Principal Component Analysis of the X-ray Spectra of Blazars

by

Dennis Gallant

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

BACHELOR OF SCIENCE

in

Honours Astrophysics

(Department of Astronomy and Physics, Dr. Luigi Gallo supervising faculty)

.....

.....

.....

.....

.....

SAINT MARY'S UNIVERSITY

April 20, 2018

Abstract

Principal Component Analysis of the X-ray Spectra of Blazars

by *Dennis Gallant*

submitted on April 20, 2018:

Principal Component Analysis (PCA) is applied to a variety of blazars to examine X-ray spectral variability. Data from nine different objects are analysed in two ways: long-term, which examines variability trends across years or decades, and short-term, which looks at variability within a single observation. The results are then compared to simulated spectra in order to identify the physical components that they correspond to. It is found that long-term variability for all objects is dominated by changes in a single power law component. The primary component is responsible for more than 84 per cent of the variability in every object, while the second component is responsible for at least 3 per cent. Short-term variability is less clear-cut, with no obvious physical analogue for some of the PCA results. We discuss the simulation process, and specifically remark on the consequences of the breakdown of the linearity assumption of PCA and how it manifests in the real data.

Contents

List of Figures

List of Tables

Chapter 1

INTRODUCTION

1.1 Active Galactic Nuclei

Most galaxies harbour a supermassive black hole at their centres. These objects can be millions or even billions of times as massive as our Sun. In most galaxies, such as our own, the black hole is dormant and not currently accreting material. However, in what are known as active galaxies, the central black hole is actively accreting nearby gas and dust. This accretion process releases a large amount of energy, as the friction in the accretion disk heats the material. This powerful emission can cause the central black hole to outshine the rest of the galaxy. The compact, bright, central region of these galaxies is known as an active galactic nucleus, or AGN.

Observational data would suggest that there are many different types of AGN, each with different spectral shapes and characteristics. However, these different objects can be unified under a single model, which describes them as a single type of object viewed from different perspectives (Antonucci, 1993). The inclination of the AGN relative to the observer gives rise to the observed characteristics.

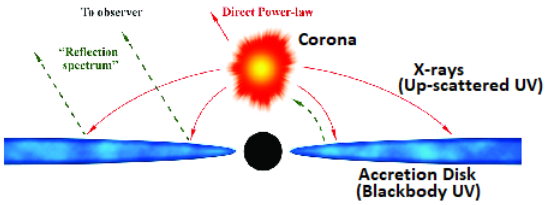

Figure 1.1: Illustration of the central engine of an AGN. The AGN is powered by accretion onto a central black hole region, which produces blackbody radiation in the UV. That radiation is up-scattered by electrons in a nearby corona of hot material. Some of the resultant X-rays reach the observer directly and appear as a power law shape, while others return to illuminate the accretion disk, producing a reflection spectrum with features such as emission lines in the X-ray band. Credit: http://inspirehep.net/record/1296802/plots

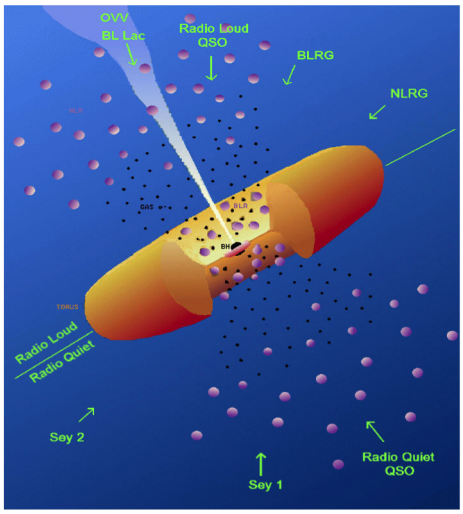

Figure 1.2: Illustration of the unified model of AGN. The classification of the object depends on the direction to the observer and on the presence of a jet. The two main categories are radio-loud and radio-quiet. Further categorizations include QSOs (quasi-stellar objects), BLRGs (broad-line radio galaxies) and NLRGs (narrow-line radio galaxies) as well as their radio-quiet counterparts Seyfert Is and Seyfert IIs. Finally, there are blazars (labelled BL Lac in the diagram after the prototype object). Note that the AGN is still embedded within a host galaxy, and the entire region within the torus is tiny in comparison to the rest of the galaxy. The narrow-line region exists on a galactic scale. Credit: Urry and Padovani (1995)

1.2 The Unified Model of AGN

All AGN have a supermassive black hole at their core. This black hole is surrounded by an accretion disk, which continuously funnels new material towards the black hole. The friction from the fast-moving disk material causes the disk to heat up and emit blackbody radiation in the UV region. This compact, inner region is known as the central engine, and it is the source of the AGN's formidable luminosity. Near the black hole is the corona, a region of hot, ionized material. The full geometry and extent of the corona is still unknown, but its presence can be deduced from the X-ray emission found in AGN spectra. UV photons from the disk reach the corona and are then up-scattered (through inverse Compton scattering, see Section 1.4) into X-rays by free electrons in the corona. Figure 1.1 illustrates the central engine of a typical AGN. Some of the up-scattered photons travel straight to the observer, giving rise to a power law spectrum (see Section 1.4). Others return to illuminate the disk, producing a reflection spectrum that contains features such as emission lines. These features provide information about the chemical make-up and dynamics of the disk.

Further out from the accretion disk, the AGN is ringed by a large dusty torus. This torus is colder than the inner regions, and emits primarily in the infrared. Because the torus is opaque to most light, it can obscure the view of the inner regions if viewed from the right angle.

The region between the central engine and the torus is known as the broad-line

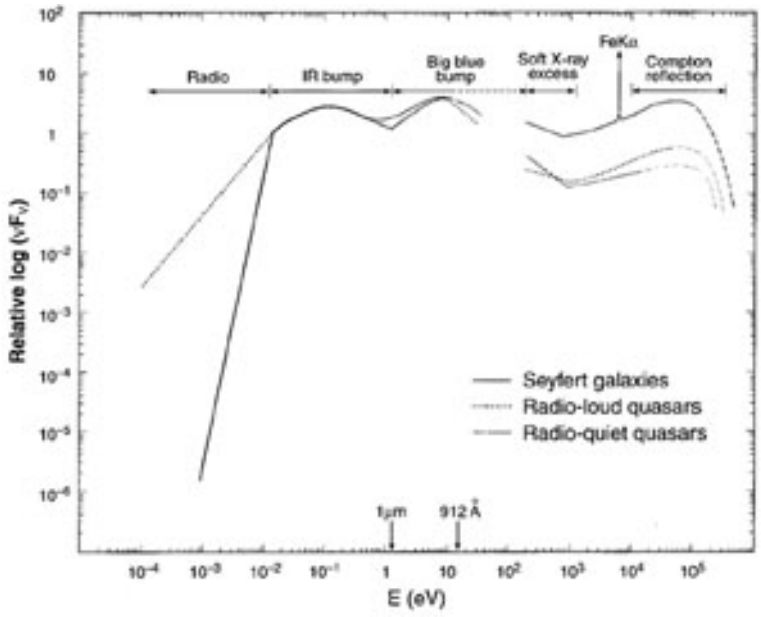

Figure 1.3: Diagram showing the broadband spectrum of a typical AGN. These objects typically show significant radiation at all wavelengths, each from a different mechanism. Radio waves are mainly generated by synchrotron radiation, the infrared light is radiated by the cool torus, the UV is thermal blackbody emission, and the X-rays are up-scattered UV photons. Credit: Koratkar and Blaes, 1999

region. This area consists of gas clouds that are heated by the central engine and produce an emission spectrum. It is called the broad-line region because these clouds are moving with high velocities (often 5000kms^{-1} or more) and therefore the emission lines are heavily broadened by the Doppler effect.

