see p. 29)

Loh, E. D. and E. J. Spillar (Apr. 1986a). "Photometric redshifts of galaxies". In: *The Astrophysical Journal* 303. SAO/NASA: 1986ApJ...303..154L, 154–161. DOI: 10.1086/164062 (see p. 40)

— (Aug. 1986b). "A measurement of the mass density of the universe". In: *The Astrophysical Journal Letters* L1–L4. SAO/NASA: 1986ApJ...307L...1L, 1986. DOI: 10.1086/184717 (see p. 40)

López-Cruz, O., W. A. Barkhouse, and H. K. Yee (Oct. 2004). "The Color-Magnitude Effect in Early-Type Cluster Galaxies". In: *The Astrophysical Journal* 614. SAO/NASA: 2004ApJ...614..679L, 679–691. DOI: 10.1086/423664 (see p. 109)

Ma, Z., W. Hu, and D. Huterer (Jan. 2006). "Effects of Photometric Redshift Uncertainties on Weak-Lensing Tomography". In: *The Astrophysical Journal* 636. SAO/NASA: 2006ApJ...636...21M, 21–29. DOI: 10.1086/497068 (see pp. 41, 43, 44, 119)

Mana, A. et al. (Sept. 2013). "Combining clustering and abundances of galaxy clusters to test cosmology and primordial non-Gaussianity". In: *Monthly Notices of the Royal Astronomical Society* 434. SAO/NASA: 2013MNRAS.434..684M, 684–695. DOI: 10.1093/mnras/stt1062 (see p. 108)

Mandelbaum, R. et al. (May 2008). "Precision photometric redshift calibration for galaxy-galaxy weak lensing". In: *Monthly Notices of the Royal Astronomical Society* 386. SAO/NASA: 2008MNRAS.386..781M, 781–806. DOI: 10.1111/j.1365-2966.2008.12947.x (see p. 47)

Meadows, J. (Jan. 1978). "Astronomy of Stonehenge". In: *Nature* 271, 189–190. DOI: 10.1038/271189a0 (see p. 2)

Mehrtens, N. et al. (June 2012). "The XMM Cluster Survey: optical analysis methodology and the first data release". In: *Monthly Notices of the Royal Astronomical Society* 423. SAO/NASA: 2012MNRAS.423.1024M, 1024–1052. DOI: 10.1111/j.1365-2966.2012.20931.x (see p. 121)

Mignoli, M. et al. (July 2005). "The K20 survey. VII. The spectroscopic catalogue: Spectral properties and evolution of the galaxy population". In: *Astronomy & Astrophysics* 437. SAO/NASA: 2005A&A...437..883M, 883–897. DOI: 10.1051/0004-6361:20042434 (see p. 99)

Milgrom, M. (1983a). "A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis". In: *Astrophysical Journal* 270, 365–370. DOI: 10.1086/161130 (see pp. 9, 26)

— (1983b). "A modification of the Newtonian dynamics - Implications for galaxies". In: *Astrophysical Journal* 270, 371–389. DOI: 10.1086/161131 (see p. 9)

Miller et al. (in prep). In: (see p. 121)

Miller, C. J. et al. (Sept. 2005). "The C4 Clustering Algorithm: Clusters of Galaxies in the Sloan Digital Sky Survey". In: *The Astronomical Journal* 130. SAO/NASA: 2005AJ....130..968M, 968–1001. DOI: 10.1086/431357 (see p. 109)

Mitchell, T. M. (1997). *Machine Learning*. Ed. by C. L. Liu and A. B. Tucker. Computer Science. New York: McGraw-Hill (see p. 45)

Mohr, J. J. et al. (Sept. 2012). "The Dark Energy Survey data processing and calibration system". In: *Society of Photo-Optical Instrumentation Engineers (SPIE)*

Conference Series. Ed. by N. M. Radziwill and G. Chiozzi. Vol. 8451. Software and Cyberinfrastructure for Astronomy II. arXiv:1207.3189. DOI: 10.1117/12.926785 (see p. 97)

Mullaney, J. (2007). *The Herschel Objects and How to Observe Them*. Ed. by M. Inglis. New York: Springer Science+Business Media, LLC (see p. 4)

