

Establishing a Multi-year Monitoring Campaign of
Low-luminosity Active Galactic Nuclei

Benjamin D. Boizelle

A senior thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of
Bachelor of Science

J. Ward Moody, Advisor

Department of Physics and Astronomy

Brigham Young University

August 2012

Copyright © 2012 Benjamin D. Boizelle

All Rights Reserved

ABSTRACT

Establishing a Multi-year Monitoring Campaign of Low-luminosity Active Galactic Nuclei

Benjamin D. Boizelle
Department of Physics and Astronomy
Bachelor of Science

Optical variability of low-luminosity Active Galactic Nuclei (LLAGN) is explored with a large subsample from the Palomar survey of nearby galaxies. Emission line ratios show that nearly half of these galaxies qualify as active, but their other properties are not usually well known. To test the activity of LLAGN, the ROVOR research group is conducting a multi-year, multi-filter monitoring campaign of 100 LLAGN and AGN. In preparation for analyzing the first few months of data, this paper explores different photometric techniques and identifies all relevant sources of error affecting the images.

Keywords: LLAGN, blazars, image subtraction, photometric accuracy, error analysis

ACKNOWLEDGMENTS

My gratitude goes to those who supported me while I prepared this project. All of the Astronomy Department professors were patient in answering questions. My gratitude especially goes to Dr. Moody, for his excellent guidance these past two years. Classmates and friends helped to keep me sane, and I must thank Valerie Lewis, Kim Bates, Marcus Holden, Nathaly Young, Thayne McCombs, and Holly Mumford for being there when the project seemed doomed. My appreciation also must go to the other members of the ROVOR group for their continued support. My best wishes for their dedication to research and the success of the project.

Contents

Table of Contents	iv
1 Introduction	1
2 Surveys	3
3 AGN Properties	6
3.1 Compactness	6
3.2 Variability	7
3.3 Polarization	8
3.4 Emission Lines	8
3.5 Large Redshifts	9
4 AGN Classifications	11
4.1 Radio-quiet AGN	12
4.1.1 Seyferts	12
4.1.2 LINERs	13
4.1.3 Transition Objects	14
4.2 Radio-loud AGN	15
4.2.1 Radio Galaxies	15
4.2.2 Blazars	16
5 Standard Model Morphology	17
5.1 AGN Evolution	18
6 The LLAGN Project	20
6.1 LINERs and Transition Objects as AGN	20
6.1.1 Con	20
6.1.2 Pro	21
6.2 Palomar Survey	22
6.3 Project Scope	22
6.3.1 Object List	23

7	ROVOR Site	24
8	Data Reduction	26
9	Photometry	28
9.1	Point-source Photometry	29
9.1.1	Fixed-aperture Photometry	29
9.1.2	Variable-aperture Photometry	29
9.1.3	Standard Star vs. Differential Photometry	30
9.2	Extended-object Photometry	31
9.2.1	PSF-matching Image Subtraction	33
9.3	Photometric Error	34
9.3.1	Signal-to-Noise	35
9.4	Initial Results and Error Analysis	36
9.4.1	Aperture Results and Analysis	37
9.4.2	Image Subtraction Results and Analysis	40
10	Conclusions	44
A	AGN Parameters	45
A.1	Explaining the Table	45
A.1.1	Redshift	46
A.1.2	Eddington Luminosity Ratio	47
A.2	AGN Monitoring Targets	49
B	Spectra to Optical Transformation	54
C	Astronomical CCD Technology	56
C.1	Manufacturing	57
C.2	Buried vs. Surface Channels	58
C.3	Photoelectrons	60
C.3.1	Dark Current	61
C.3.2	Cosmic Ray Events	61
C.4	Illumination Orientation	62
C.5	Reading Out the CCD	64
C.5.1	Gain	65
C.5.2	Read Noise	66
C.5.3	Nonlinearity	67
C.6	ROVOR Detector	67

D	CCD Reductions	69
D.1	Noise and its Removal	70
D.2	Bias Frames	71
D.3	Dark Frames	73
D.4	Flat Fields	74
D.4.1	Illumination Correction	75
D.4.2	Quantum Efficiency	76
D.4.3	Flat-Field Stability	77
D.4.4	Obtaining Flat Fields	79
D.5	Combining Images and Master Frames	80
D.5.1	Master Bias	81
D.5.2	Master Dark	83
D.5.3	Master Flat Frame	84
D.5.4	The Processed Object Frame	86
D.6	Data Reduction Afterwards	87
E	Gradients on the Sky	88
E.1	Twilight Sky	88
E.2	Moon-induced Sky Gradients	89
F	All-Sky Photometry	93
F.1	Extinction	93
F.2	Scattering Theory	96
F.3	Extinction with Filters	97
F.4	First-order Extinction	98
F.5	Second-order Extinction	99
F.5.1	Transformation Coefficients	101
F.6	All-sky Photometry at ROVOR	103
G	Using Programs in the Project	105
G.1	Mira Pro	105
G.2	IRAF	107
G.2.1	CL Commands and Reference	107
G.2.2	Flat-Field Surface Fitting	107
G.3	PSF-matching Image Subtraction Programs	110
G.3.1	Method to Determine Errors in PSF-matching Image Subtraction	110
G.3.2	Using HOTPANTS	111
G.3.3	Using Cross-Convolution Image Subtraction	116

Chapter 1

Introduction

Active Galactic Nuclei (AGN) have captivated astronomers and the public for half a century. Their massive energy output, large redshifts, and rapid variability make them one of the most exotic phenomena in the universe. AGN consist of a central engine— a supermassive black hole— accreting matter and releasing more energy at times than their entire host galaxy. Observing these systems at different orientations results in separate features; quasars, radio galaxies, Seyferts, BL Lacs, OVV's, and LINERs are classes of AGN that were first noticed as disparate objects. Connecting features (like compactness) spurred their inclusion into one theory called the Standard Model of AGN. While it has been successful, continued research needs to connect details that do not fit the current model.

Bright AGN have been well studied due to their strong features and visibility at great distances. LLAGN, the low-luminosity regions of AGN, are at a disadvantage with fewer to easily study and greater difficulty distinguishing their light from that of their resolved host galaxies. Until recently (Ho et al. 1994), LLAGN did not receive exclusive, intense study. Deep surveys have determined that nearly half of all bright, local galaxies contain fainter AGN. Along with a need to constrain the behavior of AGN at low-luminosities, this understanding increased interest in faint AGN and their properties. Nearby work with LLAGN identified them according to the relative strengths of

certain emission lines. This is insightful, this spectral work does not show any temporal variations in nuclear light. At most, only a handful of LLAGN have been monitored for changes in brightness. Data on the variability (timescales and amplitudes) of these objects is essential in connecting classical AGN properties to LLAGN. The project outlined here is to monitor a large subsample of LLAGN for long-term variability. For the time being, data for the monitoring campaign will be taken with Brigham Young University's Remote Observatory for Variable Object Research (ROVOR).

The project is a natural extension of Stephen McNeil's (2004) dissertation entitled "Detection of Small-amplitude Optical Variability in Galactic Nuclei." McNeil established one method to isolate small nuclear variations using data from large (1+ meter) telescopes. This thesis focuses on candidate selection, photometric techniques, and quantifying sources of error for the telescope and detector. Much of the material resides in appendices because the topics did not flow well in the body of the paper. The appendices are long and detail the process so future students may understand and improve this author's methods. The goal is to explore methods of data manipulation and analysis to gain precise (milli-mag) measurements with a small, remote observatory. By this, the ROVOR group seeks to contribute to the understanding of LLAGN variability.

Now, on to the history and understanding of AGN in general. In chapter six, LLAGN will be separated from their brighter counterparts and more fully explained.

Chapter 2

Surveys

Radio surveys first discovered what became known as AGN in the late 1950s. These include the Cambridge catalogue (3C and 4C), the Ohio, the PKS (Parkes) and the Arecibo Occultation (AO) surveys (Peterson 1997, 5-6). Loud, or strong, radio sources were often identified as galaxies, but some came from stellar-looking objects. As stars do not emit much radiation in radio waves, these unusual radio sources were called quasi-stellar radio sources, or quasars. Early known features of quasars included (Burbidge & Burbidge 1967):

1. Star-like objects often identified with radio sources
2. Variable light
3. Large ultraviolet (UV) flux of radiation
4. Broad emission lines in the spectra
5. Large redshifts of the spectrum lines

(quasar image-indistinguishable)

Once the unusual quasar features were found, several surveys began searching for quasars. A very effective search method uses simple Johnson filter photometry to determine the "colors" of objects. Color is the difference in magnitudes of one object from two observations in two separate filters. As AGN are typically very blue objects, early optical surveys used the color $U - B$ to distinguish

between hot stars and AGN (Krolik 24)¹ .

O and B type stars have a $U - B \cong 0.4 - 0.5$ magnitudes, while AGN had the color $U - B \cong -1$ mag² . Taking deep images and subtracting off the B magnitude from the U magnitude isolates nearby AGN candidates from stellar objects; spectral follow-up determines whether the candidates show characteristic emission line strengths. An influential color-magnitude survey is the Palomar Bright Quasar Survey (called BKS or PG for Palomar Green). Using the criteria $U - B \leq -0.44$ mag, over 1,700 candidates were found: spectroscopically, 7% of these were found to be AGN.

Color surveys assume that the nucleus outshines the host galaxy, and so a composite spectrum of the two will show a considerable blue excess. This type of survey is biased against faint AGN, which can be enshrouded in stellar light and hardly contributes to the composite spectrum. Also, single-color surveys (e.g., $U - B$) make the assumption that redshift is not important and that the AGN spectrum is smooth³ .

An alternate method places a large prism at the aperture of the telescope (called an objective prism). When exposing a photographic plate or a modern CCD in this manner, instead of point-like objects, each star and galaxy has its light spread into a spectrum. To be useful in surveys, the prism needs be low dispersing so the spectra of close objects do not overlap. A UV excess is identified by measuring the difference in flux between the far-blue and mid-blue regions, and these candidates are then imaged individually to determine spectral features and object type.

¹ U and B are the magnitudes of the objects in the Johnson U and B filters, respectively. Other Johnson-Cousins, Stromgren, and Sloan filter letters are used to denote the magnitude in that filter, such as V , I , i , u , r , g , and y .

²While this method works well for nearby objects, distant objects have their spectral features redshifted. Instead of finding the blue excess (such that $U - B$ is bright), $B - V$ or $V - R$ may be the colors which show the excess.

³Continuum AGN light originates from the accretion region, with a strong temperature gradient parallel to the disk. A star has one general temperature for its photosphere, producing a sharp Planck curve; an accretion disk has a continuum of temperatures, whose aggregate spectrum is usually smooth and extends over more wavelengths. An AGN disk spectrum usually peaks in the near UV with an effective thermal temperature of about $10^5 K$ (Peterson 43), which is hotter than the photosphere of any star. This spectrum extends both into the hard X-ray regime and in the IR out to about 1000 nm.

A major, early survey was started by Markarian in 1967 and expanded in 1981 (Markarian et al. 1981) found over 1,500 (primarily northern hemisphere) galaxies with a UV excess using an objective prism. The candidates were required to be extended objects, which have fuzziness around them which stars do not. (Krolik 24). The survey contains Seyfert and starburst galaxies, HII regions (nebulae), quasars, and BL Lac objects (Petrosian et al. 2007), with over 10% of them Seyfert AGN (Krolik 24).

Multi-color surveys are more capable of determining AGN. Most notable is the Sloan Digital Sky Survey (SDSS) begun in 2000, which uses five filters to produce four colors. This effectively eliminates stars and, by extension, galaxies with weak nuclear contributions. SDSS has identified over 120,000 quasars⁴ and provides images and follow-up spectra on most of these.

(sdss spectra of one of our AGN)

⁴www.sdss.org

Chapter 3

AGN Properties

Color surveys and radio observations help to determine positions of the potential AGN, but do not do much to determine their properties. Follow-up spectra and broadband monitoring of these candidates determine whether they are active by their redshift, the presence of broad emission lines, and variability, among other characteristic parameters.

3.1 Compactness

In high-resolution images (even with the Hubble Space Telescope), the innermost AGN components remain unresolved. For the most nearby cases, this limits the size of the nuclear structure to less than one cubic parsec (pc)¹. Kinematic studies and broadened emission lines indicate the mass contained within these regions to be on the order of 10^6 - $10^9 M_{\odot}$ ². For a black hole with mass $10^7 M_{\odot}$, the Schwarzschild radius $R_s = \frac{2Gm}{c^2}$ returns an event horizon with diameter of about 3 light-minutes, or approximately the distance between the Sun and Mercury. Additional components– the accretion disk, gas clouds, and a dusty region– increase the size of the nuclear re-

¹A parsec is equal to 3.26 light-years, and comes from a star having a parallax of one second (1/3600°) of arc.

²The symbol M_{\odot} is a solar mass, a unit which represents a multiple of the sun's mass of 1.99×10^{30} kg.

gion to several hundred parsecs (Peterson 1997). In addition, the central regions possess extremely high luminosities, rivaling or dominating the host galaxy contributions.

Supermassive black holes are widely accepted as the ubiquitous cause for this central mass. When this theory was first forwarded with respect to quasars, many scientists supported other theories such as: dense packs of star clusters, very massive, young stars, and very large collections of stellar remnants, including black holes. These compact objects are visible out to high redshifts³ ($z > 5$) and many show signs of long-term existence (like jet structures in radio galaxies). Rees (1977; 1984) argued that any arrangement other than that of a supermassive black hole would either dissipate or collapse to a singularity in a short time period. Strongly-nucleated galaxies do not allow much dissension from the standard supermassive black hole theory. Many low-luminosity AGN (LLAGN) are less massive (on the order of $10^5 M_{\odot}$) and other AGN features, such as variability, broad lines, and polarized light are not readily apparent. Alternate theories to active nuclei (see section 6.1.1) cannot be dismissed without further discussion.

3.2 Variability

In addition to radio loudness (not a property of all AGN) and UV excess, early on variability was found to be a characteristic of active nuclei (Smith & Hoeffleit 1963; Matthews & Sandage 1963). Nuclear variability is typically very dependent on wavelength and AGN spectral type. In X-rays, changes within minutes and hours are present for most luminous objects (Mushotzsky et al. 1993; Krolik 1999). It is interesting to note that while most do have compact X-ray emission, LLAGN show no significant short-term X-ray variations (Ptak et al. 1998). The UV-optical continuum changes on the order of days to weeks (Krolik 1999). Infrared (IR) light does so more slowly, over the course of years (Neugebauer et al. 1989); except for beamed radio emitters (synchrotron

³Redshift $z = \frac{\Delta\lambda}{\lambda}$ is the shift in wavelength of a rest-wavelength feature due to recession velocity between the emitter and the observer. If distance d is known, for Hubble Constant $H_0 \cong 70.8 \text{ km s}^{-1} \text{ Mpc}^{-1}$ and c the speed of light, $z = \frac{H_0 * d}{c}$.

emission— see section 4.2.2), radio variations are generally on the order of IR timescales (Ulrich et al. 1997).

(Krolik?- variable object light curve)

Radio-quiet (see section 4.1) LLAGN change in optical brightness by tens of percent in time frames of about a year. AGN 10^2 to 10^3 times brighter than these have been found to vary at about an order of magnitude less over the same time, and also to vary more slowly (Krolik 1999). The amplitude of variability, at least for the UV-IR range for typical AGN, decreases with increasing wavelength (Krolik 1999). For this reason, bluer filters (especially Johnson *B* and *V*) are traditionally used to monitor nuclear fluctuations.

3.3 Polarization

When the electric and magnetic fields are preferentially aligned in specific directions, photons passing through become slightly polarized (aligned) to the fields' orientations. By this process, stellar and galactic light is slightly polarized by the interstellar medium (ISM). Light from non-thermal processes should be polarized by several percent (Peterson 1997), but the typical AGN rarely shows values more than 0.5-2% (just barely above the contributions of the ISM), indicating that the bulk of their emission (especially in the optical) is thermal. The exceptions are strongly variable, radio-loud sources (blazars) and some AGN with only narrow emission lines: in these, around 10% of these objects' light is polarized (Krolik 1999). Those which show high polarization have a power-law spectrum associated with strong synchrotron emission.

3.4 Emission Lines

Most stars show line features only in absorption, and those that exhibit emission lines are of low magnitude relative to the strengths of emission lines from AGN (Krolik 1999). In many of these,

observable lines extend from the X-ray regime through the near-IR. These may be either from permitted or forbidden electron decay states, corresponding to broad and narrow lines respectively. Broad emission lines show rotational broadening of their lines⁴, corresponding to speeds of at least 1000-2000 km s⁻¹. Narrow lines require low-density settings in order to exist, and they have line-broadening closely corresponding to that of the nearby stellar velocity distribution. This indicates two regions— a more dense broad-line region (BLR) near a supermassive potential well, and a rarified, lower temperature narrow-line region (NLR) external to this. While the BLR is near the source of its excitation, the NLR appears to be galactic clouds lit by a more removed ionizing source.

3.5 Large Redshifts

Due to very high luminosities, quasars are the most distinct, identifiable objects in the universe. These have been observed out past a redshift of $z = 5$, corresponding to a lookback time of about 12.5 billion years. Though constrained by magnitude limitations⁵, the comoving density of bright AGN peaks just above $z = 2$, and strongly decreases beyond this point (Hewitt & Burbidge 1993). This indicates increasing mass and accretion before this point and dimming towards the present era.

With greater distance comes smaller apparent size of galaxies. At about $z = 1$, the host galaxy size shrinks to one arcsecond in diameter (Constantin 2011) and begins to be impossible to spatially resolve from the nuclear contributions in even the best space telescopes. Measuring the light (and

⁴Broadening may be from, among others: rotation, where different components of the emitter are Doppler broadening both towards and away from the line center; turbulence, where local motions within the gas Doppler shift that portion from the line center. Seth et al. (2010) is not alone in determining rotational characteristics of outer (resolved) nuclear regions and finding large-scale rotation about the galactic centers.

⁵As distance increases, so does the faintest absolute magnitude observable. Due to this, only LLAGN near the Milky Way can be studied. The density and properties of LLAGN with z cannot be studied as with bright AGN.

changes in intensity) for point sources is an easily fixed problem, but measuring changes in light from extended objects requires more care (see section 9.2). This thesis explores monitoring nearby ($z \approx 0$) AGN for variabilities, so galaxy contributions are almost always resolveable.

Chapter 4

AGN Classifications

The initial surveys above focused on finding either radio-loud or radio-quiet (optical) AGN. Note that radio-quiet objects still have a radio presence, and radio-loud nuclei are not devoid of a UV-optical continuum. Using $R_{r/o}$ as the ratio of (energy) flux in the radio to the optical¹, we get the following criteria for loud or quiet radio objects (Peterson 1997):

Radio loud has $R_{r/o} \geq 10$

Radio-quiet has $R_{r/o} \leq 1$

The radio-loud objects are radio galaxies, blazars, and up to 10-15% of quasars (Carroll & Ostlie 1996; Krolik, 1999). The rest of the quasars, along with Seyferts and LINERs make up the radio-quiet group. Radio-quiet AGN may be up to 20 times more common than their loud counterparts (Peterson 1997), making optical surveying more fruitful (also, optical surveys with plates or CCDs are both more sensitive² and much quicker than radio surveys). Quasars will not be treated as

¹The conventional measure is $R_{r/o} \equiv F_{\nu_r}/F_{\nu}(4400 \text{ \AA})$ where F_{ν} is the flux per unit frequency interval, 4400 Å is the optical section (though in angstroms and not frequency), and ν_r is the radio component, with one choice being 5 GHz.

² Even for radio-loud AGN, with $R_{r/o} = 30$, the luminosity ratio L_r/L_o at 5 GHz versus 680 THz (440 nm) is usually 2×10^{-4} , so only a small portion of the total luminosity is produced in radio frequencies (Krolik 1999).

separate objects here; they are best understood as distant, more luminous versions of either radio galaxies or Seyferts. A type of galaxy which is not active but which has some similar characteristics, is a starburst galaxy. These will be discussed in section 6.1.1.

4.1 Radio-quiet AGN

4.1.1 Seyferts

Seyfert galaxies are distinct host galaxies containing bright, stellar-like cores with strong emission lines. Almost all of these nuclei are found to be in spiral galaxies (Peterson 1997). In spectra, emission lines extend above the broad AGN continuum; these may either be broad or narrow, leading to the respective classification as type 1 and type 2. Type 1 Seyferts have both broad and narrow lines, while those of type 2 have only narrow emission features. Further distinction by Osterbrock (1981) has broken both into types³ 1, 1.2, 1.5, 1.8, 1.9, and 2.

Emission lines wider than a line's natural width⁴ are the result of some broadening mechanism. The simplest is rotation with a component along the line of sight, which Doppler shifts the approaching region blueward and the receding region redward. Typical notation gives the width of these and the forbidden lines in km s^{-1} ; for the broad lines, one may expect widths of 10^3 to 10^4 km s^{-1} . Broad

³ For instance, one showing only broad components in the $H\alpha$ line is labelled type 1.9; all other broad components are too weak to be visible. Continuum variability affects the strength– and even existence– of broad lines, and so these classifications depend on the activity level of the nucleus (Peterson 1997)

⁴ The uncertainty principle in quantum mechanics inversely couples the time between events and the uncertainty in energy. Thus an emission line will not have a clearly defined energy value but instead a range of values forming a Lorentzian profile emission line. Permitted lines come from the decays of excited electron states, while forbidden lines are from long-decay states. With a shorter decay time the spread in possible energies is larger for permitted lines than for forbidden lines. Permitted lines are called broad lines and forbidden lines are called narrow lines for this reason. The broadening effect now discussed is in addition to the natural broadening of the line.

lines are present only for permitted lines, most notably the HI series ($H\alpha$, $H\beta$, etc). This indicates that the broad emission lines are from a region of high enough density (electron density $n_e > 10^9 \text{ cm}^{-3}$) that metastable states are collisionally de-excited. The forbidden lines are produced in low-density regions ($n_e = 10^3 - 10^6 \text{ cm}^{-3}$) and have widths slightly faster—by a few hundred km s^{-1} —than non-AGN emission lines, indicating slower rotation about the nucleus (Peterson 1997).

4.1.2 LINERs

At times called dwarf Seyferts (Ho et al. 1994), low-ionization nuclear emission regions (LINERs) are on the bottom end of nuclear activity. Unlike Seyferts, which can show several broad and narrow lines simultaneously, LINERs (if they have any) only have broad $H\alpha$ lines⁵ (Ho et al. 1997). Broad emission lines have FWHM, which indicate a supermassive, superluminous and compact source (Constantin 2011). Without broad lines (i.e., for those with only narrow components), other methods are used to verify that an emission region is powered by an active nucleus. As photoionization depends on the energy source, higher-energy ionization states can distinguish between normal and active galaxies. Take oxygen, which is singly-ionized [O II] with 13.6 eV photons and only doubly-ionized [OIII] with energies above 35 eV (Ho et al. 1997; Krolik 1999). Both of these ionization states are uncommon near O and B type stars, but are abundant around bright active nuclei. Low-ionization nuclear regions, however, often have emission line strengths similar to starburst galaxies, and no single line ratio uniquely sets the two apart. Baldwin, Phillips, and Terlevich (1981) determined two sets of line ratios indicating that the region should be classified as a LINER (with clarification in Krolik 1999):

⁵Ho et al. have found that the sub-classifications 1.9 and 2 apply to LINERs and transition objects, where 2 refers to strictly narrow emission features, and 1.9 shows a (weak) broad component only in the $H\alpha$ line.

$$[\text{O III}] 5007^6 / H\beta < 3^7$$

$$[\text{O I}] 6300 / H\alpha > 0.05$$

$$[\text{S II}] 6716,6731 / H\alpha > 0.4$$

$$[\text{N II}] 6583 / H\alpha > 0.5$$

Diagnostic diagrams, which plot one object line ratio against another, determine separate spaces for separate types. Due to galaxy light contamination and the inherent ambiguity of the ratios, the type separations are best viewed as being more useful rather than physically distinguishing.

(diagnostic diagram)

There is some overlap between LINERs and normal galaxies in the diagnostic diagram, giving rise to a transitory type object.