Beyond the torus is the narrow-line region, where the gas clouds have much lower orbital velocities and therefore display a narrow emission spectrum. Since these clouds are beyond the torus, they are visible from any angle. The narrow-line region has no sharp boundary, and extends outward far into the host galaxy in a broad cone-like shape. The obscuring effect of the torus divides AGN into two main categories: Type I, or broad-line, where the view of the central engine is unobstructed, and type II, or narrow-line, where the torus blocks the observer's view into the Doppler-broadened inner region. Figure 1.2 shows the effects of inclination on observations of AGN.

Figure 1.3 shows the spectrum of a typical AGN at all wavelengths. It includes emission across the entire electromagnetic spectrum, a common feature in AGN. The radio emission comes mostly from the synchrotron process, and is much stronger in radio-loud AGN (Section 1.3). The infrared bump is light from the central engine being re-emitted by the cool torus. The spectrum peaks in the UV via blackbody radiation from the disk, but there is still a strong X-ray component comprised of up-scattered UV photons emitted by the corona. Some of these photons return to the disk and are reflected.

One of the more interesting properties of AGN is that they are highly variable objects. Due to the intense conditions within the central engine, the luminosity of an AGN can vary by orders of magnitude over the span of hours. This variability occurs in all wavebands, and effects not just the strength (or normalization) of the spectrum, but also its shape. Studying this variability allows scientists to indirectly learn about the processes that drive an AGN, as well as determine characteristic timescales and sizes of these processes and the spatial regions they take place in.

1.3 Radio-Loud AGN

Not all differences between AGN can be described as the effects of viewing angle. Some objects possess what is known as a jet, a large column of ionized material being launched out of the central region perpendicular to the plane of the disk. This material is highly relativistic, and emits a large amount of light from radio to gamma rays through synchrotron radiation and the synchrotron self-Compton effect (see Section 1.4). The mechanism responsible for these jets is still not fully understood. Objects that possess a jet are known as radio-loud AGN (due to the strong radio emission from the jet), and those that do not are called radio-quiet. While the difference between narrow-line and broad-line AGN is purely due to orientation, the difference between radio-loud and radio-quiet objects is a physical one. While some objects may have a jet that is blocked from our view, most objects possess no jet.

The presence of a jet introduces another effect of inclination. Due to the relativistic speed of the material in the jet, its emission is highly beamed in the jet's direction (see Section 1.4). As the angle between the observer and the jet decreases, the jet becomes more and more prominent. If the jet is pointed directly towards the observer, the object can appear orders of magnitude brighter than it otherwise would. When viewed directly like this, the object is known as a blazar. A blazar's spectrum is dominated by the jet, and is simpler than that of most AGN. As shown in Figure 1.4, the spectra of blazars show a broad doubly-peaked shape that range from radio waves to the X-ray or gamma regime. This shape is attributable to synchrotron emission and the synchrotron self-Compton process (Section 1.4). Features of ordinary AGN such as the infrared bump or prominent emission lines are no longer visible. Blazars are extremely variable at all wavelengths, even more so than other AGN. This work aims to study that variability across a variety of objects and timescales. Due to their extreme luminosity, relative spectral simplicity, and rapid variability, blazars are the perfect candidate for this work.

1.4 X-ray emission processes and X-ray spectra

While AGN are known for their broadband spectra that span all across the electromagnetic spectrum, this work focuses solely on X-rays. AGN are often studied in the X-ray regime because X-rays can easily penetrate the interstellar medium that would otherwise obscure our view (as it does for UV , for example). It takes extraordinary high-energy processes to produce X-ray photons. Even AGN accretion disks

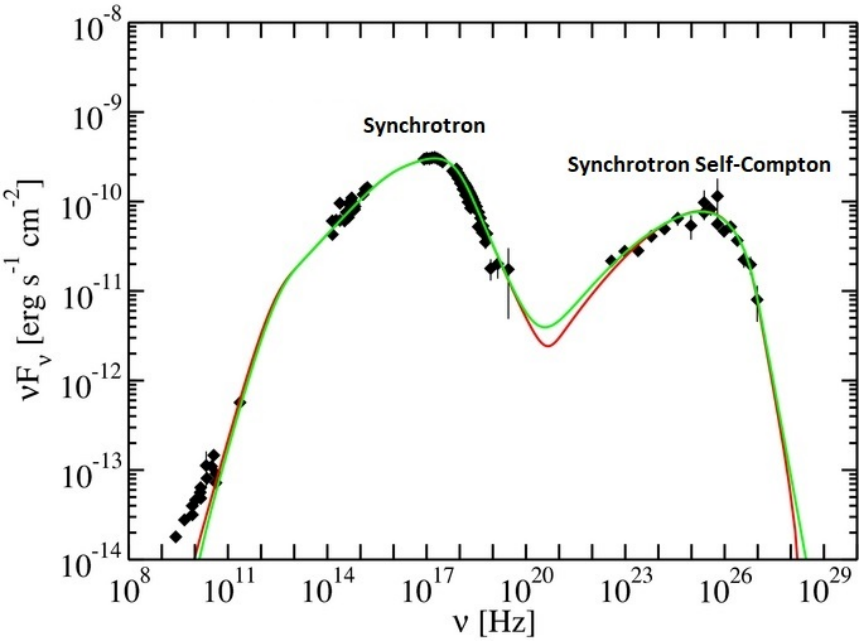

Figure 1.4: Broadband spectral energy distribution of the blazar Mrk 421, showing data from several different observatories at various wavelengths. A couple of double-peaked synchrotron + synchrotron self-Compton (SSC) models are fit to the data, showing that the spectra of blazars is dominated by the effect of the jet. The region corresponding to each emission mechanism is labelled. Credit: Abdo et al (2011).

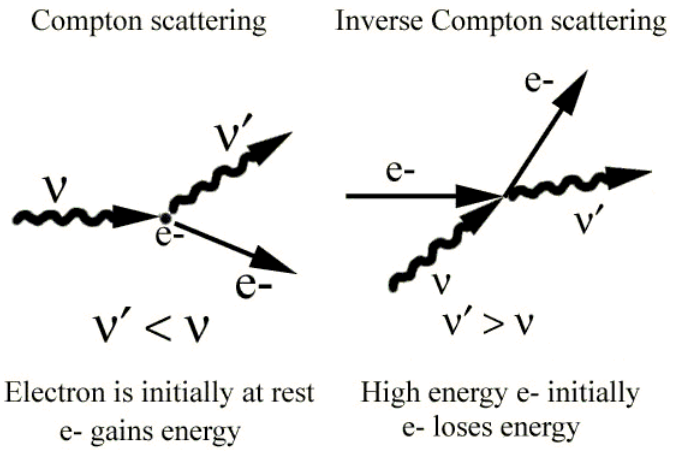

Figure 1.5: Illustration of the Compton and inverse Compton effects. An electron and a photon collide, and the photon either gains energy or loses it. Which effect takes place depends on the relative energies of the two particles. The inverse Compton effect is dominant in regions with high-energy electrons, such as in an AGN's corona. Credit: http://venables.asu.edu/quant/proj/compton.html

do not reach the temperatures required to emit X-rays through blackbody radiation. Rather, the X-rays we see from AGN are produced via inverse Compton scattering and synchrotron radiation, both of which are described in this section.