Muñoz, F. et al. (Sept. 2012). "Dark Energy Camera installation at CTIO: technical challenges". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Ed. by L. M. Stepp, R. Gilmozzi, and H. J. Hall. Vol. 8444. Ground-based and Airborne Telescopes IV. Amsterdam, Netherlands. DOI: 10.1117/12.925752 (see p. 134)

Newman, J. A. et al. (Sept. 2012). "The DEEP2 Galaxy Redshift Survey: Design, Observations, Data Reduction, and Redshifts". In: *The Astrophysical Journal Supplement Series* 208. SAO/NASA: 2013ApJS..208....5N. DOI: 10.1088/0067-0049/208/1/5 (see p. 83)

Newton, I. (1687). *Philosophiae Naturalis Principia Mathematica*. Vol. 1. Trans. Andrew Motte, 1729. London: Benjamin Motte (see p. 3)

Nobel Media AB (2006). *Press Release: The 2006 Nobel Prize in Physics*. URL: http://www.nobelprize.org/nobel_prizes/physics/laureates/2006/press.html (visited on 08/09/2013) (see p. 11)

Oort, J. H. (1932). "The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems". In: *Bulletin of the Astronomical Institutes of the Netherlands* 6, 249–287 (see p. 7)

Ostriker, J. P. and P. J. E. Peebles (1973). "A Numerical Study of the Stability of Flattened Galaxies: or, can Cold Galaxies Survive?" In: *Astrophysical Journal* 186, 467–480. DOI: 10.1086/152513 (see p. 9)

Oyaizu, H. et al. (Feb. 2008). "A Galaxy Photometric Redshift Catalog for the Sloan Digital Sky Survey Data Release 6". In: *The Astrophysical Journal* 674. SAO/NASA: 2008ApJ...674..768O, 768–783. DOI: 10.1086/523666 (see p. 40)

Pannekoek, A. (1961). *A History of Astronomy*. Mineola, New York: Dover Publications, Inc. (see p. 3)

Papovich, C. et al. (June 2010). "A Spitzer-selected Galaxy Cluster at z = 1.62". In: *The Astrophysical Journal* 716. SAO/NASA: 2010ApJ...716.1503P, 1503–1513. DOI: 10.1088/0004-637X/716/2/1503 (see p. 109)

Pearson, K. (1901). "On Lines and Planes of Closest Fit to System of Points in Space". In: *Philosophical Magazine* 2, 559–572 (see p. 62)

Penzias, A. A. and R. W. Wilson (July 1965). "A Measurement of Excess Antenna Temperature at 4080 Mc/s". In: *Astrophysical Journal Letters* 142, 419–421. DOI: 10.1086/148307 (see p. 10)

Percival, W. J. et al. (Nov. 2007). "Measuring the Baryon Acoustic Oscillation scale using the Sloan Digital Sky Survey and 2dF Galaxy Redshift Survey". In: *Monthly Notices of the Royal Astronomical Society* 381. SAO/NASA: 2007MNRAS.381.1053P, 1053–1066. DOI: 10.1111/j.1365-2966.2007.12268.x (see p. 19)

Perlmutter, S. (Apr. 2003). "Supernovae, Dark Energy, and the Accelerating Universe". In: *Physics Today* 56.4, 53. DOI: 10.1063/1.1580050 (see p. 14)

Perlmutter, S. et al. (1999). "Measurements of Omega and Lambda from 42 High-Redshift Supernovae". In: *Astrophysical Journal* 517, 565–586. DOI: 10.1086/307221 (see pp. 14, 33)

Press, W. H. and P. Schechter (Feb. 1974). "Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation". In: *The Astrophysical Journal* 187. SAO/NASA: 1974ApJ...187..425P, 425–438. DOI: 10.1086/152650 (see p. 108)

Pskovskii, Y. P. (1967). "The Photometric Properties of Supernovae". In: *Soviet Astronomy* 11, 63 (see p. 14)

Reif, F. (1965). *Fundamentals of Statistical and Thermal Physics*. Ed. by E. U. Condon. New York: McGraw-Hill (see p. 29)

Riess, A. G. et al. (1998). "Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant". In: *Astronomical Journal* 116, 1009–1038. DOI: 10.1086/300499 (see pp. 14, 33)