4.1.3 Transition Objects

Between the line ratios of HII regions and LINERs are nuclei with intermediate [OI] line strengths. Making up about 15% of nearby galaxies (Ho et al. 1997), these are called weak-[OI] LINERs or transition objects, and were first thought to be composite objects, with HII regions about a LINER nucleus (Constantin 2011). As higher-resolution spectra do not better separate a typical LINER from HII clouds (Ho 2008), this model seems inaccurate. Other, non-active models (shock excitation, starbursts, etc.) have been proposed to explain the stronger emission excitation, but these cannot explain compact radio and x-ray cores observed in many transition objects. Instead, the preferred model posits that Seyferts become LINERs and LINERs become transition objects as accretion (and therefore disk luminosity) decreases (Ho 2008).

Ho et al. (1997) have given the criteria for a transition object as one with line ratios⁸:

⁶The bracket about OIII distinguishes it as a metastable, or forbidden transition. 5007 is the central, rest-frame wavelength of the transition in angstroms (Å).

⁷Note that high-ionization nuclei (Seyferts) have $[\text{OIII}] 5007 / H\beta > 3$ (Veilleux & Osterbrock 1987).

⁸The line ratios Ho et al. use are slightly different from the Krolik clarification, but the dissimilarities are minor.

$$[\text{O III}] 5007 / H\beta < 3$$

$$0.08 \leq [\text{O I}] 6300 / H\alpha < 0.17$$

$$[\text{S II}] 6716,6731 / H\alpha \geq 0.4$$

$$[\text{N II}] 6583 / H\alpha \geq 0.6$$

See section (LINERs and Transition...AGN) for a more complete treatment.

4.2 Radio-loud AGN

4.2.1 Radio Galaxies

Strong radio emissions found in the sky are either compact or extended sources. If from a nearby galaxy, the extended radio features can be distinguished from compact features (this resolveability is unique in AGN work). More distant radio galaxies (which are only detected when they are more luminous) are classified as radio-loud quasars, and the emission types cannot usually be separated. What distinguishes a radio galaxy from a normal galaxy is the large amount of radio flux observed; typical radio emission for this type is $10^3 - 10^7$ times that of a normal galaxy (Smith 1995). If both are present, the compact (nuclear) light, or that from the galaxy proper, is not the main contributor of radio flux. On either side from the compact source, two massive lobes of radio-emitting material extend from 50 kpc to 1 Mpc⁹. In some instances, jets are found to connect the compact and extended regions. These jets appear to continually feed material to the lobes, which are formed when the jets collide with intergalactic material and spread away from the jet (Smith 1995). The spectrum of the compact and extended radio portions follow a power law relation where flux $F_\nu = \nu^{-\alpha}$ where ν is the frequency of the light and $\alpha \in [0.5, 1]$ (Krolik 1999). Add in

⁹For comparison, the Andromeda galaxy (M31) is approximately 67 kpc in diameter.

strong polarization, and the radio source is non-thermal¹⁰. The emission which characterizes radio galaxies is called synchrotron radiation: relativistic electrons spinning about magnetic field lines emit tightly beamed light in the electrons' direction of motion (Walker 1987). It is also interesting to note that most radio galaxies are elliptical galaxies (Smith 1999).

4.2.2 Blazars

Blazars exhibit rapid variability across the electromagnetic spectrum (hence the connection with "blazing") . In this category is included BL Lac objects (named after its prototype BL Lacertae¹¹) and optically violent variables (OVVs). BL Lacs are featureless in their spectra, while OVVs show some broad and narrow emission-line features, but these features are not constant in time¹². It may also be that BL Lacs have emission lines, but are overwhelmed by the continuum source (Smith 1995). In addition to being radio-loud, blazars have strong and variable polarization in the optical range (Smith 1995), indicating non-thermal mechanisms. The luminosity of blazars can change significantly within a night (by 30-40% or more) across all wavelengths (Krolik 1999) and by many orders of magnitude over days and months (Carroll & Ostlie 1996). A class of blazars emit strongly in the very high energy (VHE) regime, with gamma rays reaching TeV energies.

¹⁰High-energy AGN spectra is also non-thermal, but is the result of mechanisms other than synchrotron radiation (Krolik 1999).

¹¹This object was originally classified as a variable star from its stellar nature and wild amplitude changes.

¹²The distinction between BL Lacs and OVVs is very small and little used today. The rapid variability also changes the presence/absence of emission lines and thus a blazar's subcategory becomes, at best, amorphous.

Chapter 5

Standard Model Morphology

In 1978, Osterbrock suggested that the distinction between type 1 and type 2 active galaxies was viewing angle. Around this time, accretion was found to be maximized if the accretion flow flattened into a disk, thus breaking spherical symmetry. To unify the two types of emission lines in AGN, active galaxy symmetry is further broken by having a region of high-density gas close to the accretion disk (called the broad line region) and a low-density region (called the narrow line region) further out. The broad-line region (BLR) is highly ionized and quickly rotates about the central mass, giving rise to rotationally-broadened lines typical of type 1 Seyferts. The narrow-line region (NLR) is axisymmetric, with two cones (along either accretion disk axis) being fairly devoid of this low-density gas. The NLR is spatially much larger than the BLR, and if one views the BLR through the NLR, the broad emission lines are shrouded and only the narrow emission lines are seen. Whether one sees a type 1 or type 2 AGN depends primarily on the angle at which it is viewed with respect to the disk axis.

Also in 1978, Blandford and Rees connected the featureless synchrotron radiation from blazars and the beams seen in radio galaxies. Instead of separate phenomena, viewing angle again determines whether a strong radio source is seen as a radio galaxy or a blazar. An observer looking down the jet in AGN would see a blazar, while, if seen off-axis, a radio galaxy would be visible. With

these theories, all types of active nuclei were united into one entity distinguished only by how it is viewed. Confirmation of these ideas have caused this model to be called the Standard Model of AGN. While not able to account for every particular, the model is attractive as it accounts for many of the observations.

5.1 AGN Evolution

The density¹ of bright AGN peaks at a redshift of $z \cong 2 - 3$ (Constantin 2011) with typical luminosity about 100 types greater than those in the current epoch (Krolik 1999). This "quasar era" corresponds to higher accretion rates and greater radiative efficiency than at other evolutionary periods. While the density and brightness of AGN have changed, data show that radio-loud objects have made up about 10% of the total for several billion years (Krolik 1999).

Accretion drives luminosity, and as low-luminosity AGN greatly outnumber bright ones for $z \sim 0$, accretion has decreased since the quasar era. One theory posits that different AGN are the same phenomena represented at different luminosities (Carroll & Ostriker 1996; Ho 2008). The high-luminosity objects transition over time to as depicted below:

Radio-loud Quasars and Blazars → Radio Galaxies → Normal Ellipticals

Radio-quiet Quasars → Seyferts → LINERs → Transition Objects →

Absorption-line Nuclei / Normal Spiral Galaxies

Recent findings summarized by Ho (2008) indicate that the standard AGN unification model does not apply to LLAGN. As luminosity decreases far below 1% of the Eddington luminosity (see appendix A.1.2), the optically thick, geometrically thin disk transforms into a shortened disk with a radiatively inefficient accretion flow (RIAF) towards the center. The RIAF is optically thin and almost spherical about the black hole out to a distance of 100-1000 times R_s . The "puffed-up" RIAF

¹Comoving density indicates the same volume of space adjusted for Hubble expansion.

region surrounding the singularity enhances the magnetic field strength from the poles and supports a greater portion of jet/outflow activity than radiatively effective AGN. Ho notes that LLAGN are not starved for accretion material, even though they have lower accretion rates than their more luminous cousins. Instead, the radiative efficiency of the accretion greatly decreases, with greater amounts of energy being drawn in by either advection and convection, or being transformed to kinetic energy in particle jets and outflows.

Chapter 6

The LLAGN Project

6.1 LINERs and Transition Objects as AGN

6.1.1 Con

LINERs and transition objects are interpreted by some as regions excited by purely stellar means. Broad lines indicate fast rotation, but narrow lines are also formed by hot O and B type stars in HII regions. As mentioned in section 3.4 (Bennert et al. 1996), the NLR is spatially extended and the nucleus is not the sole ionizer of these clouds, especially far from the singularity. The Palomar survey isolates the nuclear regions ($r \lesssim 200$ pc; Ho et al. 1997); so, excepting for some incidental contributions¹, the line ratios and fluxes presented therein are approximately correct.

A galaxy may be considered a starburst if: the star-forming rate is much higher than the galaxy's average over its lifetime, the ratio of far-IR to optical luminosities is much greater than unity, or if there is evidence for a great many HII regions. These criteria are generally concurrent in most starbursts. The far-IR luminosity of some starbursts can equal the luminosity of quasars ($\sim 10^{45}$ ergs s^{-1}), and are about as prevalent as quasars (Krolik 1999).

¹Most worrisome in starburst galaxies, especially with dense, circumnuclear star-forming regions.

Table 6.1 Median Statistical Properties of LLAGNs

<i>SpectralClass</i>	$L_{H(\alpha)}$	L_X	$f_R(\%)$	$f_X(\%)$
S1	1.9×10^{40}	1.4×10^{41}	72	100
S2	1.7×10^{39}	1.3×10^{39}	30	86
L1	3.7×10^{39}	8.8×10^{39}	63	95
L2	0.5×10^{39}	1.2×10^{39}	38	74
T	0.7×10^{39}	0.5×10^{39}	16	74

6.1.2 Pro

Rees (1977) noted that if many objects (cluster of stars, etc.) instead of a single accreting black hole were to explain the high masses in galactic centers, one would expect white noise instead of random, high-amplitude changes.

Ho (2008) presents properties from a subsample ($\sim 20\%$) of the Palomar survey, and the results are summarized in the following table. Luminosities $L_{H(\alpha)}$ and L_X are the integrated $H(\alpha)$ line and x-ray intensities, respectively, in ergs s^{-1} . The x-ray luminosity covers the 2-10 keV range, and radio intensity is determined at 5GHz. Percentages f_R and f_X are the occurrences of compact radio and x-ray cores out of the sample. These are lower bounds in radio and x-ray regimes as faint cores likely exist below the sensitivity range

Nucleated radio and x-ray emission is not a regular feature of stellar components (excepting supernova remnants). By their relation to Seyferts in terms of radio-core prevalence, LINERs display AGN properties. The surprisingly high x-ray core frequency in all types confirms that LLAGN as defined in section 4.1 generally represent active phenomena.

6.2 Palomar Survey

Using the Hale 5 meter telescope at the Palomar Observatory, Ho et al. (1997) took spectra of the nearest, brightest galaxies to find LLAGN signatures. The search criteria of the Palomar survey were that the integrated Johnson B apparent magnitude of the galaxy² $B_T \leq 12.5$ and that declination $\delta > 0$. While this does exclude AGN within faint galaxies, a strong group of almost 500 nearby³ galaxies were investigated. Unlike most other redshift and AGN surveys, the spectral resolution of the Palomar survey was better (.25-.4 nm) and the aperture smaller (2"×4"), resulting in less starlight contamination and an approximate resolution of less than 200 pc about the nucleus (Ho et al. 1995). The results, based on the general definitions of section 4.1, indicated nuclear activity in nearly 40% of the sample (Ho et al. 1994).

6.3 Project Scope

The Palomar survey identified nuclear emission lines which suggest non-stellar activity. This project seeks to confirm the other main component of active nuclei: variability. Particularly, the interest is to perform a search for long-term variability across the representative sample of Seyferts, LINERs, and transition objects. Those recognized as (primarily) HII regions by Ho et al. are not monitored at this time. For either purely stellar radiation (like clusters of Wolf-Rayet stars) or large populations of HII regions, a change in brightness of one component of the ensemble would not significantly affect the overall output. If nuclear variability is observed in a LLAGN, then there exists a single, intense source in a compact region (such that energy can propagate through the components between observations).

²The values were taken from the Revised Shapley-Ames Catalog of Bright Galaxies and the Second Reference Catalogue of Bright Galaxies, with an explanation beginning with de Vaucouleurs (1977) and continuing on in series.

³Due to the limiting magnitude, the Palomar survey only reaches to $z \approx 0.027$; using a Hubble constant of value $H_0 = 75$ km/s/Mpc, this redshift corresponds to a distance of 109 Mpc or 355 million light-years.

Using the 16" telescope at the Remote Observatory for Variable Object Research (ROVOR) in Delta, UT, a monitoring project has begun to periodically observe a broad target list of both known and potential LLAGN. Due to the large size of the list (see appendix A.2) and the required sensitivity of the survey, observing each object more than a few times a years is not practical. LINERs and transition objects are not expected to significantly vary on short time scales (Barth 2012).

6.3.1 Object List

From the Palomar survey of galaxies, we have chosen the brightest⁴ 78 LLAGN to monitor for long-term changes. Of these, 24 are Seyferts, 33 are LINERs, and 21 are transition objects. As a test of the project, seven well-known, nearby blazars were added as well as fifteen well-known variable Seyfert galaxies⁵. The test objects will vary with high amplitude and quick cadence, giving an early feel for the program's sensitivity. All together, 100 nearby AGN are to be monitored for long-term variations. Appendix A.2 contains a list of these targets and some of their parameters.

⁴The division between those considered and those not was arbitrarily chosen as $\log_{10}(F(H\alpha)) > -13.5$ where $F(H\alpha)$ is the flux of the narrow component of the $H\alpha$ line in ergs s^{-1} . Then, this group was narrowed to contain the brightest of each type in roughly equal proportions. As transition objects tend to be much fainter than Seyferts and even LINERs, they are less represented in the final target list.

⁵Many famous blazars and Seyferts were excluded for being much beyond the criteria $z < 0.05$. The one exception made was for BL Lac, for which we had already been obtaining data.

Chapter 7

ROVOR Site

This survey will primarily use the Remote Observatory for Variable Object Research (ROVOR) placed a few miles west of Delta, UT. Fourteen years ago, { } the ROVOR group was established to explore the benefits and difficulties of small-aperture, remote telescopes. While in a valley floor, the location is still 4,579 feet above sea level and is in a decently dark location. Currently, the observatory houses a 0.4 meter RC Optical f/9 telescope on a German-equatorial Paramount ME pier. Attached to the back of the telescope is a Finger Lakes Instrumentation (FLI) ProLine PL003 with a SITe SI-003A 1024×1024 $24 \mu\text{m}$ pixel detector (a remade Apogee Ap8). The resultant field of view (FOV) is 23.4" square and has a plate scale of 1.375"/pixel. Seeing at ROVOR is worse than first anticipated; typical seeing is about 3" but can reach 5" on poor nights¹. The telescope can attain better resolution, but the site was chosen for simplicity. More information on the site and telescope can be found in Moody et al. (2012; in preparation) and in appendices C.6 and F.6. ROVOR is remotely controlled from BYU campus and is primarily run by undergraduate students in the research group. Remotely operating the site requires regular weather checks and looking at on-site security cameras. An observer is not physically in Delta, and the system is set up to autonomously run a script through the night. This means that clouds can cross the FOV and seeing

¹Seeing, along with extinction, appears to be a function of Alt/Az position as well as time. See appendix C.

can dramatically change without a new target being chosen or operations shut down early. Strong winds smear the point sources into elliptical rather than circular shapes². At least once, when the weather was assumed to be good, the dome for the telescope was open when some slight rain wetted the mirror. That night's data was slightly less helpful than we had hoped. The researcher must go through all the data and determine which is useful and which contains too much error.

²Non-circular stellar profiles seem to be problematic to image convolution and subtraction. See appendix G.3.1 for further information.

Chapter 8

Data Reduction

Weather mishaps (and observer error) may irrecoverably damage the data, but some sources of error can be removed. CCDs are great improvements over photographic plates in sensitivity. While the digitization of data brings its own error, and some unremoveable, many sources of error are easily characterized and removed. Here, only a short explanation is given. Appendix D is more comprehensive. Removing error is called reduction, processing, or calibration, and for ROVOR centers on three areas: bias level, dark current, and flat fielding.

Every frame has a bias, or zero, level added to it when reading out the values in each CCD pixel. A bias image is created by taking a zero-second exposure with the shutter closed. The frame has a consistent average level, called the pedestal value, and may also show pixel-to-pixel variations. About 30 bias images are combined to form a master frame that is statistically accurate. The master bias is subtracted from every non-bias image.

From the vibrations of the silicon within the CCD, electrons are ejected and can become indistinguishable from electrons released from light (photoelectrons). This is called dark current, and decreases as the detector temperature decreases. In good detectors, with proper cooling, the dark current pattern is stable with time. Up until early 2009, an Apogee Ap47 CCD was used at ROVOR, but poor temperature regulation allowed hot pixels (those with high dark current) to drift through-

out the night. Processing current data from ROVOR is much easier (and the results more accurate) with the new Finger Lakes detector.

Note that the amount of light reaching each pixel is a function of position on the chip, and the sensitivity of one pixel to light is not exactly the same as its neighbors. Flat fielding takes a uniform light source (for example, an illuminated screen or the twilight sky) and records how the telescope and detector respond to light. This correction depends greatly on the filter in which one observes, and flat fields must be taken in each filter. The difficulty in obtaining high-quality flat field frames, and the lengths some astronomers have gone to obtain their own version of them, have labelled this step a dark art.

Bias and dark corrections are additive, but flat fielding is a multiplicative effect. Take an object frame (one showing celestial objects) $O(x,y)$ with pixel designation (x,y) . If a bias frame $B(x,y)$ and a dark frame $D(x,y)$, along with the flat field $F(x,y)$, are prepared, a processed object frame $O_P(x,y)$ is made by pixel-wise operations between the frames:

$$O_P(x,y) = \frac{O(x,y) - B(x,y) - D(x,y)}{F(x,y)}$$

While the above is an oversimplification of the true process, it shows the components needed to reduce data from ROVOR (and most other sites).

Chapter 9

Photometry

The process of measuring light is called photometry. For CCD images, this involves adding the values for all the pixels in the desired region.

(PSF profile and image)

For earth-based UV to IR observations, variable seeing from atmospheric changes will alter the shape of the point-spread function (PSF). Worse seeing results in blurrier, more smeared-out stars. Measuring the charge in a CCD pixel involves some error, and this error compounds as the light from an object is spread over more pixels. Ideally, all observations would be taken in space. Cost challenges this ideal. Most small observatories avoid targets which require pristine seeing (stellar PSFs in tight globular clusters, for example, overlap more as seeing worsens). Larger observatories with deep pockets may install adaptive optics to remove most atmospheric effects, but for ROVOR (and most observatories) economics blocks this path.

Variable seeing affects pure point sources¹ the same, but extended objects vary differently. When dealing with composite objects (like point-source AGN embedded within a resolved galaxy), ei-

¹An object is considered a point source when its angular size $\theta_o < \theta_T$ where θ_T is the minimum resolveable angle defined by $\sin(\theta_T) \approx \frac{1.22\lambda}{D}$ for wavelength λ and circular aperture D . Only in a few instances, and with the largest telescopes, is the photosphere of a star resolveable.

ther: one includes all components of an extended object in a large aperture, one ignores some extended light and lives with greater error, or one uses recent algorithms to adjust images for seeing differences.

9.1 Point-source Photometry

9.1.1 Fixed-aperture Photometry

The most basic photometric method uses fixed apertures to measure light. While one aperture may be adequate for one frame, if the seeing dramatically degrades (not that uncommon at some sites), the PSF may spill over the fixed region. To keep from losing light, then, the fixed aperture is set large enough to capture the light from the worst-case seeing PSFs. When the PSF FWHM changes significantly in a night, a large number of pixels are used for the aperture. Each pixel has uncertainties in its value, and the aggregate uncertainty of the measurement increases with the number of pixels in the aperture. For nights or locations with noticeably variable seeing, fixed-aperture photometry is not capable of producing very precise results.

9.1.2 Variable-aperture Photometry

Adjusting the photometric aperture to match the seeing (the average of the stellar PSFs) preserves the ratio of a target star's light captured in each frame. However, manually checking the FWHM and altering the aperture for each frame is not reasonable for most data sets. Instead, the Astronomy group has converted to performing much of its photometry in an ingenious IRAF program. Called NightPhot4 (Ben Rose), it interactively adjusts the photometric aperture based on the seeing in each frame. The full-width of the PSF at half the maximum counts (FWHM) of each ensemble star in each frame is sampled. The user is shown a curve of growth for each object and can then decide whether to eliminate any objects when averaging the FWHM. A multiple, set in the script,

of this averaged FWHM is set as the aperture when performing photometry on the objects in the frame. Adjusting the aperture for point sources by this method retains the same relative flux in each aperture and minimizes the error from the number of pixels. While more accurate, image sets with a large number of frames can be onerous. In the original version, the user is required to both verify the aperture centers are correct and to dismiss any objects with outlier FWHM. Some additional user input is also necessary. A modified version assumes that both the images are well aligned and that all stars will be used to determine the aperture radius. With a fast computer, this automation can measure photometry for several hundred frames per minute. This adaptation is not wise for faint comparison stars or excessive user worry.

9.1.3 Standard Star vs. Differential Photometry

A night is called photometric when it is not influenced by sky gradients (e.g., the moon) and is free of faint, high clouds. On these nights, well-established standard stars may be observed in conjunction with targets (and through a range of airmasses) to determine transformation and extinction coefficients, as well as zero-point offsets. Finding these allows one to transform photometry onto a standard system compatible with other observatories². There is a downside to this—placing standard stars on a general scale carries an uncertainty of around 0.005 mag (Landolt 2009), meaning that standardizing one’s measurements brings additional error³.

The uncertainty above comes from pinning one region of the sky to another, and pinning standard stars together across all right ascensions. To circumvent this, one may compare one’s target to other stars within the field of view. These comparison stars need to be stable with time; at times no in-

²Due to different flat fielding methods and distinct quantum efficiencies of different detectors, it may still be necessary to add a slight shift to match observatories and detectors. This shift, as evidenced in Barth et al. (2011) is usually on the order of 0.01-0.02 mag.

³Error in measurement σ_M and error in standardization σ_S propagate to make a total error of $\sigma_T = \sqrt{\sigma_M^2 + \sigma_S^2}$. Standardization of an object often results in 0.02 mag error or more.

field standards exist, and candidates must be studied to remove variable stars. Called a differential magnitude, the average of the in-field standard's instrumental magnitudes is subtracted off of the target's magnitude. This avoids atmospheric extinction problems and enables large comparison ensembles (>20 stars) to promote high-accuracy differential magnitudes (often well into the milli-mag range). In addition to its simplicity, because of its high accuracy, differential photometry is preferred by those seeking small variations. The downside of this method is that it is not as easy to compare results from one telescope to another, especially as one researcher's in-field standards may not be the same as another's ensemble. If possible, differential magnitude light curves should be standardized for the best of both methods.

9.2 Extended-object Photometry

All extended objects can be decomposed into components which would be point sources to the earth-bound observer. A galaxy is composed of point-source stars, and a nearby supernova remnant can be broken up into unresolvable portions. Varying seeing will change the PSF of each component, which affects the aggregate dispersion. The atmospheric dispersion of extended objects does not follow the same pattern as do point sources. Nearby AGN (for this project, those with $z \leq 0.05$) are usually embedded in a visibly extended host galaxy. Take a model of this scenario which has the galaxy split into 100 point sources which each have the same shape as the nucleus. Then, vary the seeing by changing the shape of the point sources. Observe how the light from the host galaxy contaminates what, for the variable aperture process, should be constant flux. If the angular size of the host galaxy— or its brightness relative to the nucleus— shrinks, then the error using traditional aperture photometry diminishes.

As point-source photometry does not perform well for most LLAGN, other methods were explored. One early idea was to increase the aperture radius to include the entire galaxy. Standard stars would

also be measured (though with smaller apertures) to find the differential magnitude. A similar method takes the radius (r_{25}) of the aperture as where the surface brightness (in mag/arcsec²) falls to 25 in the B filter⁴. Values for r_{25} have likely been found for all of the galaxies in this project, but measuring counts over hundreds to thousands of pixels brings too much error. The expected small variations of the nuclei would be washed out by the uncertainty from the large pixel area.

Another, better option, is convolution. This allows proper image comparison by either broadening or tightening one frame's PSFs to match another's. Once properly convolved, various methods and programs analyze the objects. Previous work by McNeil (2004) and Pearson et al. (2010) established the method of surface brightness profiles (SBPs) as effective in finding small nuclear variations in galaxies. Another method, called PSF-matching image subtraction convolves one frame to conform to the other, and takes the difference of the two seeing-matched frames, with the expectation that any residual is a real variation in brightness between exposures. Aaron Barth (2012) suggested this as a better alternative to SBPs. As Pearson the methods are more fully explained and compared.