Compton scattering describes the scattering of a photon off of a particle, usually an electron. In ordinary Compton scattering, the photon imparts some energy to the electron, and therefore the photon's wavelength increases by an amount:

$$\Delta\lambda = \frac{h}{m_e c}(1 - cos\theta) \tag{1.1}$$

where h is Planck's constant, m_e is the mass of the electron, c is the speed of light, and θ is the angle at which the photon is scattered. The dependence on $\frac{1}{m}$ is why

astronomers are typically only concerned with the effects of electrons, since protons and neutrons are orders of magnitude more massive. If the energy of the electrons is far higher than that of the photons (i.e. $m_e c^2 >> h\nu$, the electrons are highly relativistic), then this process can happen in reverse. The electrons lose some of their energy, and the photons are up-scattered to higher energies, often orders of magnitude higher. The Compton effect, and its inverse, are illustrated in Figure 1.5. It can be shown (see, for example, Blumenthal and Gould 1970) that the power gained by a radiation field from the nearby electrons via inverse Compton scattering is:

$$P = \frac{4}{3}\sigma_T c \beta^2 \gamma^2 U_{rad} \tag{1.2}$$

where σ_T is the cross-section for Thompson scattering (the low-energy, elastic version of Compton scattering, this goes as $\frac{1}{m^2}$ and is the reason only electrons are considered), $\beta = \frac{v^2}{c^2}$ where v is the speed of the electron, the Lorentz factor is $\gamma = \frac{1}{\sqrt{1-\beta^2}}$, and U_{rad} is the energy density of the radiation field. Furthermore, it can be shown that for a power law distribution of electron energies:

$$dN \propto E^{-\Gamma} dE \tag{1.3}$$

(a reasonable assumption in most AGN), the observed spectrum is also of a power law shape. This process is the most likely explanation for the observed X-ray contin-

uum in most AGN, since it explains both the power law shape and the existence of X-ray photons with frequencies hundreds or thousands of times higher than that of the UV emission of the accretion disk.

Synchrotron radiation occurs when relativistic charged particles are accelerated in a magnetic field, causing them to emit photons along their direction of motion. Classically, any charged particle accelerated in a magnetic field will emit radiation. This is known as cyclotron emission. However, if the particles are moving at relativistic velocities, the radiation will be tightly beamed along the direction of motion, as illustrated in Figure 1.6 This beaming is the result of the Doppler effect and length contraction. The Doppler effect causes photons emitted along the particle's direction of motion to be seen at higher frequency, and length contraction stretches out the shape of the emitted radiation field in the observer's frame. These effects can cause the observed radiation to appear thousands of times stronger than if the emission was isotropic, and it is this effect that gives blazars their characteristic brightness.

As before, the contribution of protons to the luminosity is negligible due to their larger mass. The power emitted by a group of relativistic electrons in a magnetic field is:

$$P = \frac{4}{3}\sigma_T c\beta^2\gamma^2 U_B \tag{1.4}$$

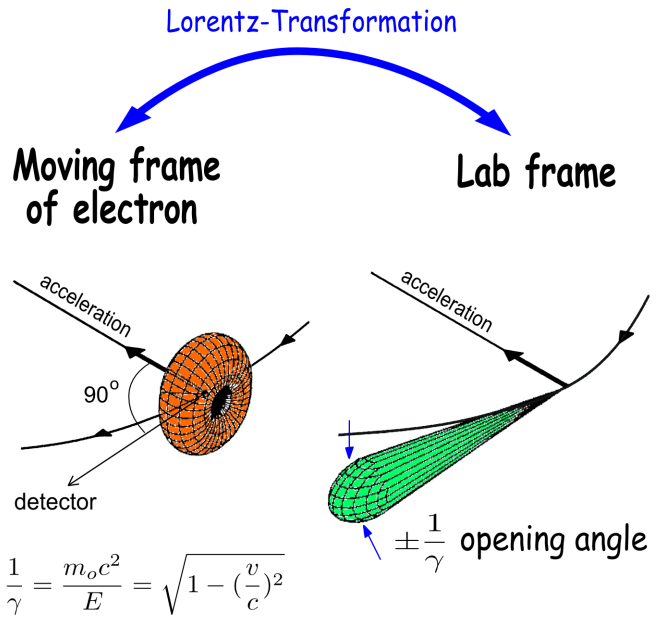

Figure 1.6: Illustration of synchrotron radiation. In the electron's rest frame, it emits power in a doughnut-like shape perpendicular to the direction of acceleration. However, due to the effects of length contraction and the Doppler effect, the radiation is shaped like a tightly-focused lobe in the observer's frame. This drastically magnifies the observed luminosity. Credit: http://photon-science.desy.de/research/students__teaching/primers/synchrotron_radiation/index_eng.html

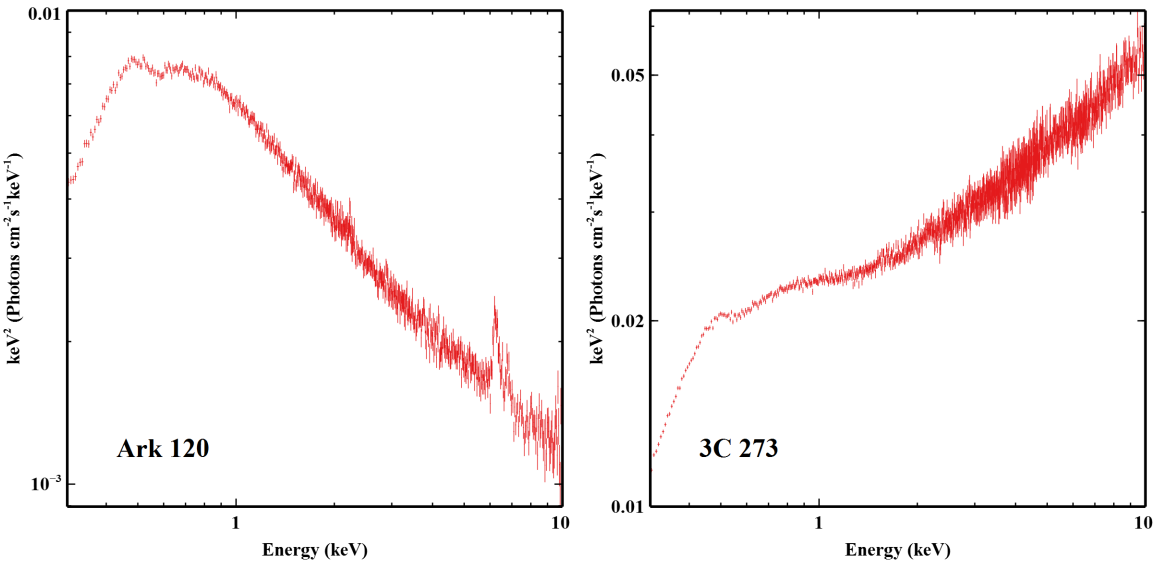

Figure 1.7: Left: X-ray spectrum of Ark 120, a type I radio-quiet AGN, unfolded against a power law with $\Gamma = 0$. While it mostly conforms to a power law shape, features such as the iron line near 6keV can still be seen. Right: Spectrum of the blazar 3C 273, in the same waveband, unfolded against the same power law. Any notable features have been drowned out by the power law shape of the jet.

where U_B is the energy density of the magnetic field. This expression is exactly the same as that for inverse Compton scattering, except with U_B replacing U_{rad}. In a sense, these processes are the same thing, with the electrons being accelerated by an electric field formed by the incident photons (in inverse Compton scattering) or by viewing a changing magnetic field from their own rest frame (in synchrotron radiation). Similarly, the spectrum produced by synchrotron radiation is also a power law.