Roe, B. P. et al. (May 2005). "Boosted decision trees as an alternative to artificial neural networks for particle identification". In: *Nuclear Instruments and Methods in Physics Research Section A* 543, 577–584. DOI: 10.1016/j.nima.2004.12.018 (see pp. 48, 51)

Rogers, J. H. (Feb. 1998). "Origins of the ancient constellations: I. The Mesopotamian traditions". In: *Journal of the British Astronomical Association* 108.1, 9–28 (see p. 2)

Roodman, A. (Sept. 2012). "Focus and Alignment of the Dark Energy Camera Using Out-Of-Focus Stars". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Ed. by I. S. McLean, S. K. Ramsay, and H. Takami. Vol. 8446. Ground-based and Airborne Instrumentation for Astronomy IV. Amsterdam, Netherlands. DOI: 10.1117/12.926671 (see p. 135)

Rubin, V. C., D. Burstein, et al. (1985). "Rotation velocities of 16 Sa galaxies and a comparison of Sa, Sb, and SC rotation properties". In: *Astrophysical Journal* 289, 81–104. DOI: 10.1086/162866 (see p. 8)

Rubin, V. C. and J. W. Kent Ford (1970). "Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions". In: *Astrophysical Journal* 159, 379–403. DOI: 10.1086/150317 (see pp. 8, 9)

Rudnick, G., M. Franx, et al. (Nov. 2001). "A K-band-Selected Photometric Redshift Catalog in the Hubble Deep Field South: Sampling the Rest-Frame V Band to z

=3". In: *The Astronomical Journal* 122. SAO/NASA: 2001AJ....122.2205R, 2205–2221. DOI: 10.1086/323542 (see p. 44)

Rudnick, G., H.-W. Rix, et al. (Dec. 2003). "The Rest-Frame Optical Luminosity Density, Color, and Stellar Mass Density of the Universe from z = 0 to z = 3". In: *The Astrophysical Journal* 599. SAO/NASA: 2003ApJ...599..847R, 847–864. DOI: 10.1086/379628 (see p. 44)

Ryden, B. S. (2003). *Introduction to Cosmology*. San Francisco, CA: Addison-Wesley (see p. 5)

Rykoff, E. S. et al. (Mar. 2013). "redMaPPer I: Algorithm and SDSS DR8 Catalog". In: *pre-print*. arXiv:1303.3562 (see pp. 109, 111, 126)

Sandage, A. and N. Visvanathan (Aug. 1978). "The color-absolute magnitude relation for E and S0 galaxies. II - New colors, magnitudes, and types for 405 galaxies". In: *The Astrophysical Journal* 223. SAO/NASA: 1978ApJ...223..707S, 707–729. DOI: 10.1086/156305 (see p. 109)

Schapire, R. E. (1990). "The Strength of Weak Learnability". In: *Machine Learning* 5, 197–227. DOI: 10.1.1.20.723 (see pp. 47, 56)

Schlafly, E. F. and D. P. Finkbeiner (Aug. 2011). "Measuring Reddening with Sloan Digital Sky Survey Stellar Spectra and Recalibrating SFD". In: *The Astrophysical Journal* 737. SAO/NASA: 2011ApJ...737..103S. DOI: 10.1088/0004-637X/737/2/103 (see p. 97)

Schlegel, D. J., D. P. Finkbeiner, and M. Davis (June 1998). "Maps of Dust Infrared Emission for Use in Estimation of Reddening and Cosmic Microwave Background Radiation Foregrounds". In: *The Astrophysical Journal* 500, 525–553 (see pp. 81, 97)

Schneider, D. P. et al. (June 2010). "The Sloan Digital Sky Survey Quasar Catalog. V. Seventh Data Release". In: *The Astronomical Journal* 139. SAO/NASA:

2010AJ....139.2360S, 2360–2373. DOI: 10.1088/0004-6256/139/6/2360 (see p. 165)

Shackelford, J. F. (2001). *Materials Science and Engineering Handbook*. Ed. by J. F. Shackelford and W. Alexander. 3rd ed. Boca Raton, FL: CRC Press (see p. 145)

Sheldon, E. S. et al. (2012). "Photometric Redshift Probability Distributions for Galaxies in the SDSS DR8". In: *The Astrophysical Journal Supplement Series* 201, 32. DOI: 10.1088/0067-0049/201/2/32 (see p. 80)