Note that for fairly nucleated galaxies, using differential aperture photometry is still acceptable as long as all the significant flux from the galaxy is included in the region. PSF-matching image subtraction generally requires a pristine template image composed of frames taken during very good seeing; when such frames are not found, aperture photometry becomes a necessary fallback. This author suggests that in-field standards are measured with variable apertures while the target galaxy region has a fixed radius large enough to contain all of the discernable galaxy. Differential photometry is calculated in the same way as before.

⁴Typical sky brightness at ROVOR is between 19.5 and 21 mag/arcsec² in B , meaning that at this radius, the outer part of the galaxy is at least 40 times fainter than the background sky.

9.2.1 PSF-matching Image Subtraction

When looking for faint objects, or small-level variations, it is preferable to check the residual light in difference images. Tomaney & Crotts (1996) used Fourier transforms to match PSFs for micro-lensing surveys. Spatially varying seeing (across a detector) and variable focus lessened the effectiveness of their method. About this same time, Kochanski et al. (1996) introduced a program called NLSFIT which performed a nonlinear convolution and subsequent subtraction on aligned frames. As the code was made to monitor low-light changes in distant AGN, it seemed to be a good match for this paper's project. For large data sets, however, finding nonlinear solutions increases computing time. As newer codes are more effective⁵ and less time-consuming, we have not sought to try NLSFIT. Alard & Lupton (1998) decomposed the convolution function, called the kernel, into a set of linear basis functions⁶. An image often has a spatially-varying PSF, especially away from the center (Alard 2000), so the kernel solution and convolution are performed on subsections of the frame. The code (called ISIS) produced by Alard & Lupton is difficult to use. Becker wrapped ISIS into a more user-friendly program named HOTPANTS. Becker's program is preferred here, and its usefulness for this project is investigated.

For high-quality data, the traditional convolution works well. The telescope at ROVOR, however, has slight tracking problems which introduce ellipticity to the PSFs of an image. Finding frames good enough to make a template has been difficult. Yuan & Akerlof (2008) note that the ISIS program does not work well with non-radially symmetric PSF. To work with less pristine data, they modified the process of Alard & Lupton to account for asymmetric data. Instead of assuming the template is good, they changed the program to convolve both the template and the input image. Their code for asymmetric PSF cross-convolution image subtraction doesn't have a flashy acronym,

⁵The best convolution and subtraction in NLSFIT appears to produce artifacts which, in addition to being not aesthetically pleasing, does not produce quite the desired accuracy.

⁶Familiar basis functions include sines and cosines in Fourier analysis and polynomials in Taylor series. The accuracy of the basis function increases with the number of terms, but so does the necessary computation time.

so we will make one for them: ACCIS. More information on ISIS, HOTPANTS, and ACCIS are provided in appendix G.3.

(ROVOR asymmetric hotpants image)

9.3 Photometric Error

Regardless of the care one takes, any measurement in any experiment contains error. These errors may be either systematic (removeable) or random (unremoveable). Systemic errors effect all measurements the same, and examples in astronomy include bias and dark addition and flat field response (see appendix D.1). Processing, or calibrating, astronomical images seeks to remove the systematic errors. Random errors depend largely on the source, and in this field photon counting is the primary random error⁷. Random errors appear as uncertainties when stating an observed value (e.g., 15.5 ± 0.1 mag, with the uncertainty as 0.1 mag).

As LLAGN variations are expected to be small, and uncertainties may be mistaken for a change in nuclear light, understanding and stating uncertainties is essential. For this project a minimum of 3σ confidence will be required to publish a change in nuclear light. This error budget allows a confidence of 99.7% that a value outside 3σ is from a real change. A useful relation to determine the significance of variations is the signal-to-noise ratio $\frac{S}{N}$. The formula for this allows for easy error quantification and inclusion. A signal-to-noise ratio of 8 corresponds to 3σ confidence.

Error sources may be independent, meaning uncorrelated, or dependent. Take a function f that is formed from independent variables a, b , and c . Error relating to the standard deviation of f is given by

$$\sigma_f^2 = \left(\frac{\partial f}{\partial a}\right)^2 \sigma_a^2 + \left(\frac{\partial f}{\partial b}\right)^2 \sigma_b^2 + \left(\frac{\partial f}{\partial c}\right)^2 \sigma_c^2$$

This is called adding in quadrature. For simple cases, like $f = \pm a \pm b \pm c$,

⁷Electronic noise, correlated with photon-counting statistics in CCD technology, is random and includes uncertainty in all calibration frames.

$$\sigma_f^2 = \sigma_a^2 + \sigma_b^2 + \sigma_c^2$$

Multiplicative functions, like $f = a^i b^j c^k$ have error

$$\left(\frac{\sigma_f}{f}\right)^2 = i^2 \left(\frac{\sigma_a}{a}\right)^2 + j^2 \left(\frac{\sigma_b}{b}\right)^2 + k^2 \left(\frac{\sigma_c}{c}\right)^2$$

9.3.1 Signal-to-Noise

After the removal of (hopefully) all systematic errors, what remains should be random uncertainty. In CCDs, signal S is the count of photons⁸ (measured in electrons e^-) from the desired object recorded in some total exposure time t and over a number of pixels $N_{pix(obj)}$. If object flux over the integrated area of $N_{pix(obj)}$ is F_{obj} , then $S = F_{obj} \times G \times t$, where t is the exposure length and G is the conversion (inverse gain) between electrons and digital counts (see appendix C.5.1). Common sources of noise⁹ are: the read noise RN , and poisson, or "shot", noises from the object (obj), the background sky (sky), and the dark current (D). When independent and random, noise σ is added in quadrature (Taylor 1997, 75), such that for the above sources σ_{obj} , σ_{sky} , σ_D , and σ_{RN} , $N = \sqrt{\sigma_{obj}^2 + \sigma_{sky}^2 + \sigma_D^2 + \sigma_{RN}^2}$. Methods and values for each σ are taken from lecture notes by Michael Bolte in 2008¹⁰, but are consistent with other sources (Howell 2006). The uncertainty in the shot noise of an observation is equal to the square root of the total integrated counts. Thus,

⁸Many prefer to use analog-to-digital units (ADUs or counts) instead of photon-electrons. While this does lessen the impact of a poorly defined inverse gain G on $\frac{S}{N}$ calculations, electron counting is instead chosen to conform with the literature.

⁹The full error calculation includes error measurements for read noise (see appendix C.5.2), flat fielding (quantum efficiency and illumination responses- see appendix D.4), finite shutter speed, dark current (appendix D.3), charge transfer, temperature nonuniformities, airmass changing across a frame, and poisson "shot" noise. Shutter speed and changing airmass are not large problems unless either the detector is very large or the telescope maps large regions of sky ($\sim 1 \text{ \AA}^2$ square) onto the CCD. Charge transfer problems and temperature fluctuations (spatial and temporal) are usually not a problem for modern detectors/camera mounts. Of those that we do not explicitly treat, only σ_{flat} is significant for ROVOR or West Mountain (WMO) data.

¹⁰www.ucolick.org/~bolte/AY257/s_n.pdf

$\sigma_{obj} = \sqrt{F_{obj} \times G \times t}$, and likewise for the sky and dark current uncertainties. For read noise, $\sigma_{RN} = \sqrt{N_{pix(obj)} \times (RN)^2}$ unless the read noise is undersampled¹¹. When summing or averaging N frames together, RN increases by a factor of $\sqrt{N_{frames}}$.

The basic signal-to-noise ratio is

$$\frac{S}{N} = \frac{F_{obj} \times G \times t}{\sqrt{F_{obj} \times G \times t + N_{pix(obj)} [F_{sky} \times G \times t + D \times G \times t + N_{frames} \times (RN)^2]}}$$

Where F_{sky} is the average sky background flux in counts $s^{-1} \text{ pix}^{-1}$. This needs adapting when placing error budgets on small variations— and not the entire value— of an object's measurement.

For a fractional variation *flare* in flux of a target, the signal-to-noise of this change is

$$\frac{S}{N} = \frac{F_{obj} \times flare \times G \times t}{\sqrt{F_{obj} \times (1 + flare) \times G \times t + N_{pix(obj)} [F_{sky} \times G \times t + D \times G \times t + N_{frames} \times (RN)^2]}}$$

It is this the monitoring campaign looks to find, and the 3σ confidence relates to the flare in intensity and not all the nuclear light in an object. The value of *flare* is decreased until, for some maximum exposure time t (currently 20 minutes a night in each Johnson V and R filters¹²), $\frac{S}{N}$ is about to sink below a value of 8. This sets the expected sensitivity limit when observing a specific target, and typical limits on *flare* are between 1 and 3 percent.

9.4 Initial Results and Error Analysis

For this project and other AGN observations, acceptable practices include a mix of fixed- and variable-aperture photometry and PSF-matching image subtraction. Which method one uses depends on the the quality of the data and the brightness of the host galaxy relative the nucleus.

¹¹This occurs when the e^- unit value for the inverse gain is greater than that of the read noise. In this case, $\sigma_{RN} = \sqrt{N_{pix(obj)} \times [(RN)^2 + \frac{G^2}{4}]}$.

¹²A flare from the disk of an AGN preferentially increases the energy flux of bluer wavelengths, which is why B and V filters are more traditional in monitoring these objects. However, LLAGN do not generally exhibit much of a big blue bump. Couple this with a lower transmission of light in the B than in the V and R filters, and higher CCD quantum efficiency into the green to red range, and the R filter easily supplanted the B in the program.

Below, this author illustrates both methods on two targets of the monitoring program: Markarian 501 and NGC 4051.

9.4.1 Aperture Results and Analysis

To test the expected sensitivity of the project, a well-known test object, Markarian 501, was chosen. The object is a blazar in the target list, and began flaring in the very high energy range at the end of May, 2012. ROVOR data from February through June of 2012 and WMO data from April through June were analyzed using a form of fixed-aperture photometry. Mrk 501 displays a resolved, but small, galaxy component with integrated V filter magnitude of about 13.3¹³. To ensure all the light from the galaxy is contained in the aperture, a fixed 6" radius was chosen. The standard stars used were all but #11 in Pearson (2010) as this star saturated in many WMO frames; their apertures were chosen to be about half the diameter as that used for Mrk 501 to match the seeing.

The differential magnitude is for a frame is $mag_{diff} = mag_{obj} - average(mag_{standards})$. Error propagation is tricky. First, error calculations will be carried out in flux, or counts, and not magnitudes. The error in measuring the object includes flat fielding errors in addition to the typical signal-to-noise requirement.

$$\sigma_{flat} = \sqrt{(C_{flat} \times N_{flat})^{-2} + \sigma_{flat,diff}^2}$$

Where the first term references shot noise and $\sigma_{flat,diff}$ is the approximate error in the illumination correction over the region of interest of the detector. This should be at least 1% (0.01), and may be two or three times this value. N_{flat} and C_{flat} are the number of processed flats combined to make the master and the average counts pixel⁻¹ of the processed frames, respectively. Now, the error in measuring the counts of an individual object or standard (replace 'obj' with 'standard,i') is;

$$\sigma_{obj}(counts) = C_{obj} \times \sqrt{(C_{obj} \times \frac{S}{N})^{-2} + \sigma_{flat}^2}$$

When the standards are averaged together, the error of the averaged value is

¹³<http://simbad.u-strasbg.fr>

$$\sigma_{standards,ave}(counts) = \sqrt{\frac{\sum \sigma_{standard,i}^2(counts)}{N_{standards}}}$$


Finally, the error of the differential magnitude is

$$\sigma_{diff}(counts) = \sqrt{\sigma_{obj}^2(counts) + \sigma_{standards,ave}^2(counts)}$$

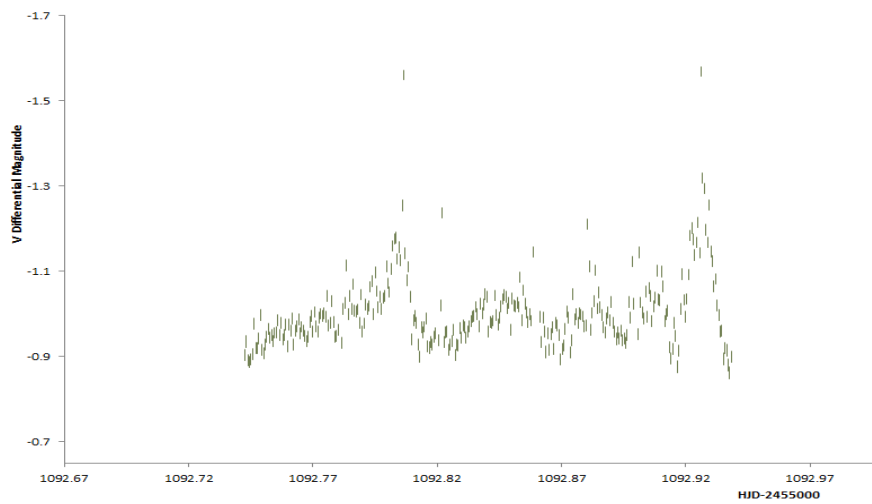

And transforming this to magnitudes,


$$\sigma_{diff}(mag) = 2.5 \times \log_{10}\left(\frac{C_{obj} + \sigma_{diff}(counts)}{C_{obj}}\right)$$

The resulting data is very positive. In addition to long-term variability, because of excellent time sampling, very rapid changes are visible in the WMO data. First, the Johnson *V* filter ROVOR data for 2012 is displayed, and shows large-scale variations of more than a tenth of a magnitude.



From WMO, the variations of Markarian 501 are even more interesting. Time sampling on the order of 25 seconds (about ten of these from read-out time) gives marvelous resolution and more precise results than ROVOR can supply. Three nights contained dedicated *V* filter data, and are wonderful examples of the rapid blazar variability which is so enjoyable to see.





9.4.2 Image Subtraction Results and Analysis

While ROVOR data does not (currently) work well in the HOTPANTS package, frames from West Mountain Observatory (WMO) run by BYU satisfy the necessary demands. The 36" fork-mounted telescope was used to take data on NGC 4051 (a classic Seyfert 1.2) on 55 nights from 27 June 2010 to 9 August 2011. Using the night of 27 July 2011 as a template for its better-than-average seeing (about 1.46"), all of the other nights were psf-matched and differenced by this template. If the input image is $I(x,y)$ and the convolved template is $T_C(x,y)$, the differenced frame is equal to $I(x,y) - T_C(x,y)$. Any residual flux of the core of NGC 4051 on a given date represents a real difference from the baseline date of the template¹⁴. Positive flux represents an excess and negative flux a dearth of light from the nucleus with respect to the intensity of the template. Before transforming to magnitudes, to avoid negative flux, a uniform offset equal to the template's nuclear flux¹⁵ was added to return each measurement to a true intensity scale, or close thereunto.

¹⁴Though this is not strictly differential photometry, which depends on stable reference stars in the field, the change in flux is tied to a reference date.

¹⁵As NGC 4051 is fairly nucleated, any error in this measuring its central, non-stellar flux is small. Doing the same for more embedded, fainter LLAGN targets will be more difficult. Care will need to be taken to determine the nuclear




Figure 9.1 WMO data from 2010, showing an active state.




Figure 9.2 WMO data from 2011, which is instead a quiescent state.

In these measurements, error was particularly difficult to measure. Note that in addition to the error of a single measurement, which is intrinsic to both the input and the template image, the contributions.

template image is then blurred to match the input, meaning that the original error measurement is no longer as accurate. Lastly, as it is the subtracted frame we measure, error propagation is necessary. Using the S/N ratio above, for the KAF 09000 detector in 2010-2011, parameters were added in: $D \sim 0.2$ counts $\text{pix}^{-1} \text{ s}^{-1}$, gain $G = 1.29$ e^-/ADU , and read noise $RN = 8.54$ e^- . Then, a relatively unimportant sky term equal to the average of the sky background (~ 200 counts pix^{-1}) was included, resulting in a per frame error of about 0.001522 mag¹⁶. The template had an error of about 0.001294 mag, and naively propagating this error in the usual way for subtraction ($\sigma_{SUB}^2 = \sigma_I^2 + \sigma_T^2$) gives an average resultant error of 0.0020 mag. Looking at the scatter of the real data, however, shows that the error bars are not large enough. Barth et al. (2011a) obtained data from WMO and other sources and found similar error values (~ 0.003 mag) which were too small, indicating that photon statistics were not the main contributor to the aggregate error¹⁷. They added another error term in quadrature with their initial values to reach the average observed scatter of $\sim 0.01 - 0.016$ mag, with WMO being more precise. Later, Barth et al. (2011b) converted to PSF-matching image subtraction and report that this method brings more accurate results than does aperture photometry.

Barth et al. never specify (and I do not believe they tried to identify) the source of error. The current author believes that the missing error is due primarily to flat fielding inaccuracies. As an object is not always centered on the same pixel across all frames, illumination and quantum efficiency responses will be different. Inasmuch as there remains error in correcting for these effects (largely shot noise from combining a limited number of flat field frames), this large source

¹⁶From Taylor (1997), traditional error measurements are given to one significant figure, but indulge the excessive accuracy for this portion.

¹⁷Their method of photometry was differential photometry using fixed apertures, which does not retain the same ratio of flux. Instead, scaling according to the FWHM of the stellar PSFs is preferred, and Rose (2006) recommends using an aperture of 2-3 times the FWHM. Their object, Zw 229-015, shows enough galaxy contribution outside this aperture that additional error is likely as the seeing of a frame changes. See section ___.

of error propagates into the object frames. Bolte says that flat fielding "can often be a significant, unquantified source of error" and it does not usually factor into the standard S/N relation. Appendix ___ documents the Boizelle-Holden-Young method, which determines the easiest representation of the error in both flux and magnitudes. In summary, take the integrated counts F of the target, the signal-to-noise relation $\frac{S}{N}$ determined from the above relation, the average counts $\text{pix}^{-1} \text{ s}^{-1}$ C_{flat} of the flat fields used to process the frames, and the number of flat frames N_{frames} combined together to form the master flat. Using this, the error in flux of an observation is

$$\sigma(flux) = \pm sum_{obj} \times \sqrt{\left(\frac{S}{N}\right)^{-2} + (N_{frames} \times C_{flat})^{-1}}$$

And the error in magnitudes is

$$\sigma(mag) = \pm 2.5 \times \log_{10} \left[1 + \sqrt{\left(\frac{S}{N}\right)^{-2} + (N_{frames} \times C_{flat})^{-1}} \right]$$

In the limiting cases, these expressions reduce to what one would expect for flux and magnitude errors.

Chapter 10

Conclusions

Both differential photometry and PSF-matching image subtraction have been shown to be effective at isolating small nuclear variations in active galaxies. Which is used depends primarily on the quality of data and the extent of a resolved galaxy. For tight, very nucleated galaxies, aperture photometry is likely adequate, especially if one does not have high-quality images with which to perform the image subtraction. For AGN with very prescient hosts, galaxy contamination necessitates using HOTPANTS or some other program to remove its contributions.

Both methods, as shown in this paper (with WMO data) are capable of detector changes of less than 0.005 mag, or changes in flux of under half of a percent. This far exceeds this authors initial expectations. Now, current ROVOR group members will need to extend this understanding to work with less-pristine– but still useable– data on LLAGN in the target list. Of special importance should be attention to error analysis and completely stating the errors involved when stating results. While extending error budgets to account for all possible sources of noise and uncertainty may negate what seemed to be changes in brightness, establishing ROVOR as a scientific contributor demands that all statements are made with integrity. If it so be that this project fails to detect variability in most LLAGN, at least upper bounds can be presented to the scientific community for their inspection.

Appendix A

AGN Parameters

Below is presented the monitoring campaign target list, along with the celestial coordinates, AGN type, black hole masses, and Eddington ratio estimates. When the monitoring campaign data is analyzed, these parameters will help analyze the properties of variability by the AGN type and accretion rates.

A.1 Explaining the Table

Objects for the monitoring campaign were chosen according to the classifications given by Ho et al. (1997). The AGN-type label is coded S=Seyfert, L=LINER, T=transition object; the number is from the Osterbrock (1981) method and references the observed strengths of broad emission lines. It should be noted that objects with an ^a afterwards were not taken from Ho et al. (1997), but are well studied objects beyond the scope of the Palomar survey. These outside sources do not have sub-type classifications (e.g., 1.2 or 1.5.) and the labels attached to them were taken from the SIMBAD Astronomical Database¹. As such, these classifications should be taken as only mostly correct. In addition, AGN type can change as the nucleus varies, so even the Palomar classifications

¹<http://simbad.u-strasbg.fr/simbad/>

change with time.

Columns for object position, redshift, central black hole mass, and Eddington ratio are also provided in the table. Position is given in right ascension and declination coordinates in the J2000 epoch.

A.1.1 Redshift

Ho et al. (1997) measured the redshifts of several emission lines in the wavelength-calibrated spectrograph images and averaged them to give the redshift of the object/galaxy. Of the 100 targets in the LLAGN monitoring project, all but one has a redshift of $z \geq 0.05$. This exception is BL Lac, and this was added (in conjunction with another project) to complete the set. The AGN from the Palomar survey strongly crowd towards a redshift of 0.004, and only a few go past 0.015. The supplementary AGN are usually further out, especially the blazars.

An alternate method of describing redshift is to determine the recession velocity. At very high redshifts, the formula for velocity is cosmological-model-dependent; the objects in the current target list are close enough to use the simplest recessional velocity $\sim cz$.

(cz histogram)

For non-Local Group objects, a pure Hubble-law relation of distance versus recessional velocity fits the data fairly well. Assuming this, distance to close objects is simply

$$d = \frac{cz}{H_0}.$$

H_0 above is the Hubble constant, and current values for this range from 70.8 km/s/Mpc to 75 km/s/Mpc. Using the measurement from WMAP ($H_0 = 70.8$ km/s/Mpc) the distance to the furthest object in the Palomar survey is just over 100 Mpc, or 350 million light-years.

A.1.2 Eddington Luminosity Ratio

As matter accretes around a black hole, energy is radiated away as light. The more matter accretes, the more the outgoing radiation pushes against infalling matter. The nucleus has reached what is called the Eddington luminosity when the radiation pressure balances accretion pressure². The Eddington luminosity $L_{Edd} = \frac{4\pi G c m_p}{\sigma_e} M$, where G is the gravitational constant, c is the speed of light, m_p is the mass of a proton, σ_e the Thompson scattering cross-section for the electron, and M the mass of the luminous body (here the central black hole). This reduces to $L_{Edd} = 3.2 \times 10^4 \left(\frac{M}{M_\odot}\right) L_\odot$ where L_\odot is the solar luminosity and $\left(\frac{M}{M_\odot}\right)$ makes it easy to use M in solar masses. For example, using reverberation mapping Bentz et al. (2006) found NGC 4151 to have a mass of about $4.57 \times 10^7 M_\odot$. This gives an Eddington luminosity $L_{Edd} = 5.6 * 10^{45}$ ergs s⁻¹, but from measured wavebands, the bolometric (total) luminosity is closer to $L_{bol} = 2 \times 10^{43}$ ergs s⁻¹ (Wang et al. 2010); the luminosity ratio $L_{bol}/L_{Edd} \approx 3.5 \times 10^{-3}$ shows that NGC 4151 is in a radiationally inefficiency state.

Ho (2008) gives the characteristic Eddington ratios for the AGN types shown below:

In addition to that from Bentz et al. (2006)^e, 33 more black hole masses and 28 Eddington ratios are found by referring to Levinson et al. (1995)^b, Gebhardt et al. (2000)^c, Peterson et al. (2004)^d, Ho, Darling, & Greene (2008)^f, Seth et al. (2010)^g, Goulding et al. (2010)^h, Woo et al. (2010)ⁱ, Denney et al. (2010)^j, McConnell et al. (2011)^k, and Grier et al. (2012a,b)^{l,m}. The black hole masses are found either by maser or stellar kinematics, or by reverberation mapping; targets with more than one method of observation can be dissimilar by nontrivial amounts, and to reconcile the

²This assumes spherical accretion; thin accretion disks are most effective, allowing greatest radiation perpendicular to the disk and a radially-varying disk temperature profile. For this model, radiation efficiency is around 0.06 – 0.42 (and typically 0.1) of the rest-mass energy (Krolik 1999). Other accretion forms, like advection-dominated flows, convection, or outflows, are most prevalent when an AGN's luminosity is far below the Eddington limit. In this case, the resultant radiative efficiencies are several orders of magnitude below the optically thick, geometrically thin model.