The synchrotron self-Compton (SSC) process occurs when photons produced by the synchrotron process are inverse Compton-scattered by the very same electrons that emitted them. This process is very common in the jets of radio-loud AGN, and is responsible for the broad double-peaked spectrum of blazars (see Figure 1.4).

Figure 1.7 shows a comparison between the X-ray spectra of radio-quiet and radio-loud objects. The right panel shows the spectrum of Ark 120, a type I radio-quiet galaxy. The most noteworthy feature is a strong iron emission line around 6keV. This line is seen in almost all AGN. The left panel is a spectrum of 3C 273, a blazar. The power law shape of the jet has completely dominated the spectrum, and obscured any trace of even the strongest features originating outside of the jet.

1.5 X-ray telescopes and XMM-Newton

Due to their extreme energies, special equipment is needed to observe X-ray photons. In an optical telescope, the light is reflected off of the telescope's mirrors at the incident angle and reaches the focal point after one reflection. However, if an X-ray photon struck such a mirror, it would penetrate into the mirror rather than be reflected. To solve this problem, X-ray telescopes use a series of many mirrors directed at very shallow angles to the incoming radiation. Figure 1.8 shows an example of this system as used in the *XMM-Newton* telescope (Jansen et al. 2001).

XMM-Newton is a space-based telescope (as all X-ray observatories must be, since X-rays can not penetrate Earth's atmosphere) that was launched in 1999 but remains operational to this day. Because of its long lifespan, *XMM-Newton* has built up a large archive of data, all of which are made publicly available online after an exclusivity period. This makes *XMM* a useful tool for a project such as this one, since no new observations were required. This project focused on the 0.3 to 10 keV energy range,

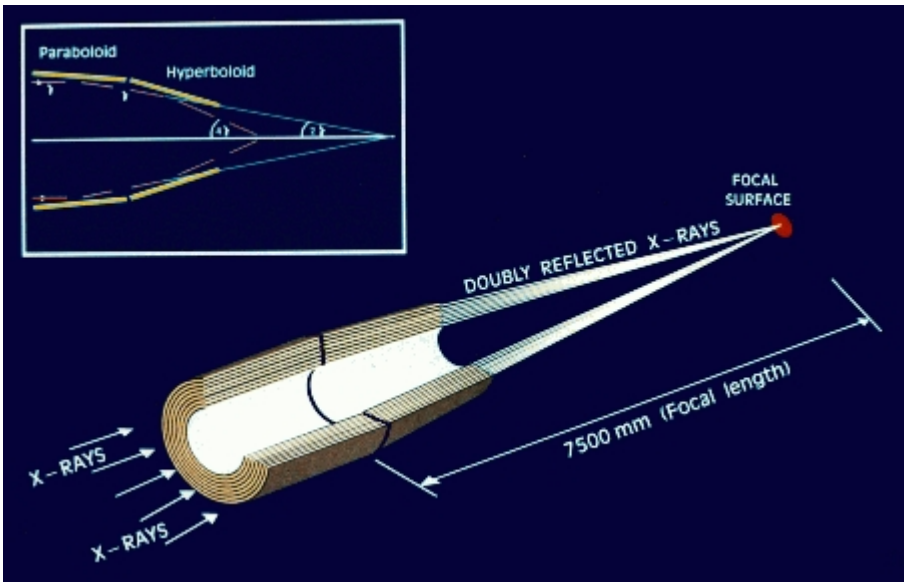

Figure 1.8: Diagram of the mirrors used in the *XMM-Newton* X-ray observatory. Nested mirror shells, both parabolic and hyperbolic, placed at shallow angles to the incident radiation, are used to carefully funnel X-rays towards the focal point. Credit: http://sci.esa.int/xmm-newton/39319-light-path-in-xmm-newton-telescope/

where the *XMM* pn-CCD camera is most efficient. This range covers the lower-energy (or softer) X-rays, as well as the power-law spectrum created by the corona emission in radio-quiet objects or synchrotron and SSC emission in the jet for blazars (although in radio-loud objects that are not blazars, often both power laws can be seen because the jet will not drown out the rest of the spectrum). Chapter 3 discusses the data sample in detail and explains how the public *XMM-Newton* data were downloaded and processed.

1.6 Principal Component Analysis

Principal Component Analysis (PCA) is a data analysis technique that is used to describe variability within a data set. It takes a data set and expresses it in a new basis

that reveals the underlying trends in the variability between data points. Because of their highly variable nature, PCA is well-suited to the study of AGN.

To picture how PCA works, imagine a system with a simple, underlying physical component, such as a spring oscillating in the x-direction. However, imagine that we do not know anything about this motion or how it can be described, and so the data we collect is not aligned with the x-axis. We suspect that there exist one (or several) vectors that describe the motion, but we do not know what they are. We need to both determine how many vectors are needed to describe this motion (i.e. account for redundancies in our data) and what those vectors are. This situation is illustrated in Figure 1.9.

When choosing a new basis for a data set, PCA looks for the directions along which the data has the largest variance and defines them as the principal components. In the example above, it is clear from each camera's data that motion along one axis has the largest variance, and therefore that direction is defined as the first principal component. However, the process does not stop there. Because this is a change of basis, and not a form of data reduction, the output data set will have the same dimensionality as it did before the PCA. In other words, if the input data set has two dimension (such as the two axes of one of our cameras), there will be two principal components.

The first component will be aligned along the direction of greatest variance, but

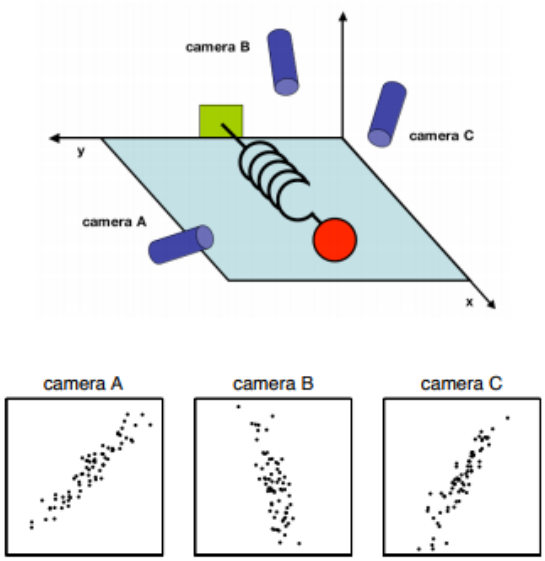

Figure 1.9: Illustration of a scenario in which PCA could be used. The object is oscillating along one direction, and its motion could be entirely described using a single basis vector aligned with the x-axis. However, we do not know this, and so our data is both redundant and taken at an angle to the motion. Credit: *A Tutorial on Principal Component Analysis* by Jonathon Shlens (2014)

the second, third, and subsequent components will not simply be the directions of second and third highest variance. If this were true, then all of the components would be virtually identical. Instead, a crucial assumption is made: the principal components are mutually orthogonal. This means that the second principal component will be the direction of greatest variance among directions orthogonal to the first, and the third will be the direction of greatest variance among those orthogonal to both of the other components, and so on. This assumption ensures that the principal components are meaningfully different from each other, and in most real situations it is a reasonable assumption. The main reason to make this assumption, though, is that it ensures that there is an analytical solution for the PCA process for any input data set, of any shape or size. This solution is presented in detail in Chapter 2.