Smoot, G. F. et al. (Sept. 1992). "Structure in the COBE differential microwave radiometer first-year maps". In: *Astrophysical Journal Letters* 396, L1–L5. DOI: 10.1086/186504 (see p. 11)

Song, J. et al. (Dec. 2012). "Redshifts, Sample Purity, and BCG Positions for the Galaxy Cluster Catalog from the First 720 Square Degrees of the South Pole Telescope Survey". In: *The Astrophysical Journal* 761. SAO/NASA: 2012ApJ...761...22S. DOI: 10.1088/0004-637X/761/1/22 (see p. 109)

Spence, K. (Nov. 2000). "Ancient Egyptian chronology and the astronomical orientation of pyramids". In: *Nature* 408, 320–324. DOI: 10.1038/35042510 (see p. 1)

Spergel, D. N. et al. (2003). "First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters". In: *Astrophysical Journal Supplement Series* 148, 175–194. DOI: 10.1086/377226 (see p. 12)

Springel, V. (2004). *The largest N-body simulation of the universe*. URL: http://www.mpa-garching.mpg.de/mpa/research/current_research/hl2004-8/hl2004-8-en.html (visited on 01/28/2014) (see p. 20)

Stefanik, A. (Dec. 2012). private communication (see p. 156)

Stott, J. P. et al. (Apr. 2009). "The evolution of the red sequence slope in massive galaxy clusters". In: *Monthly Notices of the Royal Astronomical Society* 394.

SAO/NASA: 2009MNRAS.394.2098S, 2098–2108. DOI: 10.1111/j.1365-2966.2009.14477.x (see p. 111)

Strauss, M. A. et al. (Sept. 2002). "Spectroscopic Target Selection in the Sloan Digital Sky Survey: The Main Galaxy Sample". In: *The Astronomical Journal* 124. SAO/NASA: 2002AJ....124.1810S, 1810–1824. DOI: 10.1086/342343 (see p. 82)

Sunyaev, R. A. and Y. B. Zeldovich (Nov. 1972). "The Observations of Relic Radiation as a Test of the Nature of X-Ray Radiation from the Clusters of Galaxies". In: *Comments on Astrophysics and Space Physics* 4. SAO/NASA: 1972CoASP...4..173S, 173–178 (see p. 109)

Sutherland, W. and W. Saunders (Dec. 1992). "On the likelihood ratio for source identification". In: *Monthly Notices of the Royal Astronomical Society* 259. SAO/NASA: 1992MNRAS.259..413S, 413–420 (see p. 165)

Sypniewski, A. J. and D. W. Gerdes (2011). "Estimating Photometric Redshifts in Non-Representative Galaxy Samples using Boosted Decisions Trees". In: *The 217th Meeting of the American Astronomical Society*. Large Scale Structure, Cosmic Distance Scale. Seattle, WA (see p. 84)

Tarlè, G. (2013). private communication (see p. 39)

The Dark Energy Survey (2007). *The Dark Energy Survey Science Program*. URL: https://www.darkenergysurvey.org/reports/proposal-standalone.pdf (visited on 02/01/2014) (see p. 32)

The New American Bible (1987). Iowa City, Iowa: World Publishing, Inc. (see p. 2)

Thompson, J. E. S. (May 1974). "Maya Astronomy". In: *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*. Vol. 276. The Place of Astronomy in the Ancient World 1257. The Royal Society, 83–98. DOI: 10.1098/rsta.1974.0011 (see p. 2)

Tinker, J. et al. (Dec. 2008). "Toward a Halo Mass Function for Precision Cosmology: The Limits of Universality". In: *The Astrophysical Journal* 688. SAO/NASA: 2008ApJ...688..709T, 709–728. DOI: 10.1086/591439 (see p. 108)

Türler, M. (Jan. 2006). "INTEGRAL reveals Milky Ways' supernova rate". In: *CERN Courier* 46.1 (see p. 14)

Vikhlinin, A., R. A. Burenin, et al. (Feb. 2009). "Chandra Cluster Cosmology Project. II. Samples and X-Ray Data Reduction". In: *The Astrophysical Journal* 692. SAO/NASA: 2009ApJ...692.1033V, 1033–1059. DOI: 10.1088/0004-637X/692/2/1033 (see p. 109)