<i>SpectralClass</i>	L_{bol}/L_{Edd}
S1	1.1×10^{-3}
S2	5.9×10^{-6}
L1	1.0×10^{-5}
L2	4.8×10^{-6}
T	1.2×10^{-6}

different values for M_{BH} , the final mass is the average of all supplied values. As the monitoring campaign continues, it is likely more black hole masses will be determined (or found in past literature), and the target list below can then be updated. The last portion, the bolometric luminosity L_{BOL} , requires accurate measurements from high-energy to radio regimes.

Ho (2008) reports that, for the Palomar objects, $L_{bol} \approx 220 \times L(H\alpha)$. Unless specifically determined and given in available literature, the objects presented here rely on this approximation to determine L_{BOL} .

A.2. AGN Monitoring Targets

Object	Right Ascension (α)			Declination (δ)			Type	z	$M_{BH}(\times 10^7 M_{\odot})$	$\frac{L_{bol}}{L_{Edd}}$
	00	06	19.5	+20	12	10.6				
MK 335 ^a	00	06	19.5	+20	12	10.6	S1	0.025	1.60 ^(d,f,l)	
MK 937 ^a	00	10	09.9	-04	42	37.4	S1	0.03		
NGC 404	01	09	27	+35	43	04.7	L2	0.0006	0.045	$1.71 \times 10^{-4}(g)$
QSO B0111+021 ^a	01	13	43.1	+02	22	17	BL	0.047		
MK 1044 ^a	02	30	05.5	-08	59	53	S1	0.016	0.102 ^(f)	
UGC 2024 ^a	02	33	01.2	+00	25	14	S2	0.022		
NGC 1052	02	41	04.8	-08	15	20.8	L1.9	.0004		
NGC 1068	02	42	40.8	-00	00	47.8	S1.8	0.0034	1.28	$4.96 \times 10^{-2}(c,k)$
NGC 1667	03	01	42.4	+35	12	20	S2	0.014		
3C 84 ^a	03	19	48.1	+41	30	42	S2	0.018	4.0	$8.13 \times 10^{-2}(b)$
NGC 1358	03	33	39.7	-05	05	22.4	S2	0.013		
3C 120 ^a	04	33	11	+52	11	05.6	S1	0.033	4.84	$1.53 \times 10^{-2}(d,f,i,m)$
NGC 2273	06	50	08.7	+60	50	44.5	S2	0.007	0.78	$5.89 \times 10^{-3}(k)$
UGC 3927 ^a	07	37	30.1	+59	41	03	BL	0.041		
NGC 2681	08	53	32.7	+51	18	49.4	L1.9	0.003		
NGC 2685	08	55	34.8	58	44	03.9	S2/T2	0.004		
NGC 2655	08	55	37.7	+78	13	23.1	S2	0.006		
NGC 2768	09	11	37.5	+60	02	14	L2	0.0056		
NGC 2787	09	19	18.5	+69	12	12.2	L1.9	0.003	4.3	$1.38 \times 10^{-5}(k)$
NGC 2841	09	22	02.7	+50	58	35.3	L2	0.0028		
NGC 2985	09	50	22.2	+72	16	43.1	T1.9	0.0053		
MK 1239 ^a	09	52	19.1	-01	36	44	S1	0.019	0.138 ^(f)	

Object	Right Ascension (α)			Declination (δ)			Type	z	$M_{BH}(\times 10^7 M_{\odot})$	$\frac{L_{bol}}{L_{Edd}}$
NGC 4125	12	08	06	+65	10	26.9	T2	0.0057		
NGC 4143	12	09	36.1	+42	32	03.2	L1.9	0.004		
NGC 4151	12	10	32.7	+39	24	20.7	S1.5	0.0048	3.82	$4.26 \times 10^{-3}(e, f, i)$
NGC 4150	12	10	33.7	+30	24	05.4	T2	0.0023		
NGC 4192	12	13	48.3	+14	54	01.7	T2	0.004		
NGC 4203	12	15	05	+33	11	50.1	L1.9	0.0023		
NGC 4216	12	15	54.4	+13	08	57.8	T2	0.004		
NGC 4258	12	18	57.6	+47	18	13.4	S1.9	0.0016	3.99	$1.03 \times 10^{-5}(c, k)$
NGC 4261	12	19	23.2	+05	49	29.7	L2	0.0083	54.5	$7.34 \times 10^{-6}(c, k)$
NGC 4278	12	20	06.8	+29	16	50.7	L1.9	0.0023		
NGC 4314	12	22	32	+29	53	43.1	L2	0.0023		
NGC 4374	12	25	03.7	+12	53	13.1	L2	0.004	132.5	$1.05 \times 10^{-6}(c, k)$
NGC 4388	12	25	46.8	+12	39	43.5	S1.9	0.004	0.88	$2.39 \times 10^{-3}(k)$
NGC 4414	12	26	27.1	+31	13	24.8	T2	0.0023		
NGC 4419	12	26	56.4	+15	02	50.7	T2	0.004		
NGC 4438	12	27	45.7	+13	00	30.9	L1.9	0.004		
NGC 4450	12	28	29.6	+17	05	05.8	L1.9	0.004		
NGC 4457	12	28	59	+03	34	14.2	L2	0.0041		
NGC 4459	12	29	00	13	58	42.9	T2	0.004	7.48	$1.34 \times 10^{-5}(k)$
NGC 4477	12	30	02.2	+13	38	11.2	S2	0.004		
NGC 4486	12	30	49.4	+12	23	28	L2	0.004	440	$1.12 \times 10^{-6}(c, k)$
NGC 4501	12	31	59.2	+14	25	13.5	S2	0.004		
NGC 4527	12	34	08.5	+02	39	13.7	T2	0.0032		
NGC 4552	12	35	39.9	+12	33	21.7	T2	0.004		

Object	Right Ascension (α)			Declination (δ)			Type	z	$M_{BH} (\times 10^7 M_{\odot})$	$\frac{L_{bol}}{L_{Edd}}$
NGC 4569	12	36	49.8	+13	09	46.3	T2	0.004		
NGC 4589	12	37	25	+74	11	30.8	L2	0.0071		
NGC 4579	12	37	43.5	+11	49	05.5	S1.9/L1.9	0.004		
NGC 4698	12	48	22.9	+08	29	14.1	S2	0.004		
NGC 4736	12	50	53.2	+41	07	12.6	L2	0.001	0.67	$1.51 \times 10^{-5} (k)$
NGC 4826	12	56	43.7	+21	40	57	T2	0.001	0.14	$1.90 \times 10^{-4} (k)$
NGC 5005	13	10	56.3	+37	03	32.2	L1.9	0.005		
NGC 5033	13	13	27.5	+36	35	37.1	S1.5	0.004	4.17	$8.96 \times 10^{-5} (h)$
NGC 5055	13	15	49.3	+42	01	45.4	T2	0.0017		
IC 4218 ^a	13	17	03.4	-02	15	41	S1	0.019		
NGC 5194	13	29	52.7	+47	11	42.9	S2	0.0018	0.76	$1.79 \times 10^{-4} (h)$
NGC 5273	13	42	08.4	+35	39	15.3	S1.5	0.005		
MK 273 ^a	13	44	42	+55	53	13.2	S2	0.037		
NGC 5363	13	56	07.2	+05	15	16.9	L2	0.0053		
NGC 5548	14	17	59.5	+25	08	12.5	S1.5	0.016	4.81	$1.86 \times 10^{-3} (f, i, j)$
MK 841 ^a	15	04	01.1	+10	26	16.5	S1	0.036		
NGC 5838	15	05	26.3	+02	05	57.7	T2	0.007		
NGC 5921	15	21	56.4	+05	04	11	T2	0.006		
MK 290 ^a	15	35	52.4	+57	54	09.5	S1	0.03	2.43 ^(j)	
MK 493 ^a	15	59	09.6	+35	10	15.1	S1	0.031	0.065 ^(f)	
MK 501 ^a	16	53	52.2	+39	45	36	BL	0.034		
1ES 1959+650 ^a	19	59	59.8	+65	08	54.7	BL	0.048		
BL Lac ^a	22	02	43.3	+42	16	40	BL	0.069		
NGC 7217	22	07	52.4	+31	21	33.3	L2	0.0038		

Object	Right Ascension (α)			Declination (δ)			Type	z	$M_{BH} (\times 10^7 M_{\odot})$	$\frac{L_{bol}}{L_{Edd}}$
NGC 7331	22	37	04.1	+34	24	57.3	T2	0.0034		
UGC 12138 ^a	22	40	17	+08	03	13	S1	0.025		
MK 926 ^a	23	04	43.4	-08	41	08.5	S1	0.047		
NGC 7742	23	44	15.7	+10	46	01.6	T2/L2	0.0052		
NGC 7743	23	44	21.1	+09	56	02.6	S2	0.0056		
IC 1524 ^a	23	59	10.7	-04	07	37	S1	0.019		

Appendix B

Spectra to Optical Transformation

In an attempt to predict the exposure times for the LLAGN at ROVOR, an Excel worksheet was prepared to transform spectral data into broad-band photometry. Ho et al. (1997) provide stellar-subtracted line ratios for the major emission lines (in relation to the strength of the narrow portion of the $H\alpha$ line). The spectra were obtained in the rest wavelength ranges $\sim 4230 - 5110 \text{ \AA}$ and $\sim 6210 - 6860 \text{ \AA}$ with spectral resolutions of 4 and 2.5 \AA respectively. The major lines that were obtained for most of the 486 bright galaxies are $H\alpha$, $H\beta$, $H\gamma$, [O III] 5007, [O I] 6300, [N II] 6583, and [S II] 6713, 6731. In addition, they also give the flux and equivalent width of the narrow $H\alpha$ line component and the redshift of the object. Other lines, if found, are included, such as Na I 5890, 5896 and He II 4686; the lines [O III] 4959 and [N II] 6548 are included here by considering quantum decay probabilities of the neighboring lines [O III] 5007 and [N II] 6583 (Krolik 1999, 315).

To reproduce the spectra, the workbook makes $\frac{1}{4}$ nm divisions between 350 and 950 nm to fully cover filter throughput functions for the B , V , R , and I Johnson filters. The workbook makes a Reimann approximation to the spectrum, and all other modelings (CCD quantum efficiency, night sky spectrum, etc.) will use the same divisions. Based on the supplied redshift, the resultant continuum and emission features are shifted by a factor of $(1 + z)$. The shape of the continuum

may either be flat (as are most of the LLAGN spectrum presented in Ho et al. (1997)), or have a small bump added shortward of 450 nm (the big blue bump).

Appendix C

Astronomical CCD Technology

Understanding how the data is gathered and converted is a little tedious, but as this process (as would any other) introduces error in astronomical data, an overview is warranted. An exhaustive treatment of charge-coupled devices (CCDs) may be sought in the reference book *Scientific Charge-Coupled Devices* by James R. Janesick (2001). Easier references (and very good ones) to start out with are Mackay (1986), Janesick & Elliot (1992), Massey & Jacoby (1992), and Howell (2006). It should be noted that most recent articles on the subject assume the reader is familiar with the subject. These cite old references which contain out-of-date portions. When read these early, seminal works, understand that newer processes and materials have significantly altered the effectiveness of CCDs.

A CCD is a grouping of pixels that act as light buckets. That is all the astronomer has to gather his data. The observer needs to know how the light is collected, where spurious (non-sky) electrons come from, and the nonuniformities between 'buckets'. Understanding one's detector is important as many errors do not change the data in a qualitative way. Some errors, if not fixed, introduce errors of more than 2%. Though little, ignoring these small errors is not acceptable. Now we delve into the manufacturing and science of the great astronomical development of our time. Eat some sugar and get excited.

C.1 Manufacturing

The CCD manufacturing begins with a flat chip of bulk, crystalline silicon several hundred micrometers thick which is called the substrate. On this seed silicon, an epitaxial layer¹ of only a few micrometers is grown. Depending on how one wants to transfer stored charge, a second epitaxial layer (with slightly different composition) may be grown on the first. are grown on top of the substrate. One layer produces a surface channel for charge transfer, while two layers make a buried channel.

Both the substrate and the epitaxial layer(s) are doped with impurities. When the additive is boron, the doping is called p (positive) type; if phosphorous is used, it is n (negative) doping. Other dopants are also used, but boron and phosphorous are the most common. P-type impurities accept (essentially rob) weakly held silicon valence electrons, introducing (positive) holes in the silicon lattice (where displaced electrons may recombine). N-type elements contribute excess conduction-band² electrons (though the system is still neutral).

In the epitaxial layer, the dopant is added to the silicon vapor as the crystalline structure is grown; here the doping is uniform but not as saturated as the substrate, which is highly p-doped (p^+ or p^{++}). For the reader's understanding of the differences, typical values for the regions are listed (in terms of resistance times distance) are 10-100 ohms*cm and 0.01 ohms*cm for p and p^+ doping, respectively.

Once the scientific-grade silicon wafer is grown, manufacturers add electronic paths and circuitry. This is done by etching away or adding materials by photolithography. First, a thin layer of photoresist (a type of protective coating) is deposited on the outer epitaxial layer of the chip. Then, to

¹The epitaxial layer is higher-quality crystalline silicon with different doping (added impurities) than the bulk material. The substrate is needed as a seed crystal for the higher-grade silicon. Growth is usually the result of doped vapor deposition in a high temperature environment (Janesick 1992).

² The conduction band is an intermediate energy band in a gap between fully bound and fully ejected. Here, electrons can move freely in the lattice structure and will go towards areas of high electric potential.

make a mask, an optically dark coating is laid on glass (or another optically transparent material) and has thin grooves removed. These grooves allow light to pass through the mask only in certain regions. The mask is either overlain on the chip or held at a distance and intense UV light passes through the grooves. From this light, only specific, intended parts of the photoresist are evaporated. Etching solutions, materials, or further doping respectively removes silicon, adds new materials, or changes the composition of the existing silicon only in the exposed regions. Typically, more than seven photomasks are used to produce the gate structures, potential barriers between CCD columns (channel stops), and the output register, among other features. The features produced by photolithography can be, at times, smaller than 1 micrometer (μm) wide.

On the last epitaxial silicon layer, a thin insulator (usually SiO_2) is deposited, and on this is laid multiple layers of polysilicon which photolithography separates into distinct electrodes. A simple system uses three electrodes per pixel and to them separate voltages can be applied and altered; in this case, usual voltages for the three are 0, +5, and +10 V (other orientations and potentials are now more commonly used, but are harder to explain). When the voltages are applied in this configuration, a potential energy well is produced with a maximum closest to the highest voltage. Freed electrons move towards the potential maximum and congregate in what is called a well.

Collectively, the electrodes and the insulator make up what is called the gate structure of each pixel. The front side of a CCD is defined as when the gates are facing the viewer.

C.2 Buried vs. Surface Channels

The efficiency of charge transfer was a real concern in the first astronomical CCDs. Initially, only one type of doped silicon (p) was grown. Here, when positive voltages are applied to the electrodes, an electric potential well is produced with a maximum at the electrodes; electrons (e^-) are drawn up to the edge of the insulator and the charge is transferred along the surface of the silicon-insulator

junction (making what is called a surface channel). As the insulator is deposited on the top epitaxial layer, nonuniformities in the surface structure of the doped Si and SiO₂ introduce interference traps- areas where charge may be trapped or deferred during read out (Janesick & Elliot 1992). Charge transfer efficiency (CTE), the percent of the charge moved per well per phase transfer, started in CCDs at 98% per transfer- far too low for modern applications. Also, dark current generation (see appendix ___) at the Si-SiO₂ interface is 2 to 3 times greater than within the silicon itself (Janesick & Elliot 1992). While improvements have been made to surface-channel detectors, they are not the preferred choice.

A better alternative to surface channel transfer came from buried channel processes. Before the SiO₂ region is laid down, a layer of n-type silicon is grown above the p-type layer, and then the insulator is added. This addition changes the potential well to peak within the n-type layer, away from both the n-p border and the insulator region (Janesick & Elliot 1992). Here, displaced photoelectrons are drawn to reside within the n-type silicon, and during read out they move from pixel to pixel within the silicon. Transfers became nearly lossless by this method, with CTE now reaching 0.999999 (99.9999%) per pixel in the best systems (Janesick & Elliot 1992).

Now, buried channel systems cannot retain as many electrons in their wells- about 3 to 4 times fewer than with surface-channel detectors- but are worth the cost for lower dark current and in electron retention (Howell 2006). Usual well capacity for modern buried channels is still above 300,000 e⁻. To both decrease read-out time and limit the number of transfers, many modern CCDs have multiple read-out registers, often at each corner. This becomes more important for larger arrays (e.g., a 7168x9216 pixel array). For a 1024x1024 chip, the contents of a pixel will be transferred up to 2024 times. For a CTE of 99.999% (the best for ROVOR), on average about 1% of the charge in a well is lost by the time it is recorded, with the maximum at about 2%.

C.3 Photoelectrons

When light hits the crystalline silicon, an electron (or multiple electrons) is excited out of the valence shell into a conduction band. Like in a metal, now the e^- may move freely within the atomic lattice structure of the silicon until recombining into the valence shell of another silicon atom. Typical recombination time for a free electron in lightly (p) doped silicon is $100 \mu\text{s}$ (Howell 10).³ When an electron is bumped into the conduction band and moves away a hole is produced, and the two together are called an electron-hole pair. The potential produced by the electrodes drives away majority carriers⁴, leaving a depletion region (this is the well) where displaced e^- congregate without recombining.

To reach the first conduction band, a photon need have at least 1.14 eV (or a wavelength of 1.24 μm); to produce only one electron-hole pair, the energy need by the photon is less than about 5 eV (or 250 nm). From here to 10 keV, multiple pairs are produced; past the soft X-ray range, photon interaction with the silicon becomes unlikely, and the medium becomes unfruitful as a detector (Janesick 7). This defines the useful interaction range of silicon as a detector from the near-IR to the soft X-ray. Other more exotic materials (like HgCdTe- mercury cadmium telluride- detectors for IR observations) are used for regimes where silicon is not effective.

At this point, it is important to discuss two forms of noise when measuring light with CCDs. Dark current and cosmic ray events also produce conduction-band electrons, though not with the same origins. Cosmic rays are celestial⁵, but random, occurrences. Dark current is local to the camera, and is mostly repeatable. Some further detail will be given on each of the effects.

³ Note that the highly p-doped substrate, with its abundance of holes quickly reaccepts photoelectrons. Photoelectrons reaching the well originate in the epitaxial layer(s).

⁴ For n-type Si, this means electrons; for p-type, holes are the majority carriers.

⁵Cosmic rays are easily distinguished by their large charge deposition into a decidedly non-stellar patten. Early camera glass windows emitted X-rays which the CCD readily recorded with the same profile as γ but different composition than γ cosmic rays. This initially gave several times the expected cosmic-ray rate (Mackay 1986).

C.3.1 Dark Current

Vibrations of the crystalline silicon periodically eject electrons into the conduction band. When these are produced, they are indistinguishable from photoelectrons. Thankfully, for a stable CCD, the pattern produced is likewise stable. Some pixels, either by manufacturing nonuniformities or more active electronics, are called hot pixels because of their high dark current relative to their neighbors.

Cooling the CCD significantly reduces the dark current in each pixel. This is accomplished either by liquid nitrogen in larger observatories or by on-board thermoelectric coolers at the rest. Taking a refitted Apogee Ap8, the dark current at different temperatures was tested. A curve of growth for average dark current $D(T)$ versus temperature T shows the trend to be fit best by a function of $D(T) = A \times \exp(0.1389 \times T) + C$. This exponential growth causes double the dark current with each 5° C increase in temperature. Further discussion is found in appendix ___.

C.3.2 Cosmic Ray Events

When high-energy particles (mainly protons) reach the earth's atmosphere⁶ and collide with other matter, they form cascades of particles which decay into muons. When the secondary particles reach the detector, they free about 80 electrons μm^{-1} into the conduction band (Mackay 1986). For ROVOR's thinned, backside-illuminated detector, the available silicon depth is about 8 μm . This results in between 500 and 1000 electrons (between 2500 and 5000 ADU, or counts) per muon event, usually concentrated groups affecting 1-3 pixels. Very energetic events often saturate

⁶Not all events recorded as cosmic rays are truly from without the atmosphere. Radioactive material in the CCD window and dewar (housing) produces the same effect and can greatly skew the actual count rate of cosmic rays (Howell 2006). Material improvements have decreased the "cosmic rays" originating quite close to the detector. As it doesn't much matter where these spurious photoelectrons are from, no further effort is taken to distinguish between the two.

pixels, meaning they produce more electrons than the analog-to-digital converter can make into a digital number. Mackay (1986) found that a typical cosmic-ray event rate is 1.5-2 events $\text{cm}^{-2} \text{min}^{-1}$, and data taken at ROVOR shows only a slightly higher value of ~ 6 events $\text{cm}^{-2} \text{min}^{-1}$ (see appendix ___ for details and implications). An energetic cosmic-ray event can ruin photometry on an object in a frame, and so it is recommended that instead of a single, longer exposure on an object, several shorter exposures be taken and all good images coadded together to equal the desired total exposure length⁷.

C.4 Illumination Orientation

The incident direction of light greatly impacts the sensitivity of the detector. Front-side illuminated detectors were the first on the scene, but backside-oriented ones have become more popular due to their higher sensitivity to light (especially bluer wavelengths). Unlike the charge transfer routes (surface or buried channels) which are determined during manufacturing, one may change a front-side to a backside CCD (given the funds).

For a front-side illuminated setup, light first must go through the gate structure to reach the silicon. The polysilicon and SiO_2 reflect a great deal of light. An alternate method for detecting photons with CCDs reverses the direction of the chip. Light is incident on the back of the p-doped epitaxial layer of silicon (where the substrate connected to) and the ejected electrons are collected as above. This is called a back-illuminated CCD, and can be made by removing⁸ the silicon substrate from the back of a front-illuminated CCD (Janesick 1992). As the light does not need to go through the highly-reflective gate structure, much more of the light reaches the silicon. The increased sensitivity is greater in bluer than in redder regions.

⁷For low-signal objects, such that the signal-to-noise relation in ___ is dominated by read noise (RN), note that splitting up a long exposure as explained does preserve the $\frac{S}{N}$ value from a single-exposure data point.

⁸ This is accomplished with an acid solution.

The average distance a given photon delves into the silicon layers determines whether the ejected electron reaches the well. If the average absorption depth for a certain wavelength grows much smaller than the distance from the incident surface of the silicon to the well, it becomes more probable that the electron will recombine before entering the well. For photoelectrons to reach the well, backside-illuminated detectors must have their silicon depth thinned from about 300 μm to 8 μm (removing the substrate silicon).

The benefits are not without their (initial) drawbacks. Exposing the back silicon (by acid etching) unintentionally produces what is called a native oxide layer. The oxide carries with it a positive charge and creates a potential well for electrons close to the backside surface, far from the buried channel. The expected improvement in sensitivity was lost as many photoelectrons became trapped in this second well. When it was understood, thinned-CCD producers introduced a variety of solutions: flooding the sensor before an exposure with intense UV light; adding a monolayer of gold or platinum to the back surface, and often biasing with a voltage; and, most often, doping the exposed silicon with boron to create a permanent layer with electrons to drive photoelectrons to the buried channel (Janesick 1992). The UV-light flooding is still used today as a sensitivity pre-flash on some detectors (take the West Mountain Observatory backup camera, for example). Another serious issue arose from the thinning: nonuniformities in the thinned depth (especially at the edges) plagued the focusing and quantum efficiency of earlier backside illuminated detectors (Janesick 1992). Much care went into learning acid etching rates and how to produce smooth-surface CCDs when removing the substrate.

Improvements in backside illuminated detectors, both in thinning techniques and effects, have made these devices the preferred type of imagers.

While the gate structure is more damaging to sensitivity than exposed silicon, both illumination orientations preferentially reflect higher energy photons, meaning fewer bluer detections. For Si, reflections peak in the UV around 300 nm (Janesick 2001). Certain coatings applied to the illumi-

nated side have proven to be effective in lessening this loss. This coating is usually composed of a thin layer of organic phosphors. These are transparent to optical light (400-900 nm) but down-shift UV light wavelengths into wavelengths easily seen by a CCD (even without the need for thinning). They also serve as antireflection (AR) coatings, and can increase the optical blue response by 15-20% (Howell 2006).