Applying PCA to astronomical data is slightly more complicated, but the idea is the same. Imagine a set of spectra from a particular object. These spectra could be expressed as a linear combination of various spectral components (such as a power law component, a blackbody component, and so on), but we do not know what those components are. Using PCA, the spectra can be decomposed into their composite parts. An example of this is shown in Figure 1.10. We start with three functions representing physical components of a hypothetical spectrum. Then, synthetic spectra are created by mixing together a linear combination of the three components and adding some noise. The last panel of Figure 1.10 shows the results after performing PCA on the generated spectra. The original components have been extracted from the spectrum, albeit flipped on their horizontal axis. The sign of the value on the y-axis of the principal components is arbitrary, as the process cares only about whether a component is moving towards or away from 0, not whether it is positive or negative.

For an example of PCA as applied to real data, see Figure 1.11. This is a PCA of the blazar 3C 273, and is representative of the results for most blazars. Twenty-seven spectra from various observations, spread out across more than ten years, were used to calculate this PCA. This results in twenty-seven principal components, although only the first three are plotted. The first one is uniformly below zero, indicating that all energy bands varied in the same direction and by the same amount. This is consistent with changes in a normalization factor. The second component has a pivoting shape. This means that when the flux increased at lower energies, it cor-

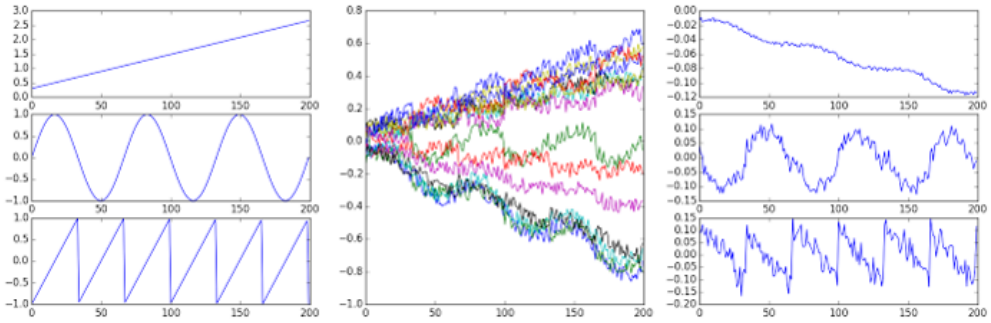

Figure 1.10: Example of the PCA process as applied to spectra. Left: Three physical components. Centre: Various spectra created by taking a linear combination of the physical components and adding some noise. Right: The three principal components returned by the PCA process. Note that the orientation of the y-axis is arbitrary; what matters is whether the component is moving towards or away from 0. Credit: PCA code by Michael Parker, http://www-xray.ast.cam.ac.uk/mlparker/.

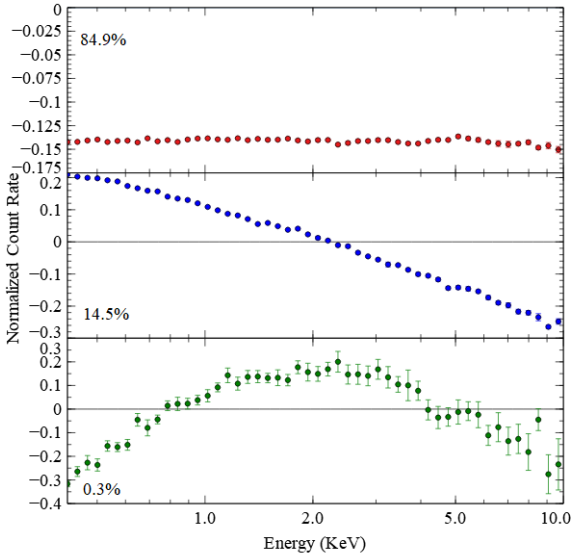

Figure 1.11: Example of PCA applied to various spectra from observations of the blazar 3C 273. The first, flat component corresponds to a change in normalization. The second, pivoting component indicates a change in the shape of the power law, and the third component is likely a mathematical artefact caused by non-linearities within the data set. The percentages given are the fraction of the total variability that each component is responsible for. The total is less than 100% because there are 24 other components not shown here.

responded to a decrease at higher energies, and vice versa. This is consistent with varying the index (exponent) of a power law. Together, these first two components account for over 99.4% of the total variability. This is expected, since blazar spectra are usually dominated by the power law from the jet. The last component has no obvious physical explanation, and is very insignificant compared to the others. It is most likely a mathematical artefact of the PCA process. In Chapter 6 it is argued that this component arises from non-linear relationships within the data set.

This work uses PCA to compare and contrast a variety of blazar spectra and examine the variability trends that drive them. Blazar spectra are relatively simple, being dominated by the power law shape of the jet. Because of that simplicity, this work will hopefully serve as an accessible introduction to the topic of PCA as it applies to astronomy, and a demonstration of the power of PCA as an analytical tool. As explained previously, the focus will be on X-ray data. The analysis is split into two types: long-term PCA, where many spectra taken over the course of years or even decades are analysed, and short-term PCA, where a single observation is split into several spectra and then analysed. The long-term analysis reveals variability trends over long timescales, whereas the short-term analysis examines variability over the course of mere hours. By comparing the results, we can determine whether these two types of variability are driven by the same processes.

PCA is used in many academic fields, and has recently seen significant application to X-ray astronomy. Francis & Willis (1999) provide an introduction to PCA of

AGN, while Grupe et al. (1999) and Grupe (2004) demonstrate some early uses of PCA applied to ROSAT data. More recently, PCA has been applied to *XMM-Newton* data to examine X-ray spectral variability in detail (for example, Vaughan & Fabian 2004, Miller et al. 2007, and Turner et al. 2007). In addition to AGN, PCA has been used to study variability in objects such as X-ray binaries (Malzac et al. 2006; Koljonen et al. 2013). Today, with almost two decades of high-quality *XMM-Newton* data available, PCA can be used to reveal variability trends within AGN over long timescales, as in Parker et al. (2015)

Parker et al. (2015) applied principal component analysis to a wide range of AGN, mostly radio-quiet, looking at long-term variability with *XMM-Newton* data. They found that most objects displayed variability in a power law continuum, but prominent variations in reflection components (in MCG-6-30-15, NGC 4051, 1H0707-495, NGC 3516, and Mrk 766) and partial covering absorption (in NGC 4395, NGC 1365, and NGC 4151) were also common. Their AGN displayed between three and five principal components, with evidence for many qualitatively different variability mechanisms. Some other sources can show long-term changes associated with emission from the distant torus (e.g. Gallo et al. 2015).