Vikhlinin, A., A. V. Kravtsov, et al. (Feb. 2009). "Chandra Cluster Cosmology Project III: Cosmological Parameter Constraints". In: *The Astrophysical Journal* 692. SAO/NASA: 2009ApJ...692.1060V, 1060–1074. DOI: 10.1088/0004-637X/692/2/1060 (see pp. 109, 110)

Wadadekar, Y. (Jan. 2005). "Estimating Photometric Redshifts Using Support Vector Machines". In: *The Publications of the Astronomical Society of the Pacific* 117. SAO/NASA: 2005PASP..117...79W, 79–85. DOI: 10.1086/427710 (see p. 45)

Wang, L. and M. Rowan-Robinson (Sept. 2009). "The Imperial IRAS-FSC Redshift Catalogue". In: *Monthly Notices of the Royal Astronomical Society* 398. SAO/NASA: 2009MNRAS.398..109W, 109–118. DOI: 10.1111/j.1365-2966.2009.15138.x (see p. 165)

Way, M. J. et al. (Nov. 2009). "New Approaches to Photometric Redshift Prediction Via Gaussian Process Regression in the Sloan Digital Sky Survey". In: *The Astrophysical Journal* 706. SAO/NASA: 2009ApJ...706..623W, 623–636. DOI: 10.1088/0004-637X/706/1/623 (see p. 45)

Weinberg, D. H. et al. (Jan. 2012). "Observational Probes of Cosmic Acceleration". In: *pre*. arXiv:1201.2434 (see p. 16)

White, S. D. M. et al. (1993). "The baryon content of galaxy clusters: a challenge to cosmological orthodoxy". In: *Nature* 366, 429–433. DOI: 10.1038/366429a0 (see pp. 10, 31)

Wirth, G. D. et al. (June 2004). "The Team Keck Treasury Redshift Survey of the GOODS-North Field". In: *The Astronomical Journal* 127. SAO/NASA: 2004AJ....127.3121W, 3121–3136. DOI: 10.1086/420999 (see p. 83)

Wolstencroft, R. D. et al. (Nov. 1986). "The Identification of IRAS Point Sources - Part One - a 304-DEGREE Field Centred on the South Galactic Pole". In: *Monthly Notices of the Royal Astronomical Society* 223. SAO/NASA: 1986MNRAS.223..279W, 279–302 (see p. 165)

Yee, H. K. C. et al. (Aug. 2000). "The CNOC2 Field Galaxy Redshift Survey. I. The Survey and the Catalog for the Patch CNOC 0223+00". In: *The Astrophysical Journal Supplement Series* 129.2, 475. DOI: 10.1086/313426 (see p. 82)

York, D. G. et al. (Sept. 2000). "The Sloan Digital Sky Survey: Technical Summary". In: *The Astronomical Journal* 120, 1579–1587. DOI: 10.1086/301513 (see pp. 39, 79)

Zeldovich, Y. B. and R. A. Sunyaev (July 1969). "The Interaction of Matter and Radiation in a Hot-Model Universe". In: *Astrophysics and Space Science* 4. SAO/NASA: 1969Ap&SS...4..301Z, 301–316. DOI: 10.1007/BF00661821 (see p. 109)

Zhang, Y. et al. (May 2013). "Estimating Photometric Redshifts of Quasars via K-nearest Neighbor Approach Based on Large Survey Databases". In: *pre-print*. arXiv:1305.5023 (see p. 45)

Zlatev, I., L. Wang, and P. J. Steinhardt (1999). "Quintessence, Cosmic Coincidence, and the Cosmological Constant". In: *Physical Review Letters* 82, 896–899. DOI: 10.1103/PhysRevLett.82.896 (see p. 15)

Zwicky, F. (1933). "Die Rotverschiebung von extragalaktischen Nebeln". In: *Helvetica Physica Acta* 6, 110–127 (see p. 8)

Zwicky, F. (1937). "On the Masses of Nebulae and of Clusters of Nebulae". In: *Astrophysical Journal* 86, 217–246. DOI: 10.1086/143864 (see p. 8)