C.5 Reading Out the CCD

Channel-stops are placed between column wells, made of highly p-doped silicon (p^+) with an excess of holes. This constrains the movements of displaced electrons to the nearest well. When the exposure is finished, a clocking mechanism shifts the three-tiered potential, creating a many potential walls that move charge towards the read-out register. One row at a time is pushed onto a read-out register, which contain nonilluminated silicon pixels mirroring the row from which it receives electrons. Before another row is pushed out of the illuminated region, the register is read out, pixel-by-pixel, through a series of amplifiers and an analog-to-digital converter (ADC); each remaining row on the chip is shifted down by one step towards the read-out register, and the next row is read out through the amplifier and ADC. Large CCDs often have multiple read-out registers (usually one at each corner) to simultaneously measure pixels and by this reduce read-out time.

Now, we return to the read-out register. Wells are transferred row by row to the register, and then pixel-by-pixel through a series of amplifiers and other electronics. The in-chip read-out amplifier takes low-level voltages and multiplies to be accessible for off-chip amplifiers (called head electronics). After this increase in strength, the signal is sampled at least twice by the analog-to-digital converter and assigned a digital number (DN) from the analog signal.

There is a data unit capacity in bits. Today, 16-bit is standard, giving 65535 integer values for the ADC to assign⁹. The former standard was 15-bit for the magnitude— up to a value of 32767— and

⁹ The number of possible values for n-bit integer systems is 2^n . For 16-bit data with possibilities starting at 0, the

one bit for the sign, positive or negative. As anti-photons do not exist, all values should be nonzero; hence, the bit used to track the sign was added to the digits available for DNs. When sampling a pixel, to ensure that no negative values are assigned nonnegative DNs, a bias, or reference, voltage is sampled and added to the pixel's voltage (in case a small, negative value may be read by slightly varying electronics). This results in adding a certain average DN to each pixel before digital conversion, with a Gaussian spread about the average due to the ADC clipping. The average effect of the bias effect (in DN) of the bias is called a pedestal value; often, slight, repeatable pixel-to-pixel variations also issue from the CCD electronics .

C.5.1 Gain

How the DN is chosen is determined by the gain of the ADC. To optimize the chip one uses, the gain is set as follows: a pixel well has some capacity (say 300,000 e^-) before nonlinearity¹⁰ in the pixel response becomes a serious problem. The gain optimizes the dynamical range of the chip if the gain $G \sim$ well depth/ possible integer values (for the capacity chosen in the 16-bit era, G should be set at about 4.58). If the gain is set at an inappropriate value, it may not allow the full dynamical range of the CCD¹¹ .

If the gain is set too high, it introduces great error into the digital data. The units of gain (sometimes called inverse gain) are e^-/ADU (analog-to-digital units). From this, a source of irreversible error is introduced— the ADC applies a flooring mechanism to the signal it receives. That is, if the gain $G = 10 e^-/ADU$ and 321 e^- are detected, 32 ADU (alternately called counts or DN) are maximum assignable integer is $2^{16}-1=65535$.

¹⁰ Nonlinearity is caused by the stored electrons in a well lessening the potential from the electrodes. A conduction-band electron in this case is less likely to come to rest in the depletion region near the gate structure.

¹¹ Most worrying would be if the gain allowed serious nonlinearity to be recorded. While some nonlinear response is expected at $> \sim 60,000$ ADU, if instead this threshold were much lower, an observer may obtain nonlinear data and unknowingly skew results.

recorded. Any remainder is lost. While loss by the ADC is inevitable, if the gain is high (20-100 e^-/ADU), large amounts of signal are lost. This effect is especially grievous in low-signal sources, where this remainder becomes a significant portion of the signal. Many modern detectors have low gain set between 4 and 4.5 e^-/ADU (Howell 2006)¹², but others are closer to unity.

C.5.2 Read Noise

Reading out a pixel will introduce noise to the measurement. This is called read noise (RN), and, from Howell (2006), it is a composite of two effects. First, the analog-to-digital conversion is not perfectly repeatable: a real gain for a conversion will be one from a distribution with the mean value equal to the ideal, reported gain. Second, electronic noise introduces new electrons into the charge transfer and measurement processes.

Read noise is quoted in terms of e^- , with reference to the average number of additional electrons introduced in a measurement (per pixel). As current increases when reading out a frame, the amplifier temperature increases and swings. Temperature fluctuations significantly affect the amplifier's sensitivity. The faster the read-out speed, the more the noise is introduced. Slower read-out times introduce less RN, but take up precious observing time.

From Massey & Jacoby (1992),

$$G = \frac{(\bar{F}_1 + \bar{F}_2) - (\bar{B}_1 + \bar{B}_2)}{\sigma_{F_1 - F_2}^2 - \sigma_{B_1 - B_2}^2}$$

$$RN = 2^{-\frac{1}{2}} \times G \times \sigma_{B_1 - B_2}$$

For F_1 , F_2 two good (reduced) flat frames and B_1 and B_2 two good bias frames. \bar{F} and \bar{B} are the average of these frames and $\sigma_{F_1 - F_2}$, $\sigma_{B_1 - B_2}$ are the standard deviation of the subtracted frames. Rather than using the entire frame, a subsection containing the central 5-10% seems to produce the most consistent results.

¹² Note that noninverted gain is usually listed in terms of $\mu\text{V}/e^-$ instead of ADU/e^- .

C.5.3 Nonlinearity

It is easy to check if a pixel is saturated (reaches the maximum attainable value). Certainly, this affects measurements, and the overflowed well often spills into adjacent wells. More worrying is reaching nonlinearity before saturation; as a greater number of electrons fill the well, more energy is required for each additional electron to be retained. The nonlinear regime usually comes unannounced and hampers one's ability to produce quality data. It is not enough to assume that the CCD behaves linearly. If the value for the gain is correctly set, this nonlinearity usually sets in at high ADUs (if the maximum is 65535, nonlinearity beginning at 55000 or higher isn't unusual).

Nonlinearity may also affect the initial range at low light levels. Early surface-channel CCDs required an addition to each well before transfer called a fat zero (similar in ways to a UV preflash), or the charge would often be deferred in transfer (). Developments in buried-channel technology since their inception have all but eliminated low-signal nonlinearity. Neither nonlinearities in CCDs, though, are as serious as those plaguing photographic plates used for most of the 20th century (Howell 2006).

C.6 ROVOR Detector

The Remote Observatory for Variable Object Research currently uses a SITe SI-003A detector¹³, 1024x1024 pixel dimensions with square 24 μm pixels and a full-well signal of up to 350,000 e^- . It is housed in a Finger Lakes Instruments ProLine PL003 with on-board thermoelectric cooler. The chip is back-illuminated, with a standard anti-reflection coating, and has a buried channel for .99995 to .99999 CTE. Its gain $G \approx (4.0, 4.5) e^-/\text{ADU}$ and its $RN \sim 17 e^-$. Quantum efficiency peaks at 85% around 650 nm for the back-illuminated version, but only gets up to 42% if the

¹³ This was an Apogee Ap8 which Finger Lakes Instruments refitted.

detector is front-illuminated.

The linearity of the detector was tested using a three non-variable stars in the same field (HD 6717, 67109, and 67283) and using an aperture which integrates over the full PSF of the object¹⁴. Multiple exposures were taken at seventeen different lengths from 3 to 120 seconds and each saturated at a different point. The night this was tested showed variable seeing and the airmasses ran from 1.25 to 1.84 airmasses, so plotting the peak value against the exposure length was not viable. Instead, what was plotted was the fraction of light in the central pixel to that of the whole aperture, versus the exposure length. A typical k_V coefficient was used to correct for atmospheric extinction. When saturation sets in, the central fraction of the whole strongly decrease as the wings of the PSF increase, Non-linearity occurs when the range of fractions begins to be constrained, but visually saturation has not begun. From the data, it appears that the detector is linear until around 60,000 counts. Further study¹⁵ during a photometric night may be warranted to tighten up this figure.

¹⁴The program used to find the photometry was NightPhot4, which preserves photometric ratios between frames despite changes in seeing. See appendix ___.

¹⁵An interesting alternative method is suggested – take flat-images (see appendix ___) in rapid cadence as the morning sky brightens until the entire frame is saturated. Each linear, quality flat frame should have close to the same pattern. Thus, after processing the frames, divide each i^{th} by a good-quality master flat, such that $F_{div} = \frac{F_i}{F_M}$. Observe the point where the brightest regions of F_{div} begin to dim. At this point in F_i , and consequently for the detector, nonlinearity has begun. Nonlinearity is a function of the properties of each pixel (like well capacity), and so it will set in at slightly different values at different pixels.

Appendix D

CCD Reductions

By tradition, knowledge has passed, from astronomer to apprentice, generations untold. Professional astronomers have several duties that occupy their time, including their own research, teaching courses, and getting funding. When mentoring a student, an astronomer cannot fully explain the purpose and method of every principle. Students are expected to delve into and learn the material by themselves. However, if students are not properly mentored in basic procedures, the student's own progress will be delayed.

apparatus.

The purpose of this appendix is to introduce the student to both the purposes and methods of data reduction. Removing error introduced by charge-coupled devices (CCDs) (see Figure ?? for an example) is alternately called reducing, processing, or calibrating data. Much of the interesting (and relatively unstudied) phenomena in astronomy require accuracies 1 to 2 % or less (Ivezic et al. 2011), while systemic errors from using telescopes and modern detectors can be several times this limit. The author hopes that the material will illuminate methods which the reader may have performed without understanding. Also, the reader may find that some prior methods need adjusting. The approach will be mostly general, but at times will focus the process based on the system at ROVOR.

D.1 Noise and its Removal

Some error, or noise, is unavoidable. Noise may be either random or systemic. Randomness returns repeated measurements clustered about a central value; systemic error is the offset from what is considered the correct average value. For instance, read noise indicates the random noise from measurements made with imperfect electronics and includes the unavoidable error from digitizing analog charge (related to the gain of the detector; see appendix ___). Other random noise sources include (Janesick 2001):

- Intrinsic noise, which covers random noise in measurement (read noise) and photon "shot" noise, among others.
- Man-made noise, such as street light emissions and electronic discharges from power lines.
- Natural events, which include cosmic ray hits and lightning.

Other sources are systemic, such as:

- Fixed-pattern bias structures.
- Dark current.
- Quantum efficiency and illumination nonuniformities.

Random noise limits an observer's ability to detect and measure small amounts of charge in a pixel. Systemic errors, if not processed out of data, adds extra noise to the unavoidable sources.

For CCDs, the primary calibration steps are the zero (or bias), dark, and flat field corrections. The different steps are applied one at a time to the object frame (the image from which data values will be extracted). Each step corrects a separate effect in the object frame, and the correction type determines how to apply the correction. For instance, the zeros and the darks are additive

corrections while flat fielding is multiplicative. A calibration step may be a single-value correction (one value applied to all pixels in the object frame) or a full-image correction (each pixel in the object frame is corrected by the corresponding pixel in the calibration frame).

Object frames are not the only frames which need processing. A raw (unprocessed) dark frame contains the bias signature beneath the dark signal. Also, any images of nonzero exposure length will have dark current on a per-pixel basis. Lastly, any images taken with the shutter open need to have its pixels corrected for nonuniform illumination and sensitivity by a flat field. The general formula for reducing an object frame $O(x,y)$ is:

$$O_P(x,y) = \frac{O(x,y) - B(x,y) - D(x,y)}{F(x,y)}$$

Where B is the bias, D is the dark frame, F is the flat field. For the object frame O, the subscript P means processed.

The errors introduced by failing to properly calibrate data are not easy to see. Take the example of a light curve for the Delta Scuti-type star YZ Bootis (see Figure ??). While the errors are on the order of 1%, the error depends the position on the detector.

are data for unprocessed and processed images.

D.2 Bias Frames

A zero second exposure, with the shutter closed, does not return a blank image. Every image has a bias level regardless of image type or exposure length. This bias level is spatially dependent—each pixel will vary slightly from an average, pedestal value. The pedestal is often set at 100-200 analog-to-digital units (ADUs, or counts), but higher values aren't excluded: both the Fairchild 3041 at West Mountain Observatory and the SITE detector used at ROVOR have their pedestals at around 2300-2400 counts (this decreases the dynamical range of the detector but has no other real effect).

When reading out a CCD, statistically important structures may also appear in addition to a pedestal value. The image from a zero second exposure (with the CCD shutter closed) may show lines of slightly higher values, among other patterns (refer to Figure ??). Called a bias or zero frame, it is generally the first of the calibration processes undertaken. Upon powering up and cooling down a detector, the structure of this and the dark frame (explanation forthcoming) often change, and the pedestal value is also not consistent from night to night. Taking bias and dark frames is, then, a nightly necessity.

When reading a pixel's voltage, slightly varying electronics can register small, negative voltages. To ensure that no negative values are assigned positive ADUs, a reference voltage is sampled and added to each pixel's voltage. This adds a certain average DN to each pixel after analog-to-digital conversion, with a Gaussian spread about the average due to the analog-to-digital converter (ADC) flooring command¹. This is the pedestal level, with small pixel to pixel variations, and a pedestal value as the average of the entire bias frame.

The ADC clipping of voltages into DN's introduces uncertainty called read noise. The read noise RN is equal to the standard deviation of a bias frame. Combining images together constrains the noise from a measurement. Methods of combining which are valid include: mean, median, mean with the minimum and maximum values per pixel clipped, and mean with values above and below a certain standard deviation clipped. For calibration frames, a rejection technique is usually needed desired. For all these frames, the ROVOR group has chosen to combine frames using the mean after the minimum and maximum points are rejected..

¹The ADC takes a voltage measured with great precision and clips it to some integer multiple of G^{-1} . If the gain is set at $10 \text{ e}^-/\text{ADU}$ and 321 electrons are in the well when measured, 32 counts will be output. The ADC will also output 32 counts for 325, 327, and 329 e^- in the well. This is the source of the read noise uncertainty in a bias frame.

D.3 Dark Frames

The two primary sources of electrons in a pixel's well are photoelectrons and thermal electrons. Motions of the silicon within its lattice structure will excite valence electrons into the conduction band, which may reach the well before recombination (this is called dark current). Once in the well, a thermo-electron is indistinguishable from a photoelectron. Efforts to remove this effect do have uncertainty, but increase the overall accuracy of an exposure.

Silicon thermal motion is greatly reduced as temperature is lowered. Around -80°C , the dark current (over the length of an exposure) is negligible, and the dark frame is not needed. Most large observatories fill the CCD housing unit (called a dewar) with liquid nitrogen. This has a boiling point of -196°C , and so an onboard unit will heat the CCD to a specified, higher level for temperature stability (Howell 2006).

ROVOR uses a simple thermoelectric cooler, which is rated to drop the temperature down by 60°C from the ambient². Usual operating temperature during the winter is -40°C , and during the summer is set to -30°C . While these values still necessitate the dark current removal, the cost and hassle of liquid cooling—especially for a remote site—are not worthwhile. Taking dark frames at many different temperatures, a graph of average dark counts per pixel per second versus temperature shows an exponentially increasing curve. For the setup in Delta, as shown in Figure ??, the average dark current doubles with each 5.08°C increase. When the ROVOR detector is cooled to -40°C , typical dark current is $0.13\text{ ADU pix}^{-1}\text{ second}^{-1}$. With the gain set at about $4.5\text{ e}^{-}/\text{ADU}$, this corresponds to around $0.585\text{ e}^{-}\text{ pix}^{-1}\text{ second}^{-1}$ into each well³.

To determine the effect of dark current, one exposes the CCD with the shutter closed for nonzero

²www.flicamera.com/proline/index.html

³Dark current is usually reported in pico-amps/cm² in journals and CCD specifications to put all detectors on the same spatial level. For $24\text{ }\mu\text{m}$ square pixels, this gives 1.622×10^{-2} pico-Amps cm⁻². As this value is usually quoted at room temperature (20°C), the ROVOR detector expects 66.7 pico-Amps cm⁻², which is exactly in the range the SITE specifications provide.

time. After subtracting the bias frame, what remains is called the dark frame. Usually, pixel-to-pixel variations are evident and of much greater importance here than with bias frames. Most salient when studying a dark frame is that some pixels exhibit very large dark current with respect to its neighbors. Called hot pixels, these may significantly misconstrue the signal received in an object frame if not corrected. As with zero frames, dark frames are additive in nature; the correction is to subtract a dark exposure of the same length as one's object frame. Dark exposures are usually long to obtain good statistics on low-level dark current. The dark image is scaled to the object frame's exposure during calibration.

D.4 Flat Fields

The most difficult calibration frame to accurately obtain is the flat field frame. By its name, its purpose is to flatten (remove) nonuniformities in the chip's response. As Craig Mackay (1986) noted, "The only uniform CCD is a dead CCD". Illumination nonuniformities are the largest errors fixed by flat fielding. As one looks away from the center of the image, the amount of light per pixel decreases. This is called vignetting, or natural illumination falloff, and the change in illumination between the center of the chip and portions near the edge is often several percent. The natural falloff may be approximated by a $\cos(d)$ function (Stubbs & Tonry 2006), but becomes more complicated when using filters which interrupt the natural light pattern before it reaches the detector. Interference patterns and internal reflections from the glass filters alter this natural falloff slightly, especially at the edges of the frame.

In addition to illumination differences on the CCD, each pixel responds slightly differently to light due to manufacturing variations and photolithography inaccuracies (Janesick 2001; see appendix __). One can find these changes, called the quantum efficiency (QE) per pixel, by removing the illumination response (see appendix __ for details). For the SITE SI-003A CCD, the average QE

is found empirically to vary less than half a percent across the CCD. Serious defects are not very common in modern detectors, but may be introduced by damage from very high-energy cosmic rays. In one such region in the ROVOR chip, light sensitivity has been reduced by half (see Figure ??). Without adjusting for this, if one measures an object that lies on this part of the detector, serious errors follow. As a typical flat field for ROVOR exhibits between seven and ten percent variation across a chip, compared to the half percent QE, achieving 1% photometry primarily depends on determining the correct illumination response of the telescope.

D.4.1 Illumination Correction

Correcting for the illumination is a complicated process. The response of the telescope to uniform illumination greatly depends on the filter being used. Most filters of the same kind do follow a general pattern for illumination, and only differ by slight (but still significant) amounts. Traditional Johnson-Cousins filters (*UBVRI*) and Sloan filters (*u'g'r'i'z'*) are uniform glass filters and show similar patterns. These do not compare to Stromgren filters (*vby*, but not *u*) or the newer Johnson-Cousins interference filters, which are multiple-cavity interference filters and have thin layers of dielectric and metallic materials deposited to isolate transmitted wavelengths (Crawford 1987).

Achieving uniform illumination is a fiction in practice. Uniform illumination requires constant light intensity over the entire aperture of the telescope. The problem becomes intractable⁴ when

⁴ Take ϕ as the intensity of light on the detector, I as the intensity of a light source above the telescope, QE as the quantum efficiency of the detector, and H as the transfer function which describes the way the light is mapped onto the detector. Note that these all depend on wavelength λ . Also, assume $I(\lambda)$ is uniform across the field of view (for an ideal flat field) and ignore any polarization; if we take (x,y) as the detector positions and (x',y') as the field positions, we may write $\phi(x,y,\lambda) = \int I(\lambda) * H(x,y,x',y',\lambda) dx' dy'$. For each pixel, then, the response, or signal, S is $S(x,y) = \int \phi(x,y,\lambda) * QE(x,y,\lambda) d\lambda$. The field ϕ on the detector is wavelength-dependent, and it reacts with the QE differently than would another ϕ function. In addition to being of lower amplitude, I for the night sky is a separate function for the twilight sky and an illuminated screen inside a dome. Any illumination I without the same spectrum

one considers that the illumination and quantum efficiency field corrections depend on wavelength⁵. The spectrum used to make a flat field frame (be it with lamps or the lit sky) is not the same spectrum one observes in either the night sky or one's objects (which are superimposed on each other).

In choosing the method of correcting for the response of the system, the observer cannot satisfy the demands of both effects, and cannot reproduce the effect of each of his objects for each pixel on his detector. Especially grievous are strong variable sky emission lines in the far red through the IR range, which dramatically change the night sky contributions per wavelength (Joyce 1992) and plague the goal of 1% accurate photometry. New methods of calibrating promise much greater precision: The LSST Science Collaborations (2009) outline combining lamp-lit dome flats with images from photodiodes that have been calibrated to 0.1%. This method will help to establish the illumination response and quantum efficiency of their detectors with less uncertainty than known before. This does not solve the intractable problem of flat fielding— it merely reduces the final error.

D.4.2 Quantum Efficiency

Due to manufacturing nonuniformities, interpixel quantum efficiency can vary on the order of $\frac{1}{2}\%$ in the ROVOR chip. Intrapixel variations (sensitivity changes within a single pixel) in some CCDs (specifically older models) and at some wavelengths are nearly 30%⁶ (Mackay 1986). Though im-

as the night sky will produce a dissimilar signal signature and a different ideal flat field. Matching the night sky is not the solution— stars, galaxies, and other objects contribute their own ϕ functions— and no amount of matching with composite spectrum will produce a flat field for all these objects simultaneously and be free from errors.

⁵From Mackay (1986), nonuniform response of the detector is inherently wavelength-dependent as absorption length is wavelength-dependent. See ___ for absorption length.

⁶The source of this is a dissertation with Mackay as the graduate advisor for the student. The text of it doesn't seem to exist, and so the method and arguments behind the statement are not known. This author assumes that the detector(s) used to produce the result was front-illuminated given the date of the source. Different portions of the gate

provements have been made in manufacturing silicon wafers and forming CCDs, intrapixel variations are still expected on at least the same order as the interpixel variations.

Both the illumination and quantum efficiency corrections are multiplicative, meaning that one divides, pixel-by-pixel, an object frame by the effects of both illumination and quantum efficiency. To obtain the QE response per pixel, one takes a flat frame (which, for this, doesn't need not be obtained under uniform illumination) and performs a high-order polynomial fit to the surface. This process is fully described in appendix __. Just as with illumination corrections, QE is on a per-wavelength basis and needs be corrected for each filter one uses. While quantum efficiency does not usually change more than half of a percent at ROVOR, some damaged areas do show significant difference in responsivity (see Figure ??). Usually, quantum efficiency and illumination effects are bundled in one inclusive flat fielding. Disentangling the effects does present some danger- each calibration frame has inherent error that cannot be avoided, and dividing by two (or more) separate frames enhances these errors (Stubbs & Tonry 2006).

sensitivity is either the result of manufacturing defect or damage by a very energetic cosmic ray.

D.4.3 Flat-Field Stability

The importance of an accurate flat field is preeminate in the data calibration process. Due to cloud coverage or time limitations, obtaining twilight flat fields in all filters every night observed is not reasonable. This is especially true in remote observing, where the observer cannot precisely determine the night's quality. Also, for a filter wheel containing 6 to 8 filters, obtaining twilight flats in all these filters often limits the number of flats to no more than six in each filter. The limits structure reflect light differently; if tight beams of light go through single features of the polysilicon and SiO₂ making up the front of a pixel, the amount of transmitted light (and hence signal) will drastically change. As ROVOR's CCD is back-illuminated and light does not pass through a gate, these intrapixel nonuniformities may not exist to the same extent.

on the number of frames per filter restricts the statistical accuracy when combining frames (see section ___). Instead, for the past two and a half years, the ROVOR group has limited itself to one filter per night, and often more than twenty frames are obtained in that filter. Each new night focuses on a new filter, and a filter is not repeated until the full set has been run in the interim.

Conventional observer wisdom is that flat fields do not change appreciably over the course of days or weeks (LSST Science Collaborations 2009)⁷. To test this, Johnson V filter flat fields from September 2 to November 5, 2011 were taken from the West Mountain Observatory (WMO) repository. After careful calibration, each frame was divided by the initial frame, and the standard deviation (to measure the error in the frame) was computed for the central 5%, then 10%, and by tens up to the full frame. The standard deviation of the quotient of two flat frames is due to photon counting uncertainty, and primarily not differences in the flat field quality⁸. The data suggests no

⁷Of course, if the illumination is not uniform, either from sky gradients or clouds, the resultant flat field will not accurately represent the response of the detector to light. Obtaining flat fields in consistently cloud-free regions with minimal twilight gradients is essential.