1.7 This Work

This work applies a similar analysis to that of Parker et al (2015) to a smaller sample of objects, specifically blazars, across both long and short timescales. Blazars were chosen due to their simple spectra and rapid variability, as well as the lack of PCA results for many well-known blazars. Blazar spectra are dominated by the effects of the jet, which follows a synchrotron self-Compton shape (Mastichiadis 1997), resulting in X-ray spectra that conform closely to a single power law model. By applying PCA to objects that are already known to be spectrally simple, we can better understand the intricacies of this technique and examine what drives blazar variability.

Chapter 2 provides a detailed explanation of the math behind PCA. In Chapter 3, we describe our sample and data analysis. Chapter 4 presents the long-term, or multi-epoch, PCA results. These PCAs use all available observations of an object to describe variability over the span of years. Chapter 5 presents the short-term, or single-observation, results. These PCAs take a longer observation of an object and divide it into several spectra in order to observe variability over the span of hours. In Chapter 6, we discuss PCAs applied to simulated data in order to compare models to the real results. Lastly, Chapters 7 and 8 summarize our results and present the conclusions.

Chapter 2

PCA In Detail

Section 1.6 described the PCA process qualitatively. This section will examine the math behind it. This solution is adapted from Shlens (2014).

We will start with a data set X, an m x n matrix where m is the number of different measurements (a number of energy bins, in our case), and n is the total number of data points in each of those bins (one for each spectrum). We want to perform a change of basis that will transform this data set such that there are as few correlations within the data as possible. An important concept for this analysis is *covariance*. The covariance between two variables x and y, each with N data points, is defined as:

$$cov(x, y) = \sum_{i=1}^{N} \frac{(x_i - \bar{x})(y_i - \bar{y})}{N} \tag{2.1}$$

where $\bar{\text{x}}$ and $\bar{\text{y}}$ are the means of x and y, respectively. The covariance is a measure of how correlated the two variables are. The *covariance matrix* of our data set X is:

$$C_X = \frac{1}{n} X X^T \tag{2.2}$$

where X^T is the transpose of X. The diagonal terms of C_X are the variances of particular measurement types. The off-diagonal terms are the covariances between measurement types. If the covariance between two variables is zero, those variables are uncorrelated. Our goal is to take X and transform it into a set of principal components, P, with as little redundancy between them as possible. This corresponds to diagonalizing C_X, since a diagonal covariance matrix would have no correlation between any of its variables.

Diagonalizing the covariance matrix amounts to finding an orthonormal matrix P such that C_y is diagonalized, where $Y = PX$. The orthogonality of P is chosen for reasons outlined earlier. The rows of P are the principal components of X, and the diagonals of C_y are the fractional variability of the corresponding principal component. Let us express C_y in terms of P:

$$\begin{aligned} C_y &= \tfrac{1}{n} Y Y^T \\ &= \tfrac{1}{n} (PX)(PX)^T \\ &= \tfrac{1}{n} P X X^T P^T \\ &= P(\tfrac{1}{n} X X^T) P^T \\ C_y &= P C_x P^T \end{aligned}$$

This has revealed C_x, the covariance of the data set.

Continuing, recognize that for any symmetric matrix A, $A = EDE^T$, where D is a diagonal matrix and E is a matrix of the eigenvectors of A arranged as columns. If we choose P to be a matrix where each row is an eigenvector of C_x, then $P = E^T$ with respect to C_x. This lets us show that:

$$\begin{aligned} C_y &= PC_xP^T \\ &= P(E^TDE)P^T \\ &= P(P^TDP)P^T \\ &= (PP^T)D(PP^T) \\ C_y &= D \end{aligned}$$

We have shown that this choice of P diagonalizes C_y. Since the rows of P are the principal components of X, the eigenvectors of C_x are the principal components. Therefore, calculating the principal components of a data set X amounts to finding the eigenvectors of that data set's covariance matrix.

The eigenvectors of C_x are calculated using singular value decomposition (SVD). SVD takes a matrix Z and express it in the form:

$$Z = UDV^T$$

where D is a diagonal matrix and U and V are orthogonal matrices. Furthermore, the columns of V are the eigenvectors of $Z^T Z$. Now, to use SVD to perform our PCA, we define Z as $Z = \frac{1}{\sqrt{n}} X^T$ and notice that this means $Z^T Z = C_x$. Then we can easily calculate the eigenvectors of C_x (and therefore the principal components) by performing SVD on Z. The eigenvectors we're looking for will be the columns of V. In this way, it becomes clear that PCA and SVD are intimately related techniques.

Chapter 3

Sample and Data Processing

Objects were selected from among those in Costamante & Ghisellini (2002) with publicly available *XMM-Newton* data (Jansen et al 2001), as well as the well-known object 3C 273. Observation dates ranged from May 2000 to May 2017. Only the highest signal-to-noise EPIC-pn instrument (Struder et al. 2001) was used.

Observations where the target was significantly off-axis were excluded. The data were collected in a variety of window modes and optical filters. Data in timing mode were not used, due to the uncertain calibration of this mode. For each observation, the observation data files (ODFs) were downloaded from the *XMM-Newton* Science Archives and processed to create spectra using the Science Analysis System (SAS) version 15.0.0

Epchain was used to generate event lists from the ODFs, and the spectra were made using a source region with a radius of 35 arcseconds. Background subtraction was performed using a background region of radius 50 arcseconds located near the source.

Each observation was checked for pileup, and some showed significant amounts of

it. This was corrected for by extracting the source spectrum from an annulus with the same outer radius, and an inner radius of 8 arcseconds, which excludes the most highly piled-up light from the centre of the object. To ensure that this corrective technique did not influence the results, PCAs of piled-up objects were performed both with and without the piled-up observations. Other than showing more noise due to the lower sample size, this caused no major difference in the shapes or significance of the principal components.

Some observations displayed high levels of background flaring at certain times. These observations were filtered through a good time interval (GTI) that excluded the times when the flaring occurred. Response matrices and ancillary response files were created using RMFGEN and ARFGEN, respectively.

Figure 3.1 presents representative spectra for each object, unfolded against a power law with $\Gamma = 0$. The complete list of observations is shown in Table 3.1.

Table 3.1: Complete list of observations

Object	Obs ID	Revolution	Start Time	Duration	GTI	0.3-10 keV Count Rate (s^{-1})	Pileup Correction?	Window Mode
3C 273	0112770101	370	2001-12-16 15:35:23	6399	3507	64.3	N	Small
	0112770201	373	2001-12-22 00: 19:58	6399	3471	62.23	N	Small
	0112770501	655	2003-07-08 10:33:51	8553	5631	62.67	N	Small
	0112770601	472	2002-07-07 14:25:05	5996	3504	47.91	N	Small
	0112770701	563	2003-01-05 17:24:04	5630	3503	58.32	N	Small
	0112770801	554	2002-12-17 22:24:56	5624	3503	69.37	N	Small
	0112771001	645	2003-06-18 01:07:13	5950	3861	70.72	N	Small
	0112771101	735	2003-12-14 19:23:21	12849	5928	47.87	N	Small
	0126700301	94	2000-06-13 23:39:53	73556	45260	42.02	N	Small
	0126700601	95	2000-06-15 12:58:18	31032	20820	40.46	N	Small
	0126700701	95	2000-06-15 23:32:02	36346	21030	39.29	N	Small
	0126700801	96	2000-06-17 23:24:14	73561	42510	45.52	N	Small
	0136550101	277	2001-06-13 07:14:26	89765	62000	53.65	N	Small
	0136550501	563	2003-01-05 14:17:24	8951	5965	66.58	N	Small
	0136550801	835	2004-06-30 13:02:25	62913	13910	40.40	N	Small
	0136551001	1023	2005-07-10 13:51:19	28111	19330	44.24	N	Small
	0159960101	655	2003-07-07 17:40:27	58557	40600	63.48	N	Small
	0414190101	1299	2007-01-12 07:13:55	78566	53710	49.47	N	Small
	0414190301	1381	2007-06-25 05:08:14	32511	22440	40.83	N	Small
	0414190401	1465	2007-12-08 20:11:25	35875	24820	81.17	N	Small
	0414190501	1649	2008-12-09 20:12:31	41015	28420	57.01	N	Small
	0414190601	1837	2009-12-20 03:42:44	31912	22030	62.46	N	Small
	0414190701	2015	2010-12-10 01:37:45	36414	25210	46.88	N	Small
	0414190801	2199	2011-12-12 17:44:21	43915	30380	42.22	N	Small
	0414191001	2308	2012-07-16 11:59:23	38918	17760	36.70	N	Small
	0414191101	2856	2015-07-13 21:03:55	72400	49680	31.73	N	Small
	0414191201	3031	2016-06-26 20:22:08	67200	46030	55.51	N	Small
3C 279	0651610101	2035	2011-01-18 16:49:52	126346	86960	49.48	N	Small
H1426+428	0111850201	278	2001-06-16 00:49:21	68574	45770	16.83	N	Small
	0165770101	852	2004-08-04 00:59:26	67866	45860	20.13	N	Small
	0165770201	853	2004-08-06 00:32:43	68920	47980	20.06	N	Small
	0212090201	939	2005-01-24 14:44:40	30417	20960	25.13	N	Small
	0310190101	1012	2005-06-19 07:39:40	47034	32680	36.97	N	Small
	0310190201	1015	2005-06-25 06:03:28	49505	31140	28.79	N	Small
	0310190501	1035	2005-08-04 04:52:10	47542	32410	28.47	N	Small
Mrk 421	0099280101	84	2000-05-25 03:17:11	66497	21160	216.8	N	Small
	0099280201	165	2000-11-01 23:47:51	40115	24240	112.5	Y	Small
	0099280301	171	2000-11-13 22:00:29	49811	25640	279.7	N	Small
	0136540101	259	2001-05-08 09:09:35	39007	25730	144.5	Y	Small
	0136540301	532	2002-11-04 00:44:59	23913	13830	23.40	N	Full-frame
	0136540401	532	2002-11-04 07:41:43	23917	14180	43.51	Y	Full-frame
	0136540701	537	2002-11-14 00:07:35	71520	37970	97.45	Y	Large
	0153950601	440	2002-05-04 16:09:17	39727	34330	25.75	Y	Large
	0153950701	440	2002-05-04 03:51:30	19982	15940	16.55	Y	Large
	0158970101	637	2003-06-01 11:33:26	47538	24920	103.3	Y	Small
	0162960101	733	2003-12-10 21:23:14	50755	16470	119.8	Y	Small
	0411081301	1358	2007-05-10 03:37:41	18913	13960	37.50	Y	Full-frame
	0411083201	1820	2009-11-16 17:37:59	58070	7526	112.3	Y	Large
	0560980101	1640	2008-11-22 14:07:29	71318	8479	51.67	Y	Large
	0560983301	1732	2009-05-25 03:37:32	64173	8468	63.64	Y	Large
	0656380101	1904	2010-05-03 07:19:29	51169	6619	91.76	Y	Large
	0656380801	2001	2010-11-12 20:51:05	42669	7628	66.98	Y	Large
	0658800101	2094	2011-05-19 10:02:48	35074	8941	38.94	Y	Large
	0658801301	2837	2015-06-05 23:48:35	29000	19270	105.6	Y	Small
	0658801801	2915	2015-11-08 13:42:37	33600	21200	81.51	Y	Small
	0658802301	3005	2016-05-06 03:38:20	29400	19540	72.91	Y	Small
	0791780101	3096	2016-11-03 13:15:45	17500	11210	61.34	N	Small
	0791780601	3187	2017-05-04 04:01:33	12500	7708	153.7	N	Small

Table 3.1: Continued

Object	Obs ID	Revolution	Start Time	Duration	GTI	Count Rate (s^{-1})	Pileup Correction?	Window Mode
Mrk 501	0113060401	475	2002-07-14 17:02:39	15769	2945	0.1304	N	Small
	0652570101	1969	2010-09-08 23:50:27	44912	31160	26.47	N	Small
	0652570201	1970	2010-09-10 23:42:24	44919	31160	27.39	N	Small
	0652570301	2047	2011-02-11 14:43:25	40914	28350	28.86	N	Small
	0652570401	2049	2011-02-15 14:18:29	40715	28220	37.72	N	Small
OJ 287	0300480201	978	2005-04-12 13:13:21	38913	9918	1.388	N	Large
	0300480301	1081	2005-11-03 21:16:31	48059	28800	1.073	N	Large
	0401060201	1271	2006-11-17 00:33:10	47211	41360	0.8658	N	Large
	0502630201	1533	2008-04-22 17:13:34	55815	48100	0.8675	N	Large
	0679380701	2170	2011-10-15 08:18:19	23917	20150	2.938	N	Large
	0761500201	2822	2015-05-07 05:23:25	129200	94890	1.871	N	Large
PG 1553+113	0656990101	1952	2010-08-06 12:38:17	21914	15050	15.19	N	Small
	0727780101	2495	2013-07-24 14:57:49	34500	23120	28.98	N	Small
	0727780201	2680	2014-07-28 04:00:06	36300	24380	17.25	N	Small
	0727780301	2882	2015-09-04 18:23:24	29999	19960	10.24	N	Small
	0727780401	3057	2016-08-17 21:56:06	30000	19960	12.33	N	Small
	0761100101	2864	2015-07-29 19:57:33	138400	119700	5.569	Y	Full-frame
	0761100201	2866	2015-08-02 19:40:00	138900	119000	4.575	Y	Full-frame
	0761100301	2867	2015-08-04 19:32:00	138900	19960	10.24	N	Small
	0761100401	2869	2015-08-08 19:12:07	138900	117700	4.348	Y	Full-frame
	0761100701	2873	2015-08-30 18:52:06	90000	62010	8.622	N	Small
	0761101001	2880	2015-08-30 17:52:29	139000	117200	5.798	Y	Full-frame
PKS 2155-304	0080940101	174	2000-11-19 18:38:20	60511	40190	16.75	Y	Small
	0080940301	174	2000-11-20 12:53:01	61411	40810	58.16	N	Small
	0124930201	87	2000-05-31 00:30:51	72558	41580	77.13	N	Small
	0124930301	362	2001-11-30 02:36:09	92617	31260	79.04	Y	Small
	0124930501	450	2002-05-24 09:31:02	104868	22300	55.68	N	Small
	0124930601	545	2002-11-29 23:27:28	114675	39790	29.71	N	Small
	0158960101	724	2003-11-23 00:46:22	27159	18670	27.55	N	Small
	0158960901	908	2004-11-22 21:35:30	28919	19960	30.90	N	Small
	0158961001	908	2004-11-23 19:45:55	40419	27960	40.12	N	Small
	0158961101	993	2005-05-12 12:51:06	28910	19250	69.63	N	Small
	0158961301	1095	2005-11-30 20:34:03	60415	41900	76.16	N	Small
	0411780101	1266	2006-11-07 00:22:47	101012	20870	42.28	N	Small
	0411780201	1349	2007-04-22 04:07:23	67911	43360	74.47	N	Small
	0411780301	1543	2008-05-12 15:02:34	61216	42600	89.11	N	Small
	0411780401	1734	2009-05-28 08:08:42	64820	45100	62.32	N	Small
	0411780501	1902	2010-04-29 20:26:00	74298	47730	31.87	N	Small
	0411780601	2084	2011-04-26 13:50:40	63818	44400	49.04	N	Small
	0411780701	2268	2012-04-28 00:48:26	68735	38660	12.56	N	Small
	0411782101	2449	2013-04-23 22:31:38	76015	48830	27.62	N	Small
	0727770901	2633	2014-04-25 03:14:56	65000	44500	29.68	N	Small
S5 0716+714	0502271401	1427	2007-09-24 16:23:32	73917	50120	4.269	N	Small