⁸Signal-dominated electronic measurements have variance well described by shot noise, where the uncertainty is equal to \sqrt{N} for N the number of counts per measurement (here, per pixel). Flat frames are usually taken between 25,000 and 40,000 counts, or generally about 50-75% of the full well capacity. From the above, this yields an uncertainty of about 0.0052 in a normalized flat frame. Combining 6 frames together with a minmax rejection decreases this by a factor of \sqrt{n} , where n is the number of frames combined. A typical value for the uncertainty would then be 0.0022 for the pixels within a quality, normalized master flat frame. When measuring the change in a flat field, and dividing one frame by another, the error from the input frames propagates. Take the output to be F_{out} , and the two inputs F_1 and F_2 , such that $F_{out} = \frac{F_1}{F_2}$, and the uncertainties per frame as σ_{out} , σ_1 , and σ_2 . The propagated error σ_{out} is related by $(\frac{\sigma_{out}}{F_{out}})^2 = (\frac{\sigma_1}{F_1})^2 + (\frac{\sigma_2}{F_2})^2$. The ratio of normalized flat frames are expected to return uncertainties of $\sigma_{out} \approx 0.0028$, which coincides very closely with the standard deviation of the inner portions of the F_{out} frame. This author suspects that discrepancies, though small, are due to QE(x,y) differences. Variations in either the illumination flatness or the spectrum (color) of the twilight sky will produce differences in the resultant flat fields from night to night (and with a night, as the sky changes color). This may account for the additional 0.01-.02% error not solely the result of shot noise. As long as the flat frames are taken correctly, the uncertainties in F_{out} are due primarily to photon

trend in the ratio of flat fields with time (just over two months), indicating the relative stability of a flat field. Instead of poorer statistics with only 6-8 flat frames in a filter per night, one is justified in taking many more flat fields in only one filter in a night⁹.

D.4.4 Obtaining Flat Fields

Major observatories usually shine light on large screens (affixed to the interior of the domes) to make their flat fields. They attempt to make the light reaching the telescope aperture uniform, using clusters of projectors and diffusers. Some use special blends of paint on these screens to reflect light more uniformly across wavelengths (Massey & Jacoby 1992). Others worry about the spectrum of the light output by the bulbs in their projectors (many bulbs produce discrete emission lines, a continuum of light is preferable). Many smaller observatories use as a backdrop the sky after twilight or before sunrise, but even here there are variations in procedure. The methods used seem, at times, more akin to rituals than science, and little has been done to compare or reconcile disparate methods. In defense against these dark arts, and with a desire to find a better way to conduct flat fielding in remote observatories, we have investigated other possible methods.

Astronomers using twilight flat frames (or, alternately, those taken before sunrise) often make the assumption that the sky is uniform over the projected aperture of the telescope. (This is not the case, but for smaller aperture telescopes, the change is often small enough that such flats are approximately accurate. Large aperture telescopes do not have this luxury, and instead work to create uniform illumination on domes and screens.) Earlier studies done with a large scale camera have determined that the sky brightness gradient is at a minimum (much less than 1% per degree) just around the zenith during twilight (Chromey & Hasselbacher 1992). This is support for the stability of flat fields.

⁹Taking 24 frames in one filter (with good signal), one reduces the shot noise to almost 0.1%. The greater number of frames allows more minmax pairs (and hence residual stars) to be removed without seriously affecting the averaging effectiveness. This treatment does not focus on ROVOR data, and is only performed on one filter, but there is no reason why different flat fields, properly obtained, would respond differently.

1996). Dr. Joner of BYU, and Dr. Laney, visiting astronomer, have said that CTIO have performed studies (unpublished) which determined the sky 1-2 hours east of the zenith at twilight is sufficiently flat.

For all of this, differences exist between the flat field structure and the night sky background. This is due to either poor execution of twilight flats or gradients across the night sky (especially from the moon or encroaching civilization). The correct response of the telescope should not include any gradient on the sky, but it is difficult to determine whether the region of the sky or the screen produces truly uniform illumination.

An experiment conducted at ROVOR has used a thick, white sheet wrapped around both the secondary-mirror support trusses and the entire aperture to diffuse light when taking flat fields. This method is at least as accurate, if not a little more so, than traditional twilight flats (see Figure ??). However, as the site is remote, an observer must travel down and back to attach the sheet, obtain the flat fields, and then remove the sheet. This is possible only every month or two at best. The general stability of flat fields over the course of at least two months indicates that this is a viable option. Research continues to determine whether these sheet flats will replace traditional twilight flats at ROVOR.

D.5 Combining Images and Master Frames

Each measurement has uncertainty. To constrain this error to its lowest values, one combines many frames of each calibration type to form a "master" frame. This master frame is then applied at each subsequent step. If enough frames are taken, the spread of values for a given pixel will cluster around a mean value. The shape of this distribution is approximately gaussian for good data.

If all calibration frames were good, all of the data points would be used. Real data have error which contaminate these statistics. For instance, a cosmic ray hitting the detector can contribute

thousands of counts to just a few pixels, but where and when it will hit is random. An average of 0.61 cosmic rays per second reach the CCD in Delta and seriously affect an average of 2 pixels. A second error occurs when flat fields are taken using the twilight sky: as the calibration images are obtained, the sky dims and stars become more prominent. Like cosmic rays, stars in the background of flat fields add light to pixels, and these pixels become outliers in the distribution.

Choices for combining images include median, mean, mode, sum, with possible rejection techniques like minmax, sigma clipping, percentage clipping, and various others. Only the most applicable of methods will be explained here; the reader is referred to Massey (1997) for a more complete treatment. Ideal median combining automatically ignores possible outliers (either high or low), but the distribution is not always gaussian in shape, especially for small data sets. Hence, the median value may not represent the true center value as well as the mean. The mode is more useful in scaling flat fields to the same level, as will be discussed, than in comparing the same pixel from multiple images. Especially with small data sets, the most used value may be one far removed from the center by a fluke. The mean is usually a more accurate method, but must have any outliers removed first. Rejecting these high pixels in sets of data is common; to do this in an efficient manner, the highest value from each set of pixels is discarded when combining the images. If the highest value is neglected the lowest value must also be rejected, or the computed mean will be biased lower from the ideal mean. This choice is called a minmax rejection, and more than one minmax pair can be rejected. As the distribution population grows large, the values for the median, the mode, and the mean (with a minmax reject) converge.

D.5.1 Master Bias

A tutorial by Massey (1997) states that 25 is the minimum number of bias frames needed to produce a good master bias frame. Any significant underlying structure should become visible when combining this many images. The master bias B_M is made by taking the mean value of each pixel

from the 25+ bias image taken. Minmax rejection is used, and typically only one pair ($n = 1$) is needed. The equation to produce a master bias frame, for the pixel position (x,y) , is:

$$B_M(x,y) = \frac{\sum_{i=1}^N B_i(x,y) - \min[B_i(x,y),n] - \max[B_i(x,y),n]}{N - 2*n}$$

The functions $\min[B_i(x,y),n]$ and $\max[B_i(x,y),n]$ are the n minimum and n maximum data points $i = 1 : N$ for pixel (x,y) .

At ROVOR, at least 30 contiguous bias frames are taken every night. The standard deviation of the master image is usually less than one count. A complication was discovered when experimenting with nightly bias stability. The average value of a bias frame—the pedestal level—varies with the time of night and is not consistent from night to night. The pedestal tends to increase through the night, and an additional 40-50 counts by the end is common¹⁰. It is thought that due to decreased power use, a floating electrical ground relaxes as the night continues. The pedestal level is the result of the analog-to-digital converter sampling a bias voltage (which biases every pixel's voltage upwards; hence the name bias frame. See appendix _), and the bias voltage changes with the ground level. To temporarily fix the reduction problem, three additional bias frames are taken every twenty minutes to resample the pedestal. The average of these three frames is the new pedestal, and the change (positive or negative) is added to the original master bias to produce a master bias every twenty minutes of observations. Non-bias frames are calibrated with the master bias closest in observation time to them. A floating pedestal value is largely immaterial for object frames, as aperture photometry subtracts off the background level about each object measured. For the other calibration frames, however, using a master bias far removed in time can seriously reduce calibration accuracy¹¹.

¹⁰Compared to the base pedestal level of 2300, adding 50 counts seems insignificant. However, this is more than either the sky contributions or the average dark current per pixel for a typical 60 second exposure.

¹¹The bias drift in a night may seem insignificant compared to the pedestal, but it can almost match the dark current per pixel in 300 seconds (at ROVOR with the CCD at -40°C). Using a poorly-chosen master bias can double the true dark current in the master dark frame. This will reduce the object frame's pixel values by around ten counts, but

D.5.2 Master Dark

For dark frames, between 5 and 10 frames (with very long exposures) are needed to produce a quality master dark. Each dark frame should be at least as long as the longest exposure taken that night (Massey 1997). A common practice has been to take a series of dark exposures at each exposure length at which object frames were taken, and creating a master frame for each length. The bias signature is found in the raw dark frame, and so often the bias and dark effects are lumped into one master frame for each exposure length. The bias/dark master is subtracted directly from the appropriate object frames. This reduces the number of calibration steps and, as each calibration step contains error, slightly reduces the uncertainty of the reduction. While not bad in general, very short dark exposures (generally less than two minutes for ROVOR) heavily sample the uncertainty of a bias frame. Consider that the average dark current (at -40°C) is about 0.13 ADU/pixel/second, and the uncertainty in the background of a good (raw) bias is about 4 ADU. If one wants the uncertainty of the bias frame to be less than 1% of the average dark current, one must expose for around 3000 seconds. One is hard pressed to take as many bias/dark images for each exposure time as one would with separate bias frames (which only require read-out time). Fewer available frames means a less accurate bias signature and a less accurate master dark. This is to be avoided. Current practice at ROVOR is to take seven long exposures (five minutes each), and combining them for the same effect.

Where P_i is the i^{th} processed frame in the set, the master dark D_M is formed first removing the master bias and then combining the resultant frames together:

$$D_{P,i}(x,y) = D_i(x,y) - B_M(x,y)$$

$$D_M(x,y) = \frac{\sum_{i=1}^N D_{P,i}(x,y) - \min[D_{P,i}(x,y),n] - \max[D_{P,i}(x,y),n]}{N - 2*n}$$

not quite uniformly. Then, recall that flat fields are multiplicative while bias frames are additive. Subtracting off an incorrect bias will change the flat field signature by at most about 0.2%. These changes are small, but, if not dealt with, they increase the necessary error budget for each affected object frame.

Dark frames are combined not only to improve the accuracy of the master dark, but also to eliminate cosmic ray events. Due to low probability of these events, and short (if non-zero) exposure length of biases and flats, only darks are significantly affected by cosmic rays. An average cosmic ray affects two pixels, and in an average dark frame of 300 seconds, the probability of a cosmic ray affecting a single pixel on the detector in Delta is about 0.035%, or about 1.15×10^{-4} percent second⁻¹. Therefore, about three times should two events affect common (x,y) pixels in 7.5-minute exposures¹². Most cosmic rays are of lower energy, and only produce several hundred additional counts above the dark current. When mean combining pixels with unrejected cosmic-ray hits, the measured mean is only a couple hundred counts above the ideal mean. The probability of three events on the same patch (in seven five-minute exposures) falls to the remote 0.0015, or occurring in only 3 out of 2000 sets. To avoid the chance that double cosmic-ray hits will bias some master frame pixels much higher, take n=2 rejection pairs when creating the master dark¹³.

D.5.3 Master Flat Frame

Proper flat fielding requires not only enough signal to reproduce the illumination response, but to distinguish small QE differences between pixels. As the response is color-dependent, this high-signal flat field needs to be found in each filter. Exposures need at least 20,000 counts per pixel,

¹²These values were found by master bias and master dark subtracting two raw dark frames. Any peaks (not duplicated in the two frames) far above the background noise were taken to be cosmic rays. This does assume linear, repeatable dark current affected only by random shot noise. In the “reduced” dark frames, four negative “peaks” were visible, corresponding to four strong, double cosmic ray events in that set; the master dark was created using only one minmax rejection, so the lower of the two double events was recorded and severely increased the mean. This squares well with the three expected doubles, though it is likely that other cosmic-ray-induced negative peaks were present, but were of lesser magnitude and were therefore difficult to see.

¹³The sample for this was 10% of a dark frame, and so the statistics might be slightly off. Until recently, only 6 frames were used in the combination, meaning that the number of double events is only about two per dark set (which, based on observation is about right). Still, using a method other than n=1 minmax rejection is suggested.

and are best above 40,000¹⁴. From each raw flat F_i the master bias is removed as is the master dark (scaled to the flat exposure length).

$$F_{P,i}(x,y) = F_i(x,y) - B_M(x,y) - \frac{\text{exptime}(F_i)}{\text{exptime}(D_M)} * D_M(x,y)$$

Whether using the twilight sky (which quickly dims), an illuminated dome or screen, or any other method, the intensity of light per frame is usually never constant. The flats must first be scaled to each other before combination. Background stars will push the mean value of a flat image upward, and as crowded star fields will strongly push the median value high, the mode is the preferred scaling tool. In the manner of Mackay (1986), to normalize each frame, multiply each frame by a ratio of the modes of the current frame to the highest-signal raw flat (usually the first).

$$F_{scaled,P,i}(x,y) = \frac{\text{mode}(F_{P,1})}{\text{mode}(F_{P,i})} \times F_{P,i}(x,y)$$

ROVOR uses twilight exposures for these frames, and, after the sun is below the horizon, stars increasing become visible in images. Between these exposures, then, the telescope should be jogged several arcseconds in any direction so that no star covers the same pixels in consecutive exposures¹⁵. Scrolling through the flat images, the stars will appear to move across the field of view. Minmax rejection should remove these stars when combining frames, but at times two stars will lie on the same pixels in two frames. As long as many exposures (11+) are taken in a filter, n=2 pairs should be rejected. For very large sets (17+ or so) rejecting n=3+ pairs is preferred, especially for twilight flats taken in low-light levels or when looking at a crowded field.

$$F_{scaled,P}(x,y) = \frac{\sum_{i=1}^N F_{scaled,P,i}(x,y) - \min[F_{scaled,P,i}(x,y),n] - \max[F_{scaled,P,i}(x,y),n]}{N - 2*n}$$

When performing calibrations, the object frame is, at the last, divided by the master flat frame for that filter. While not strictly required for differential aperture photometry, when dividing through by the response of the system (the master flat), it is best to divide through by a factor of one. As

¹⁴Higher is better until nonlinearity begins. The current setup at ROVOR has nonlinearity starting no sooner than 60,000 counts, but to be safe do not begin twilight flats until the average pixel is below 55,000 ADU.

¹⁵It is important that the telescope is tracking when exposing flat fields, or much larger swaths in a frame will be unusable.

each pixel does not respond the same, we seek to divide the object frame by an image with a mean value of one. The processed, combined flat $F_{scaled,P}$ is divided by its average to form the master frame for that filter.

$$F_M(x,y) = \frac{F_{scaled,P}(x,y)}{\sum_{x=1, y=1}^{N_x, N_y} F_{scaled,P}(x,y) / Npix}$$

Above, $Npix$ is the total number of pixels in an image, while N_x and N_y are the number of pixels in each dimension x and y . When combining these frames, no frames are used which are not like the others, or those with unusual, large-scale gradients across the image (which is a sign of improper illumination on the aperture).

D.5.4 The Processed Object Frame

Bias and dark subtraction are carried out before dividing by the master flat field. In removing the response of the detector, order is specific. To adjust an astronomical image to a detector-independent state, calibration is in the reverse order as the error was incurred. As light reaches the telescope, vignetting skews the proportions of light reaching each pixel. Filter throughput and transmission, along with QE considerations, next affect the original light. While the photoelectrons are collecting in the well, dark current confuses the signal from the sky and that from the detector. Lastly, during read-out, analog-to-digital conversion necessitates the addition of a bias level to each pixel.

Returning along this path, the master bias frame $B_M(x,y)$ is first removed, followed by the scaled master dark $D_M(x,y)$. Next, CCD QE nonuniformities $QE_M(x,y)$ are peeled from the images, and lastly the illumination response $I_M(x,y)$ (containing vignetting, filter transmission, and dust diffraction) is taken out. The ordered process will be:

$$O_{P,i}(x,y) = \left\{ \frac{[O_i(x,y) - B_M(x,y)] - \frac{exptime(O_i)}{exptime(D_M)} D_M(x,y)}{QE_M(x,y)} \right\} / I_M(x,y)$$

Due to the associativity of the master bias and dark frames, and the commutivity of the quantum efficiency and illumination frames, the bias and dark operations can be switched, and the QE and

illumination corrections can also be done in any order. As the QE and illumination corrections are bundled into one frame by most observers, only division by the composite master flat is needed.

As before, the bias and a scaled dark are subtracted off the raw object frame O_i , and then is divided by the master flat created above.

$$O_{P,i}(x,y) = \frac{O_i(x,y) - B_M(x,y) - \frac{\text{exptime}(O_i)}{\text{exptime}(D_M)} \times D_M(x,y)}{F_M(x,y)}$$

In most cases, object frames are not combined by the mean with any rejection. If images are combined, it is usually by the sum function, where each pixel (x,y) in all i frames is added together.

$$O_{sum,P}(x,y) = \sum_{i=1}^N O_{P,i}(x,y)$$

None of the calibration frames need to be aligned (they are detector-dependent), whereas images of celestial objects (which are sky-dependent) must be aligned before their combination. When producing public-friendly images, multiple frames are often combined using a median filter (with a median or mode scaling for the sky background) or by taking the mean with a minmax rejection. These options reduce the likelihood that cosmic rays or unfixed hot pixels will deface the final product.

D.6 Data Reduction Afterwards

This primer should serve as a good basis with which to start astronomical reductions. A last note of caution: data will not be perfect. All data is ugly and has problems, but ugly data is real data (Massey & Jacoby 1992). Do not let some ugliness keep you from keeping usable data. Error is inevitable, and perhaps more so from humans than on the detectors. Know the sources of error, remove all that you can, and honestly account for any that you cannot (Janesick & Elliot 1992).

Appendix E

Gradients on the Sky

E.1 Twilight Sky

When obtaining flat fields, most smaller observatories will use the dimming sky at twilight for a uniform source of illumination. While powerhouses like LSST will be/are using projectors and screens to flatten their data, their analysis (LSST Science Collaborations 2009) shows that the necessary process is far more than any small observatory could afford, and is especially unsound for a remote site. And they are right to fear using the sky. A quick glance before dusk should convince even the inastute observer that sky brightness, as well as color, is not constant across the sky. However, for small-field imagers (≤ 30 arcminutes), the change across the region of interest can be small enough to be called flat. Joneer (2011) indicates that CTIO has studied the sky gradient at twilight, and they hold (unpublished) that 1-2 hours away from sun's direction and down from the zenith produces the most flat regions of the sky. Another source, Chromey & Hasslebacher (1996), finds that, at KPNO, being 90° away from the sun (with declination equal to the observer's latitude) has the smallest gradients. In an effort to conform to these results, the ROVOR group has decided to take flat fields 5° away from the zenith, opposite the sun. An estimation from the KPNO

data suggests that this location will have a gradient of less than one fifth of a percent per degree, or an accuracy of about 5

It should also be noted that sky gradients, especially from the moon during observations, should not be removed by a surface fitting such as is done with flat fielding. The moon contributes additional flux and does not affect the response of the telescope and detector to light. When performing photometry, the annulus about each object should be able to remove the extra (additive) sky counts. Making a fit of the additional sky background and dividing each frame through by this will not preserve the response of the system. Even subtracting the fit is suspect because during photometry the sky background (and hence the error from the sky) will be underestimated.

E.2 Moon-induced Sky Gradients

Modeling the pattern of the moon gradient is a complicated process- for a complete description, the modeler must take into account the above-atmosphere moonlight spectrum, the fraction of the moon lit as well as the intensity of that portion¹, molecular absorption, atmospheric refraction and the Rayleigh scattering pattern produced for a specified (non-zenith) sky position², the altitude of the observer, the humidity, and then the telescope, filter transmission, and perhaps even the CCD quantum efficiency. Other factors, such as Mie scattering for aerosols when Rayleigh scattering breaks down, play a small but still important role in the model completeness. Due to the extreme difficulty of the problem (and the small benefit from such lengthy exercises), an observational route was taken to ascertain the effect of a bright moon on the sky. A good theoretical primer is presented in Krisciunas & Schaefer (1991). It should be noted that, while these results can be applied to other observatories at other sites, differences in location and atmospheric conditions will

¹One may also add, if they are so inclined, solar intensity fluctuations and the changes in earth-sun and earth-moon distances with time.

²Upper atmospheric conditions and the conditions of inversion layers also play a large part of this analysis.

Sky Brightness at CTIO (mag arcsec⁻²)

<i>lunar age (days)</i>	<i>U</i>	<i>B</i>	<i>V</i>	<i>R</i>	<i>I</i>
0	22.0	22.7	21.8	20.9	19.9
3	21.5	22.4	21.7	20.8	19.9
7	19.9	21.6	21.4	20.6	19.7
10	18.5	20.7	20.7	20.3	19.5
14	17.0	19.5	20.0	19.9	19.2

affect the applicability of the results.

Sky brightness is a strong function of both wavelength range and lunar phase. Data from Alistair Walker in an NOAO newsletter in 1987³ gives the following sky brightness values for CTIO:

Note that lunar day is zero at the new moon phase. Also, while the sky is intrinsically (moon-free) brightest in I and faintest in U, the moonlight affects the U filter the most and the I filter the least. While the data is qualitatively interesting, of more importance is how the scattered light falls off with angular distance from its source. Sky brightness aside, one must also reduce the gradient across one's chip from the moon. The goal of this side project is to determine at what angular distance the differential effect across a region of interest falls to 1% or less. Then, in observations at ROVOR, the increase in noise from one side of that region to the other should remain below this threshold.

On March 18, 2011 (UT), observations at ROVOR measured flux of the sky in B, V, and I filters and at intervals away from the almost full moon. The central position of the orb was just a few degrees above the celestial equator, so not as high in the sky as it could get (on average the moon's maximum declination is about 23.5°). At different declinations, the minimum safe angular ob-

³Haven't found the newsletter online, but several documents cite this. For example, see Bolte's lecture notes at www.ucolick.org/~bolte/AY257/s_n.pdf.

	<i>B</i>	<i>V</i>	<i>I</i>
<i>a</i>	678.3	2007.5	2425.3
<i>b</i>	15.2	15.14	15.1
<i>c</i>	-1.38	-1.42	-1.51
<i>d</i>	101.1	31.3	40.5

Table E.1 Coefficients for Sky Brightness $B = a(x - c)^c + d$ for Johnson-Cousins Filters

serving distance will be different, but it is hoped that the change is very small. The measurements were taken with the moon near the meridian to keep observations at approximately the same air-mass. Photometry was performed on a uniform region at the center of each frame, and due to its constancy no flat fielding was performed. Distance measurements were calculated using the lunar ephemerides in the *Astronomical Almanac*⁴ of the U.S. Naval Observatory and collaborators. As the moon is about $\frac{1}{2}$ of a degree in diameter, the separation values were taken to be from the edge, and not the center, of the moon's disk.

Frames were bias and dark subtracted before photometry. The sky brightness for a frame's angular distance (the central pixels) was averaged from the central region that avoided stars. Error in the angular distance was estimated at a maximum of $\frac{1}{6}$ arcminute, and error in the sky brightness as typical shot noise. Using a non-linear least-squares fit in Mathematica, the best-fit function for each filter was found to be of the form $a(x - b)^c + d$. Coefficient solutions are shown below:

Note that, as expected from the newsletter by Walker, the moonlight dies away quickest in redder filters⁵. A quick calculation shows that a 1% gradient across the entire current ROVOR detector

⁴See the online version at <http://asa.usno.navy.mil>

⁵The derivative is equal to $ac(x - b)^{c-1}$, and so the amplitude a does play a role in the slope, and a increases with the filter's reddening. So while the decay depends on both a and c , the decay parameter c does increase with redness. Due to this, the V filter moon gradient is actually the last to level off below a slope of 1%.

($\sim 24'$) is reached at 54° from the full moon center. As the outer portions of a detector have less accurate flat fielding (see D.4.3), using an inner portion is usually preferred. At half the frame's pixels (75% of each side, or $18'$), the desired level of moon-gradient flatness is reached in 50° . Before these values are taken and used, note that on the night observed: the moon was not quite full; also, the lunar declination was low, and at its highest points the gradient will be spread to larger values. Expect that 60° or slightly more will safely secure 1% or less gradient from the moon.

Appendix F

All-Sky Photometry

On photometric nights (meaning without any moonlight interference or thin, high clouds), all-sky photometry can be performed. Due to unique sensitivities at various wavelengths, each telescope and detector combination will record different magnitudes for the same objects. Standardization, or all-sky photometry, allows the observer to compare results from his system to those from other observers at other telescopes. To understand the methods of this type of photometry, an introduction into extinction is first needed. The explanations and formulas are heavily borrowed from A. A. Henden and R. H. Kaitchuck in their 1982 book *Astronomical Photometry*. The reader is strongly encouraged to reference this book for a more in-depth discussion. Further, Cameron Pace (2010) provided a tutorial, albeit in a different (and no less valid) manner, and the reader is also encouraged to reference his work with standardization at ROVOR (some caution should be exercised with the ROVOR site).

F.1 Extinction

As light passes through a medium, light is irretrievably scattered from its original direction. In mathematical form,

$$dF_\lambda = -F_\lambda \times \alpha_\lambda \times ds$$

Above, α_λ is the absorption coefficient, or the fraction of light at wavelength λ lost through a medium per unit of length (depth). F_λ is the flux of light (usually in terms of energy) of a certain wavelength λ per unit area (perpendicular to the incident direction of light). As α_λ and F_λ are always nonnegative, the negative sign in front of F_λ indicates more loss of flux with distance into the medium. Lastly, the terms dF_λ and ds are differential forms of flux and depth, respectively. Solving the above,

$$\frac{dF_\lambda}{F_\lambda} = -\alpha_\lambda \times ds$$

Integrating this over a length s through the medium,

$$\ln\left(\frac{F_\lambda}{F_{0\lambda}}\right) = -\int_0^s \alpha_\lambda ds$$

Here, $F_{0\lambda}$ is the flux above the absorbing region. This leaves

$$\frac{F_\lambda}{F_{0\lambda}} = \exp\left(-\int_0^s \alpha_\lambda ds\right)$$

Often, for convenience, the definition

$$\tau_\lambda \equiv \int_0^s \alpha_\lambda ds$$

Where τ_λ is called the optical depth through a path of length s . For flux F and magnitude m for the two objects or situations 1 and 2, recall that the magnitude relation in astronomy is,

$$m_1 - m_2 = -2.5 \times \log_{10}\left(\frac{F_1}{F_2}\right)$$

Using the above two relations, and the fact that

$$\log_n(\exp(x)) = x \times \log_n(e)$$

We can rewrite the magnitude relation as

$$m_{0\lambda} = m_\lambda - 1.086 \times \tau_\lambda$$

The term $m_{0\lambda}$ is the magnitude outside of the atmosphere (corrected for atmospheric extinction). Using y as the thickness of the atmosphere at the zenith directly overhead (called one airmass), the angle z of an observation from the zenith relates to the distance through the atmosphere by

$$\cos(z) \cong \frac{y}{s}$$

Or,

$$s \cong y \times \sec(z)$$

At zenith angles much greater than 60° , the atmosphere is not a plane slab, but bends with the curvature of the earth. Also, density of the atmosphere is not consistent with altitude, and a viewing angle close to the horizon will not go the same amounts of a certain density as will a view close to the zenith. For angles much beyond 60° (corresponding to an airmass s of 2), the approximation $s = \sec(z)$ begins to overstate the value for the airmass. A more correct form is from Hardie in 1904, where the airmass is again approximated as a polynomial series in $(\sec(z) - 1)$:

$$s = \sec(z) - 0.0018167 \times [\sec(z) - 1] - 0.002875 \times [\sec(z) - 1]^2 - 0.0008083 \times [\sec(z) - 1]^3$$

The above statement is accurate to about 10° from the horizon (an airmass of almost 7).

Taking the simplest approximation and differentiating both sides (with constant zenith angle),

$$ds = dy \times \sec(z)$$

Rewriting the equation for τ_λ , we get

$$\tau_\lambda = \sec(z) \times \int_0^y \alpha_\lambda dy$$

Which reduces the above magnitude relation to

$$m_{0\lambda} = m_{\lambda} - k'_{\lambda} \times \sec(z)$$

The integral above is called the principle extinction coefficient k'_{λ} (in units of magnitudes per airmass) and varies from night to night (and even within a night) depending on atmospheric conditions. The figure below is an experimental extinction curve. Traditional units for astronomy (magnitudes/airmass) are used instead of a fraction of flux scattered.

F.2 Scattering Theory

Extinction is an inclusive effect containing both scattering and absorption of light. Scattering by particles is modeled by spherical particles (not a very good approximation for complex molecules). Called Mie scattering, this effect is best modeled numerically. For radius r of the particle sphere, if the wavelength $\lambda \gg 2\pi r$, we may approximate Mie scattering by Rayleigh scattering. The form for Rayleigh scattering contains a factor of $\frac{1}{\lambda^4}$, meaning that, for constant particulate size, shorter wavelengths are preferentially scattered. This applies primarily to atoms and diatomic molecules. Larger particles (relative to the wavelength) require a more rigorous Mie theory calculation, where scattering goes as λ^{-1} or λ^0 (constant with respect to wavelength). Primary among these are aerosols, which are fine solid or liquid particles suspended in the atmosphere (examples are dust, pollen, and clouds). Other major effects of extinction include strong absorption features near resonance frequencies, especially for O_3 (in the near-UV) and water vapor (in the near-IR). Below, as a further clarification of the previous graph, components of extinction are shown by dotted lines, while the composite effect is represented by the solid line. (Note: the near-IR regime is plagued with a great many features, but mainly in emission rather than in absorption.)

F.3 Extinction with Filters

Astronomical observations frequently use filters to observe only a portion of an object's spectrum. Filters are often made to either key in on a portion of a standard object's spectrum or to avoid night-sky emission features. To correct for the atmosphere when using a filter, an extinction coefficient is found for that certain filter. For the Johnson V filter, for example,

$$m_{0V} = m_V - k'_V * X$$

We start now to introduce the common terminology of X for the airmass instead of s or $\sec(z)$. To determine the value for the extinction coefficient one plots the instrumental magnitudes m_V for a standard star versus airmass X . Typically, as the magnitude curve becomes nonlinear with X much past a value of two (note this in the graph below), only values with an airmass between one and two are included in the fit (an observer may choose to extend his observations just into the two airmass territory).

The slope of the least-squares fit is the solution for k'_V . For accurate fitting, measurements across a range of airmasses are required. There is jeopardy in failing to take observations at the limits of acceptable X values (as close to one and as close to two): measurements carry with them inherent error, and the best-fit slope for one standard star will not likely reflect the true first-order extinction coefficient for that filter. Observe below both a graph showing the best fit line for one star, and then a table containing this and ten other stars' best-fit slopes.

Averaging several slopes together will yield a more accurate extinction coefficient. This gives the value for $k'_V = 0.2319654$. Here, a note should be given concerning which values to average together and which to avoid. Using graphs and intuition, some values, if any, may reasonably be rejected. For example, extinction is not a function of magnitude. The above standard stars have standard magnitudes, and if we plot the first-order extinction coefficient versus the magnitude of the star, no pattern should be discernible. This is not what happens, as shown below.

The last three points are out of the norm by showing a downward trend in the extinction coef-

ficient. It is possible that these points are just fluctuations about the mean extinction, but the probability that this random is extremely small. Excluding these points, a more correct value is $k'_V = 0.2391886$. Caution should be exhibited, for accepting this may bias the solution away from the true value. Not excluding poor data points will also skew solutions, so with care exclusions should be made.

Wider, broadband filters (30+ nm equivalent width) generally provide greater signal-to-noise of continuum light; while agreeable for many general studies, extinction can change significantly over the course of the filter. To a first approximation, a single extinction coefficient works. For more accurate results, one must take into account the slope of the extinction curve and the relative shape of the object's spectrum.

F.4 First-order Extinction

The first-order extinction coefficients which are needed, for the Johnson-Cousins filters *UBVRI*, are $k'_V, k'_{BV}, k'_{UB}, k'_{VR},$ and k'_{RI} . Combination terms, here also called cross terms, follow the form $k'_{BV} = k'_B - k'_V$. We may write, using first-order terms for each filter,

$$v_0 = v - k'_V * X$$

$$b_0 = b - k'_B * X$$

$$u_0 = u - k'_U * X$$

$$r_0 = r - k'_R * X$$

$$i_0 = i - k'_I * X$$

Above, u is the instrumental magnitude and u_0 the magnitude corrected in the U filter for atmospheric extinction (effectively placing it above the atmosphere), and so on for the other filters. Due to slight processing differences (Joner 2011), it is not always true that $(b - v)_0 = b_0 - v_0$. Instead, the last four of the above equations are often found using:

$$(u - b)_0 = (u - b) - k'_{UB} * X_{ave}$$

$$(b - v)_0 = (b - v) - k'_{BV} * X_{ave}$$

$$(v - r)_0 = (v - r) - k'_{VR} * X_{ave}$$

$$(r - i)_0 = (r - i) - k'_{RI} * X_{ave}$$

One needs to average the airmass X between observations in separate filters to find the cross-term coefficients above. As airmass changes rapidly as z approaches higher values, observations in adjacent filters should be next to each other to avoid any large changes in X . For this, the Palindrome sequence is then used: for the above filters, with the letter representing an observation, the sequence is *UBVRIIRVBU*, and repeated for set for each object one observes.

Integration times should be long enough to average out small atmospheric changes, but short enough that airmass does not appreciably change over the course of a single observation. Usual upper and lower limits are 30 to 240 seconds (). ROVOR has the additional complication that its telescope does not currently track accurately enough to allow for long exposures. Either the encoders on the mount are not endowed with enough precision, or the pointing model needs complete revision. Methinks the latter is the culprit, and is most easily fixed.

F.5 Second-order Extinction

From previous equations, you may notice that the V filter is the base from which the other Johnson-Cousins filters are stepped off. In standardized photometry, or all-sky photometry, color terms are used to correct for different spectral types. Each of the first-order coefficients may have an additional term added to enhance accuracy, with a color term included:

$$k'_V \Rightarrow k'_V + k''_V * (b - v)$$

$$k'_{BV} \Rightarrow k'_{BV} + k''_{BV} * (b - v)$$

$$k'_{UB} \Rightarrow k'_{UB} + k''_{UB} * (u - b)$$

$$k'_{VR} \Rightarrow k'_{VR} + k''_{VR} * (v - r)$$

$$k'_{RI} \Rightarrow k'_{RI} + k''_{RI} * (r - i)$$

Now, the corrected magnitudes read

$$v_0 = v - [k'_V + k''_V * (b - v)] * X_{ave}$$

$$(u - b)_0 = (u - b) * [1 - k''_{UB} * X_{ave}] - k'_{UB} * X_{ave}$$

$$(b - v)_0 = (b - v) * [1 - k''_{BV} * X_{ave}] - k'_{BV} * X_{ave}$$

$$(v - r)_0 = (v - r) * [1 - k''_{VR} * X_{ave}] - k'_{VR} * X_{ave}$$

$$(r - i)_0 = (r - i) * [1 - k''_{RI} * X_{ave}] - k'_{RI} * X_{ave}$$

Depending on the quality of the data, k''_V may be ignored.

As shown above, finding the first-order terms requires an appropriate range of airmasses for the best-fit line. Determining the second-order terms not only requires a wide range of airmasses, but wide range of colors. As the atmosphere-corrected colors should be constant with airmass, the difference of two of these colors (for any given airmass) should also be a constant. Take two standard stars 1 and 2, which are spaced close together. Then, for a set of observations in two filters (for standards in the same frame, just two images make up a set):

$$(b - v)_{0,1} - (b - v)_{0,2} = (b - v)_1 * [1 - k''_{BV} * X_{ave,1}] - k'_{BV} * X_{ave,1} - (b - v)_2 * [1 - k''_{BV} * X_{ave,2}] + k'_{BV} * X_{ave,2}$$

For adjacent observations of these stars (so that changes in airmass are negligible) $X_{ave,1} \cong X_{ave,2}$, and we use \bar{X}_{ave} as the average of these two values. (Note: standard stars within a single frame have the same airmass, and greatly simplify the process to find second-order coefficients. The only difficulty is finding same-frame standards with a wide difference in color.) We now rewrite this using Δ as the change the color from one star to another.

$$\Delta(b - v)_0 = \Delta(b - v) * [1 - k''_{BV} * \bar{X}_{ave}]$$

As $\Delta(b - v)_0$ is constant, the value for k''_{BV} is the best-fit slope of the graph of $\Delta(b - v)$ versus $\bar{X}_{ave} * \Delta(b - v)$, as shown below:

The other second-order coefficients (k''_{RI} , etc.) are generally ignored, but if one desires, the equations are in the same form as above.

If the colors are not dissimilar enough when solving for the second-order terms, the best-fit slope will not reflect the coefficient, but rather scatter in the instrumental magnitudes due to error and uncertainty. See the below graph, where the color difference between the two standards is too small to produce a consistent regression line. This graph, with the same range and domain as the one above, is attempting to find k''_{BV} .

First-order extinction terms (k'_V , k'_{VR} , etc.) change on a per night basis, and need to be determined every night when performing standardized photometry. However, second-order terms (like k''_{BV}) are both small and relatively stable. Determining these terms may only be necessary a couple of times a year.

F.5.1 Transformation Coefficients

Due to the differences of each imaging system (comprising the detector and filters), in order to compare values between systems, one's values must be transformed into a standard reference which was found using a specific detector and filter set. All other standard values are tethered to these primary standard stars, which have a usual error around 0.02 magnitudes. The accuracy of standardizing an object cannot improve much from that of the standard stars.

Common practice in determining standard star values is to use photomultiplier tubes instead of using CCDs. This removes error introduced by unflattened detectors and increases the uniformity of the measurements. Key in the field of standard stars is Arnold Landolt, whose work to calibrate stars using uniform detectors has allowed more accurate and stable standard values for standard stars. To be useful to most of both hemispheres, most standard fields crowd around the celestial equator. Often called Selected Areas, these fields are useful to the Utah area as they reach a minimum airmass of about 1.3.

Recapping the main equations of the last two sections,

$$\begin{aligned}
 v_0 &= v - k'_V \times X \\
 (b - v)_0 &= (b - v) \times [1 - k''_{BV} \times X_{ave}] - k'_{BV} \times X_{ave} \\
 (u - b)_0 &= (u - b) - k'_{UB} \times X_{ave} \\
 (v - r)_0 &= (v - r) - k'_{VR} \times X_{ave} \\
 (r - i)_0 &= (r - i) - k'_{RI} \times X_{ave}
 \end{aligned}$$

Lower-case letters indicate magnitudes from a specific (non-standard) system, while uppercase indicates a standard value. Transformation equations to the standard system are:

$$\begin{aligned}
 V &= v_0 + \varepsilon \times (B - V) + \zeta_V \\
 (B - V) &= \mu \times (b - v)_0 + \zeta_{BV} \\
 (U - B) &= \psi \times (u - b)_0 + \zeta_{UB} \\
 (V - R) &= \phi \times (v - r)_0 + \zeta_{VR} \\
 (R - I) &= \gamma \times (r - i)_0 + \zeta_{RI}
 \end{aligned}$$

Terms ε , μ , ψ , ϕ , and γ are transformation coefficients which relate the sensitivities of a system to the standard system, and are relatively constant in time. Usually, ε is very small, while all other transformation coefficients are quite close to one. The various ζ values are zero-point corrections and depend greatly on the conditions of the night- these need be determined each night all-sky photometry is attempted. In addition to changes from night to night, both the zero-point and transformation values alter when the sensitivity of the system changes (e.g.- from cleaned or realuminized mirrors).

To determine the transformation coefficients and zero-point corrections, stars with a range of colors and previously standardized magnitudes $UBVRI$ are observed. Their instrumental magnitudes are corrected for extinction by the method outlined above. Lastly, the following best-fit lines are needed to determine, for:

$$\varepsilon \text{ and } \zeta_V, \text{ the slope and intercept, respectively, of the graph of } (V - v_0) \text{ vs. } (B - V)$$

μ and ζ_{BV} , the slope and intercept of the graph of $(B - V) - (b - v)_0$ vs. $(B - V)$

ψ and ζ_{UB} , the slope and intercept of the graph of $(U - B) - (u - b)_0$ vs. $(U - B)$

ϕ and ζ_{VR} , the slope and intercept of the graph of $(V - R) - (v - r)_0$ vs. $(V - B)$

γ and ζ_{RI} , the slope and intercept of the graph of $(R - I) - (r - i)_0$ vs. $(R - I)$

As with before, if standard stars do not exhibit a wide enough range in colors, the regression lines will not yield tight correlations. Choosing standard stars properly will determine how easy it is to calculate standard magnitudes for unstandardized objects. Once good transformations are found from the standard stars, the equations are applied to the unstandardized object.

F.6 All-sky Photometry at ROVOR

Cameron Pace's thesis, submitted in 2010, dealt primarily with expanding and solidifying the number of stable reference stars nearby high-energy blazars (AGN whose jets are oriented towards the observer). Previous standardization of on-field standards were well reproduced by Cameron; still, discrepancies influenced him to revisit the data. In mid-2012, his follow-up work found a curious problem in conducting all-sky photometry at ROVOR. The secondary standards (also called in-field standards) exhibit significantly different extinction coefficients than do Landolt standard stars.

Taking blazar Markarian 501 as an example, on 10 August 2009, with 11 in-field standards, $k'_V = 0.239$ for the secondary standards while for the primary standards $k'_V = 0.318$ (which is, by itself, unusually high for the V filter). This disparity occurs over several photometric nights. The best explanation is as follows: ROVOR is based in a dusty valley floor, and near the site is a coal power plant which produces constant emissions. These heavy, low-lying aerosols (dust and coal) may be causing much greater extinction towards the horizon than usual. The Landolt fields are lower in the sky towards the south, while Mrk 501 travels close to the zenith. Air currents may also render

this non-uniform extinction a function of both the time of night and azimuth angle, but as of now no study of this has been completed.

Instead of using primary standards, then, one may choose to use in-field standards to find extinction and transformation coefficients and the zero-point values for an object for a night. This does require that one tracks each object (that will be standardized) through as broad a range in airmasses as possible; in-field standards may (as is the case of Mrk 501) have only small differences in color, making second-order and transformation coefficients more difficult to determine.

Appendix G

Using Programs in the Project

Several programs have been made to process and analyze astronomical data, and are capable of reading and writing in the .fits, .fit, .fts, or .FIT formats. These include MaxIm DL, Mira Pro, CCDSoft, and IRAF. CCDSoft and MaxIm DL are fundamentally data acquisition programs (used at ROVOR and WMO) and will not be explicitly discussed in this thesis. Instead, Mira Pro will be discussed as being prototypical of the other two programs. All three are GUI-based (graphical user interface), are compatible with Windows operating systems, and all three are proprietary software. IRAF is run on UNIX platforms and requires experience in command line operations¹, but is free (if you can get the system set up and all libraries in the correct locations).

G.1 Mira Pro

Like the other two above, Mira Pro 7UE is fully capable of handling large data sets. Images are easily viewed to see both the raw data and the effects of processing and combining frames.

¹PyRAF combines the Python language with IRAF scripts and capabilities. GUIs in PyRAF make IRAF tasks more manageable, but still requires IRAF installation. It is a promising direction for serious astronomers who are tired of CL grief.

This aids in discriminating between good and poor data, especially in preparing master calibration frames. As is typical in Windows programs, unsatisfactory changes are quickly undone (the visual aspect of Mira comes in handy here), reducing time otherwise spent recalibrating data. Mira (and MaxIm) are well maintained and have easy to use references. Benefits of having customers financially support upkeep and expansion.

For all the good, Mira and the others are not as robust in their calculations. Image processing primarily consists of arithmetic operations on whole pixels, and, with little exception, any of the mentioned programs are acceptable to process data. Photometry, however, requires circular apertures which measure fractions of pixels rather than the entire value. Mira has more limited accuracy in this than IRAF ($\frac{1}{10}^{th}$ of a pixel increments versus $\frac{1}{100}^{th}$) when approximating a pixel's contributions. As such, though the effect may be small, differences arise between the two programs. Photometry was done on 9 nights of ROVOR data of Markarian 501 standard stars in both Mira and IRAF to compare results. Prior research by Fiorucci & Tosti (1996), Villata et al. (1998), Pace(2010) and others have confirmed that these stars are stable. For each of these objects, the standard deviation of their differential photometry on these nights is plotted for each of the methods. First, photometry is done with Mira using an aperture equal to half of the average FWHM of the stars. Then, IRAF is used with the same apertures. Finally, a script in IRAF, NightPHOT4, is used to adjust the aperture and optimize the signal-to-noise of the measurement. Comparing Mira and IRAF with the same apertures, Mira values show almost twice as much scatter as those from IRAF. While Mira is convenient, its photometry is not as accurate. For low-signal sources especially, such as what we expect for the LLAGN targets, photometry need be done with IRAF.

G.2 IRAF

G.2.1 CL Commands and Reference

G.2.2 Flat-Field Surface Fitting

These sections document IRAF processes, and no attempt will be made to explain IRAF or command line programs explained here. For general help on these topics, find someone familiar with them and become friends. Most online sources (and official IRAF documentation) may sway a novice away from the basic astronomical packages.

In an attempt to make a better flat field, image surface-fitting was played with. Mira has limited surface-fitting capabilities, but as these are of low order², and do not as accurately reproduce the flat field pattern, nothing more of this will be discussed.

As discussed in appendix D.4, flat fields are composite quantum efficiency (QE) and illumination responses for the telescope-camera setup. A master flat contains both signatures, and at the beginning of this process only the illumination correction will be explored. First, make one or two copies of a good master flat (here called flatV.fits) as backups. The surface fit will average over small pixel-to-pixel variations, but manufacturing defects and cosmic-ray damage can severely alter the QE of affected pixels with respect to their neighbors. When a high-order polynomial tries to fit about these regions, it will try to fit the QE at the expense of the illumination frame. To smooth over these purely QE concerns, create a bad pixel mask (explained below) and smooth over the

²While 10th order cross-terms are as high as Mira will go in performing a surface fit, its process is designed to fit a background and ignore stars. This is a feature not easily available in IRAF's `imsurfit` process, and one collaborator gets good results using the background fit instead of twilight flats. As (twilight) flat fields are not more accurate than about 1%, the error from performing a low-order background fit is close to that of traditional flats. The only primary difference this author has seen is a poorer fit at the outer portions of a fitted frame. As long as photometry is kept to within the central 50-60% of a frame the edge fringing is negligible.

strongly deviating pixels.

First, take the master flat and boxcar smooth each pixel with respect to its eight immediate neighbors. The command is:

```
boxcar flatV.fits flatVboxcar.fits 3 3
```

Where "flatVboxcar.fits" is the output image the program creates and "3 3" are the pixel dimensions, x and y, of the averaging region, centered on each pixel. The purpose of this step is to slightly blur bad regions into neighboring pixels. Then, to isolate these regions, make a difference frame:

```
imarith flatV.fits / flatVboxcar.fits flatVdiff.fits
```

Any distinct features in the output "flatVdiff.fits" file correspond to real QE hot/dead spots on the detector (or possible stars that made it through the processing describe in appendix D.5.3. In either case, these will be removed and will not be a concern for the illumination frame). The quickest method to determine the pixel (physical) positions of these points is to load the image into a ds9 window and make a "~.reg" file by centering created regions on the distinct features individually. Save the "~.reg" file in "xy" format and with physical coordinates. Then, open this file to retrieve the coordinates. Each line will contain an x and a y pixel coordinate, in that order. Copy and paste these into a Excel-like worksheet. Make a four-column list by first subtracting 3 from the x coordinates, then adding 3, and then subtracting 3 from the y coordinates, and then adding 3. The pattern will look something like:

```
257 263 592 598
```

This defines a 3×3 box about the affected pixel(s) which will have its values replaced with an interpolation from the neighboring pixel values. Dump this four-column list into a text file and call it badpix.txt (or anything else, just so long as "bpm" are not in the name. The fixpix program looks by default for a file of this name). To smooth over these pixels, take an original master flat and perform the following:

```
fixpix flatV.fits badpix.txt
```

Rename "flatV.fits" something else, like "flatVgoodpix.fits" to avoid confusion later on.

```
boxcar flatVgoodpix.fits flatVboxgoodpix.fits 3 3
```

Now the serious QE effects should be removed. The command to fit the illumination response is:

```
imsurfit flatVboxgoodpix.fits flatVsurf.63.fits 63 63 function=chebyshev
```

The output file "flatVsurf.80.fits" has 80 within it to keep track of the order of the polynomial fit. Function can also be "legendre", and both appear to work as effectively. "80 80" indicates the x-order and y-order, respectively, for the polynomial terms. Which order to go to is a tricky question, and will depend on the filter. At low orders (<30), the fit to the vignetting/filter transmission signature will not be accurately represented. At high orders (>80), the fitting begins to include dust diffraction patterns and 'rings' at sharp edges, like the dust patterns and the frame edges. Using only the V filter, and the above process on the SITE detector, it was determined that the differences between frames reached a minimum at a 63rd order polynomial (indicating where 'ringing' began to be more detrimental than the increased illumination accuracy).

A QE frame is now easily produced by dividing an untouched master flat frame by the best-order polynomial fit:

```
imarith flatV.fits / flatVsurf.63.fits flatVQE.fits
```

In addition to random scatter in pixel-to-pixel values, which is expected due to manufacturing inequalities, damaged regions will now very quickly appear. Note that "flatVQE.fits" is normalized, and only represents QE with respect to the average of the frame. True QE is lower- around the 70-85% range for a standard CCD in the V filter. Over the course of 3 months, the illumination frame changed pattern slightly, indicating either a difference in twilight gradient or thin, non-isotropic clouds. QE did not change in this range; in fact, only the dust pattern changed over the course of two years, and these change the detector response for affected pixels by about $\frac{1}{2}\%$ at most. This adds support to the conclusion in appendix D.4.3 that flat fields are stable over long periods of

time. Further studies could investigate the different QE of different wavelength ranges, and could find best-order polynomial fits for the old ROVOR Ap47 CCD used in 2009.

G.3 PSF-matching Image Subtraction Programs

HOTPANTS stands for High Order Transformation of PSF ANd Template Subtraction. While it is not alone in performing image convolution and subtraction, HOTPANTS is the most user-friendly program available. Other programs to look into are ppSub (), and the Automated Image Subtraction Algorithm, or ISIS (Alard & Lupton 1998; Alard 2000). ISIS is used for finding supernovae, as in the Supernova Legacy Survey (Astier et al. 2006) and follow-up work on microlensing surveys like OGLE (Alard 2000). ppSub is used by the Pan-STARRS all-sky survey for has a variety of options which tailor the image subtraction to the input data. The program runs by command line (CL), and it is easiest to run on a Linux-based system. Currently, this program is run by the group out of Cygwin terminal. CL programs require that the files to be used are in the working directory, or that the path name is defined.

G.3.1 Method to Determine Errors in PSF-matching Image Subtraction

This section follows the Boizelle-Holden-Young derivation for including flat-field errors.

Once a S/N is obtained for both input and template images by the method of section 9.3.1, compute the on-frame error by

$$\sigma(mag) = 2.5 \times \log_{10}\left(1 + \frac{1}{S/N}\right)$$

Naively, for input and template magnitude errors σ_{input} and $\sigma_{template}$, the image subtraction $I(x, y) - T_C(x, y)$ would produce propagated error $\sigma_{conv,sub}^2 = \sigma_{input}^2 + \sigma_{template}^2$. Each frame, however, is complicated by the addition of flat fielding uncertainties. Locally, flat fielding errors are small- 0.2 to 0.3 percent due to shot noise. Only over hundreds of pixels does the illumination response

change by several percent. Transform uncertainty $\sigma(mag)$ to the uncertainty $\sigma(counts)$ measured in ADU:

$$\begin{aligned}\sigma_{initial}(mag) &= 2.5 \times \log_{10}\left(1 + \frac{N}{S}\right) = -2.5 \times \log_{10}\left(\frac{sum_1}{sum_2}\right) \\ &= -2.5 \times \log_{10}\left(1 + \frac{\sigma_{initial}(counts)}{sum_{obj}}\right)\end{aligned}$$

Where F is flux.

$$\begin{aligned}\left(\frac{\sigma_{final}(counts)}{sum_{obj}}\right)^2 &= \left(\frac{\sigma_{initial}(counts)}{sum_{obj}}\right)^2 + \left(\frac{\sigma_{flat}(counts)}{C_{flat}}\right)^2 \\ \sigma_{flat}(counts) &= \sqrt{C_{flat}} / \sqrt{N_{frames}} = \sqrt{\frac{C_{flat}}{N_{frames}}} \\ \frac{\sigma_{initial}(counts)}{sum_{obj}} &= 1 - 10^{\left(\frac{-\sigma_{initial}(mag)}{2.5}\right)} \\ \frac{\sigma_{final}(counts)}{sum_{obj}} &= \sqrt{\left(1 - 10^{\left(\frac{-\sigma_{initial}(mag)}{2.5}\right)}\right)^2 + (C_{flat} \times N_{frames})^{-1}} \\ 1 - 10^{\left(\frac{-\sigma_{initial}(mag)}{2.5}\right)} &= 1 - 10^{\frac{-2.5 \times \log_{10}\left(1 + \frac{N}{S}\right)}{2.5}} = \frac{N}{S} \text{ or } \left(\frac{S}{N}\right)^{-1} \\ \sigma_{final}(counts) &= sum_{obj} \times \sqrt{\left(\frac{S}{N}\right)^{-2} + (C_{flat} \times N_{frames})^{-1}} \\ \sigma_{final}(mag) &= 2.5 \times \log_{10}\left[1 + \sqrt{\left(\frac{S}{N}\right)^{-2} + (C_{flat} \times N_{frames})^{-1}}\right]\end{aligned}$$

The final error calculation (in counts) gives the error of the PSF-matching and subtraction routine.

$$\sigma_{conv,sub}(counts) = \sqrt{(\sigma_{input,final}(counts))^2 + (\sigma_{template,final}(counts))^2}$$

G.3.2 Using HOTPANTS

When using HOTPANTS (version 5.1.10), the CL code to type is:

```
hotpants [options]
```

The [options] above includes the required input, template, and output file names, and the optional tweaks to customize the program to a specific object and telescope. Below is a more indepth explanation of how the ISIS and HOTPANTS code works, taken from Alard (2000) and Becker, and also the required and optional input parameters. To perform a simple image subtraction, the code one might use is:

```
hotpants -inim mk501.1.fits -tmplim mk501.template.fits -outim mk501.diff.fits -nsx 7 -nsy 7 -tu
110000 -iu 55000
```

One must be especially careful to make the template as high-quality as possible. This will affect the image subtraction more than any other step.

Options `-tu` and `-iu` allow PSFs to be used in the kernel fitting for objects with counts other than the default 25000. As most detectors do not saturate or become non-linear near this range, to expand the number of objects measured and used, use a value just under where the detector becomes unreliable (for combined images, adjust this value to the sum of the useable dynamic range of the frames). The parameters `-nsx` and `-nsy` define the number of "stamps" to make for the x and y directions. An image is divided up into a number of regions which contain stamps (subregions) for which the kernel solution is performed independently. The goal of the program is to seek a convolution kernel K which minimizes the equation $\sum([T \otimes K](x,y) - I(x,y))^2$ for $I(x,y)$ the input image, $T(x,y)$ the template image, and K the kernel transformation. The kernel is composed of basis functions such that

$$K(u, v) = \sum_{i,j,k} A_{i,j,k} K_{i,j,k}(u, v), \text{ with}$$

$$K_{i,j,k}(u, v) = u^i v^j \exp\left(-\frac{(u^2+v^2)}{2\sigma_k}\right)$$

While the powers of dummy variables u and v above $i, j = 0$ do not produce uniform gaussians, the sum of enough of these functions should make an appropriate kernel function.

Unless the user changes the parameters, the defaults are $k = 3$, with:

$$\sigma_1 = 0.7 \text{ pixels}; i + j \leq 6$$

$$\sigma_2 = 1.5 \text{ pixels}; i + j \leq 4$$

$$\sigma_3 = 3.0 \text{ pixels}; i + j \leq 2$$

The first, σ_1 , is allowed to spatially vary, but contains a narrow gaussian, while the last, σ_3 , is confined spatially but has a broad gaussian component. The point-spread function (PSF) of an image is shown in section 9.

Finally, the program returns an image equal to $I(x,y) - [T \otimes K](x,y)$. Note that the image of higher quality is blurred to match the seeing of the input image. This is essential for the quality of the output image. When choosing the template image, it is best to combine high-quality images. High-quality demands that the PSFs are very tight, with the FWHM the smallest of the set to be difference. Also, do not use any calibrated images which have gradients across the frame, either from poor processing or moon contamination (Pancoast 2012). If need be, a low-order background removal may be used, but how to remove the gradient depends on the source: if the flat fielding is wrong, the effect is multiplicative; if from a non-uniform light source, such as the moon, the effect is additive. As these effects may be compounded, it is best to avoid removing a gradient from potential template images and choose gradient-free frames instead³.

The required and additional options for the command line hotpants code are given below for convenience. Unless the data is better served with different parameters, the example code above is generally sufficient. If not specifically altered, each option defaults to the values in parentheses.

Required options:

[-inim fitsfile] : comparison image to be differenced

[-tmplim fitsfile]: template image

[-outim fitsfile] : output difference image

Additional options:

[-tu tuthresh] : upper valid data count, template (25000)

[-tuk tuthresh] : upper valid data count for kernel, template (tuthresh)

[-tl tlthresh] : lower valid data count, template (0)

[-tg tgain] : gain in template (1)

[-tr trdnoise] : e- readnoise in template (0)

[-tp tpedestal] : ADU pedestal in template (0)

[-tni fitsfile] : input template noise array (undef)

[-tmi fitsfile] : input template mask image (undef)

[-iu iuthresh] : upper valid data count, image (25000)

³Note that the input images can have some gradient issues and still produce proper data.

- [-iuk iuthresh] : upper valid data count for kernel, image (iuthresh)
- [-il iltthresh] : lower valid data count, image (0)
- [-ig igain] : gain in image (1)
- [-ir irdnoise] : e- readnoise in image (0)
- [-ip ipedestal] : ADU pedestal in image (0)
- [-ini fitsfile] : input image noise array (undef)
- [-imi fitsfile] : input image mask image (undef)
- [-ki fitsfile] : use kernel table in image header (undef)
- [-r rkernel] : convolution kernel half width (10)
- [-kcs step] : size of step for spatial convolution ($2*rkernel + 1$)
- [-ft fitthresh] : RMS threshold for good centroid in kernel fit (20.0)
- [-sft scale] : scale fitthresh by this fraction if... (0.5)
- [-nft fraction] : this fraction of stamps are not filled (0.1)
- [-mins spread] : Fraction of kernel half width to spread input mask (1.0)
- [-mous spread] : Ditto output mask, negative = no diffim masking (1.0)
- [-omi fitsfile] : Output bad pixel mask (undef)
- [-gd xmin xmax ymin ymax] : only use subsection of full image (full image)
- [-nrx xregion] : number of image regions in x dimension (1)
- [-nry yregion] : number of image regions in y dimension (1)
- OR –
- [-rf regionfile] : ascii file with image regions 'xmin:xmax,ymin:ymax'
- OR –
- [-rkw keyword num]: header 'keyword[0->(num-1)]' indicates valid regions
- [-nsx xstamp] : number of each region's stamps in x dimension (10)
- [-nsy ystamp] : number of each region's stamps in y dimension (10)
- OR –
- [-ssf stampfile] : ascii file indicating substamp centers 'x y'
- OR –
- [-cmp cmpfile] : .cmp file indicating substamp centers 'x y'
- [-afssc find] : autofind stamp centers so #=-nss when -ssf,-cmp (1)

[-nss substamps] : number of centroids to use for each stamp (3)
[-rss radius] : half width substamp to extract around each centroid (15)
[-savexy file] : save positions of stamps for convolution kernel (undef)
[-c toconvolve] : force convolution on (t)emplate or (i)mage (undef)
[-n normalize] : normalize to (t)emplate, (i)mage, or (u)nconvolved (t)
[-fom figmerit] : (v)ariance, (s)igma or (h)istogram convolution merit (v)
[-sconv] : all regions convolved in same direction (0)
[-ko kernelorder] : spatial order of kernel variation within region (2)
[-bgo bgorder] : spatial order of background variation within region (1)
[-ssig statsig] : threshold for sigma clipping statistics (3.0)
[-ks badkernel sig] : high sigma rejection for bad stamps in kernel fit (2.0)
[-kfm kerfracmask] : fraction of abs(kernel) sum for ok pixel (0.990)
[-okn] : rescale noise for 'ok' pixels (0)
[-fi fill] : value for invalid (bad) pixels (1.0e-30)
[-fin fill] : noise image only fillvalue (0.0e+00)
[-convvar] : convolve variance not noise (0)
[-oni fitsfile] : output noise image (undef)
[-ond fitsfile] : output noise scaled difference image (undef)
[-nim] : add noise image as layer to sub image (0)
[-ndm] : add noise-scaled sub image as layer to sub image (0)
[-oci fitsfile] : output convolved image (undef)
[-cim] : add convolved image as layer to sub image (0)
[-allm] : output all possible image layers
[-nc] : do not clobber output image (0)
[-hki] : print extensive kernel info to output image header (0)
[-oki fitsfile] : new fitsfile with kernel info (under)
[-sht] : output images 16 bitpix int, vs -32 bitpix float (0)
[-obs bscale] : if -sht, output image BSCALE, overrides -inim (1.0)
[-obz bzero] : if -sht, output image BZERO, overrides -inim (0.0)
[-nsht] : output noise image 16 bitpix int, vs -32 bitpix float (0)

[-nbs bscale] : noise image only BSCALE, overrides -obs (1.0)
 [-nbz bzero] : noise image only BZERO, overrides -obz (0.0)
 [-ng ngauss degree0 sigma0 .. degreeN sigmaN]
 : ngauss = number of gaussians which compose kernel (3)
 : degree = degree of polynomial associated with gaussian # (6 4 2)
 : sigma = width of gaussian # (0.70 1.50 3.00) : N = 0 .. ngauss - 1
 : (3 6 0.70 4 1.50 2 3.00)
 [-pca nk k0.fits ... n(k-1).fits]
 : nk = number of input basis functions
 : k?.fits = name of fitsfile holding basis function
 : Since this uses input basis functions, it will fix
 : hwKernel
 [-v] verbosity : level of verbosity, 0-2 (1)

NOTE: Fits header params will be added to the difference image

COMMAND (what was called on the command line)
 NREGION (number of regions in image)
 PHOTNORM (to which system the difference image is normalized)
 TARGET (image which was differenced)
 TEMPLATE (template for the difference imaging)
 DIFFIM (output difference image)
 MASKVAL (value for masked pixels)
 REGION?? (IRAF-format limits for each region in the image)
 CONVOL?? (which image was convolved for each region)
 KSUM?? (sum of the convolution kernel for each region)

G.3.3 Using Cross-Convolution Image Subtraction

Yuan & Akerlof (2008) have developed a symmetric psf-matching convolution program (called cross-convolution) based on the ISIS code by Alard (2000). Using the same nomenclature as above, with $T(x,y)$ the template and $I(x,y)$ the input image, now there are two respective kernel

transformations $K_T(u, v)$ and $K_I(u, v)$. Cross-convolution does not assume that the template is of good quality. Instead, a symmetric approach is taken: $T(x, y) \otimes K_T(u, v) \approx I(x, y) \otimes K_I(x, y)$. To solve, as for the ISIS method, a nonnegative function is constructed and then minimized. The natural choice is $D(x, y) = \{T(x, y) \otimes K_T(u, v)\} - \{I(x, y) \otimes K_I(x, y)\}$ with the minimization being of $\sum D(x, y)^2$. Minimizing the spatial sum of $D(x, y)^2$, as it is above, the kernels would seek to almost infinitely blur the template and input images. This would match the images well, but make extracting data impossible. The end minimization equation has some additional components which push the convolutions towards less-broad solutions.

A brief photometry comparison showed that using an illumination frame and a QE frame instead of the inclusive flat frame when processing produced almost identical results. If there is much difference, the illumination/QE reductions produce slightly more scatter in standard stars than do traditional twilight flats.

Bibliography

Alard, C. 2000, A&AS, 144, 363

Alard, C., & Lupton, R. H. 1998, AJ, 503, 325

Astier, P., Guy, J., Regnault, N. et al. 2006, A&A, 471, 31

Baldwin, J. A., Phillips, M. M., & Terlevich, R. 1981, PASP, 93, 5

Barth, A. 2012, personal communication

Barth, A. J., Nguyen, M. L., Malkan, M. A., et al. 2011a, ApJ, 732, 121

Barth, A. J., Pancoast, A., Thorman, S. J., et al. 2011b, ApJL, 734, L4

Bentz, M., Denney, K. D., Cackett, E. M., et al. 2006, ApJ, 651, 775

Blandford, R. D. & Rees, M. J. 1978, PhyS, 17, 265

Burbidge, G. & Burbidge, M. 1967, Quasi-Stellar Objects (San Francisco, CA: W. H. Freeman and Company), 10

Carroll, B. W., & Ostlie, D. A. 1996, An Introduction to Modern Astrophysics (Reading, Massachusetts: Addison-Wesley Publishing Company, Inc.)

Chromey, F. R., & Hasselbacher, D. A. 1996, PASP, 108, 944

- Constantin, A. 2011, *Supermassive Black Hole Accretion Across the Age of the Universe* (Lexington, Kentucky: Lambert Academic Press)
- Crawford, D. L. 1987, in *New Generation Small Telescopes*, ed. D. S. Hayes et al. (Mesa, AZ: Fairborn Press), 345
- Denney, K. D., Peterson, B. M., Pogge, R. W., et al. 2010, *ApJ*, 721, 715
- Fiorucci, M. & Tosti, G. 1996, *A&AS*, 116, 403
- Gebhardt, K., Bender, R., Bower, G., et al. 2000, *ApJ*, 539, L13
- Goulding, A. D., Alexander, D. M., Lehmer, B. D., & Mullaney, J. R. 2010, *MNRAS*, 406, 597
- Grier, C. J., Peterson, B. M., Pogge, R. W., et al. 2012, *ApJL*, 744, L4
- Grier, C. J., Peterson, B. M., Pogge, R. W., et al. 2012, arXiv:1206.6523
- Henden, A. A., & Kaitchuck, R. H. 1982, *Astronomical Photometry*. (New York, New York: Van Nostrand Reinhold Company)
- Hewitt, A. & Burbidge, G. 1993, *ApJS*, 87, 451
- Ho, L. C., Filippenko, A. V., & Sargent, W. L. W. 1994, *IAUS*, 159, 275
- Ho, L. C., Filippenko, A. V., & Sargent, W. L. W. 1995, *ApJS*, 98, 477
- Ho, L. C., Filippenko, A. V., & Sargent, W. L. W. 1997, *ApJS*, 112, 337
- Ho, L. C. 2008, *ARA&A*, 46, 475-540
- Ho, L. C., Darling, J., & Greene, J. E. 2008, *ApJS*, 177, 103
- Howell, S. B. 2006, *Handbook of CCD Astronomy* (2nd ed; Cambridge, England: Cambridge University Press), 13

- Ivezic, Z., et al. 2011, The LSST System Science Requirements Document, v5.2.3, The LSST Science Council, <http://www.lsst.org/files/docs/SRD.pdf>.
- Janesick, J., & Elliott, T. 1992, in ASP Conf. Ser. 23, *Astronomical CCD Observing and Reduction Techniques*, ed. S. B. Howell (San Francisco, CA: ASP), 1
- Janesick, J. R. 2001, *Scientific Charge-Coupled Devices* (Bellingham, Washington; SPIE Press)
- Joner, M. D., 2011, in personal communication
- Joyce, R. R. 1992, in ASP Conf. Ser. 23, *Astronomical CCD Observing and Reduction Techniques*, ed. S. B. Howell (San Francisco, CA: ASP), 258
- Kochanski, G. P., Tyson, J. A., & Fischer, P. 1996, *AJ*, 111, 1444
- Krisciunas, K., & Schaefer, B. E. 1991, *PASP*, 103, 1033
- Krolik, J. H. 1999, *Active Galactic Nuclei: From the Central Black Hole to the Galactic Environment*. (Princeton, New Jersey: Princeton University Press)
- Landolt, A. U. 2009, *AJ*, 137, 4186
- Levinson, A., Laor, A., & Vermeulen, R. C. 1995, 448, 589
- LSST Science Collaborations. 2009, *LSST Science Book*, version 2.0, <http://www.lsst.org>
- Mackay, C. D. 1986, *ARA&A*, 24, 255-83
- Massey, P., & Jacoby, G. H. 1992, in ASP Conf. Ser. 23, *Astronomical CCD Observing and Reduction Techniques*, ed. S. B. Howell (San Francisco, CA: ASP), 240
- Markarian, B. E., Lipovetsky, V. A., & Stepanian, D. A. 1981, *Astrophysics*, 16, 4
- Matthews, T. A., & Sandage, A. R. 1963, *ApJ*, 138, 30

McConnell, N. J., Ma, C., Gebhardt, K., et al. 2011, *Nature*, 480, 215

McNeil, S. 2004, dissertation

Moody, J. W., Boizelle, B., Bates, K., Little, B., McCombs, T., Nelson, J., Pace, C., Pearson III, R. L., Harrison, J., Brown, P. J., & Barnes, J. 2012, *PASP*, submitted

Mushotzky, R. F., Done, C., & Pounds, K. A. 1993, *ARA&A*, 31, 717-61

Nesvadba, N. P. H., Lehnert, M. D., Breuck, C. De, Gilbert, A. M., & van Breugel, W. 2008, *A&A*, 491, 407

Neugebauer, G., Soifer, B. T., Matthews, K., & Elias, J. H. 1989, *AJ*, 97, 4

Osterbrock, D. E. 1978, *PNAS*, 75 no 2, 540

Osterbrock, D. E. 1981, *ApJ*, 249, 462

Pace, C. 2010, thesis

Pancoast, A. 2012, personal communication

Pearson, R., Pace, C., & Moody, J. W. 2010, *BAAS*, 42, 378

Peterson, B. M. 1997, *An Introduction to Active Galactic Nuclei*. (Cambridge, England: Cambridge University Press)

Petrosian, A., McLean, B., Allen, R. J., & MacKenty, J. W. 2007, *ApJS*, 170, 33-70

Rees, M. J. 1977, *QJRAS*, 18, 429

Seth, A. C., Cappellari, M., Neumayer, N., et al. 2010, *ApJ*, 714, 713

Smith, H. J., & Hoffleit, D. 1963, *Nat*, 198, 650

- Smith, R. C. 1995, *Observational Astrophysics* (Cambridge, England: Cambridge University Press)
- Stubbs, C. W., & Tonry, J. L. 2006, *AJ*, 646, 1436
- Taylor, J. R. 1997, *An Introduction to Error Analysis* (2nd ed; Sausalito, CA: University Science Books)
- Tomaney, A. B. & Crotts, A. P. S. 1996, *AJ*, 112, 2872
- Ulrich, M., Maraschi, L., & Urry, C. M. 1997, *ARA&A*, 35:445-502
- de Vaucouleurs, G. 1977, *ApJS*, 33, 211
- Vaz, A. L. 2011, *Building a Better Flat-Field: An Instrumental Calibration Projector for the Large Synoptic Survey Telescope*, Undergraduate Thesis, MIT
- Villata, M., Raiteri, C. M., Lanteri, L., Sorbrito, G., & Cavallone, M. 1998, *A&AS*, 130, 305
- Veilleux, S., & Osterbrock, D. E. 1987, *ApJS*, 63, 295
- Walker, G. W. 1987, *Astronomical Observations: an Optical Perspective* (Cambridge, England: Cambridge University Press)
- Wang, J., Risaliti, G., Fabbiano, G., et al. 2010, *ApJ*, 714, 1497
- Woo, J., Treu, T., Barth, A. J., et al. 2010, *ApJ*, 716, 269
- Yuan, F., & Akerlof, C. W. 2008, *AJ*, 677, 808