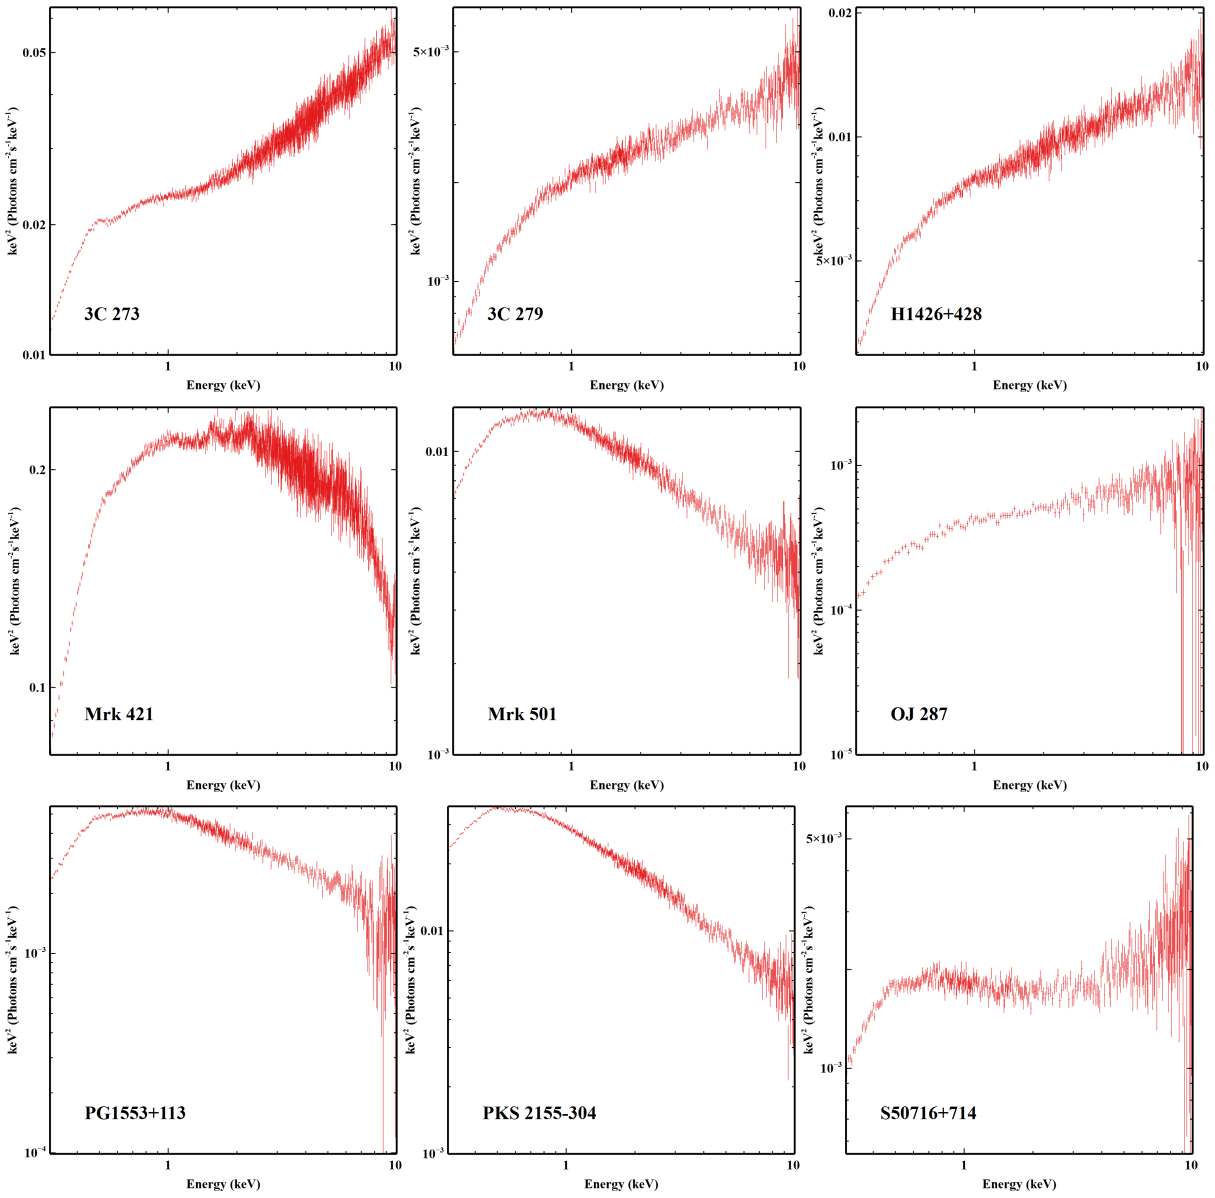

Figure 3.1: Representative spectra for each object, unfolded against a power law with $\Gamma = 0$.

Chapter 4

Long-term (multi-epoch) variability

This section presents the results of PCA performed on a single object across many different observations. For objects with more than four separate observations (of any duration), this PCA was calculated in order to examine long-term variability. Seven objects from our sample met this criteria: 3C 273, H1426+428, Mrk 421, Mrk 501, OJ 287, PG1553+113, and PKS 2155-304.

These observations span a minimum of four years (for H1426+428) to a maximum of seventeen years (for Mrk 421). The observations were not evenly spread out in time, with some sources being observed much more frequently than others. Each observation for a given object corresponded to a single spectrum used in the PCA, as opposed to splitting each observation into chunks as in Parker et al.(2015) or the short-term PCAs presented in Section 5. 3C 273 was the most sampled object, with 27 observations over fifteen years. Figure 4.1 shows the spectral variability of this object by comparing all 27 observations to an average power law fit.

Figure 4.2 shows the results of the long-term PCA analysis for each object. The first three principal components are plotted in decreasing order of the fraction of the

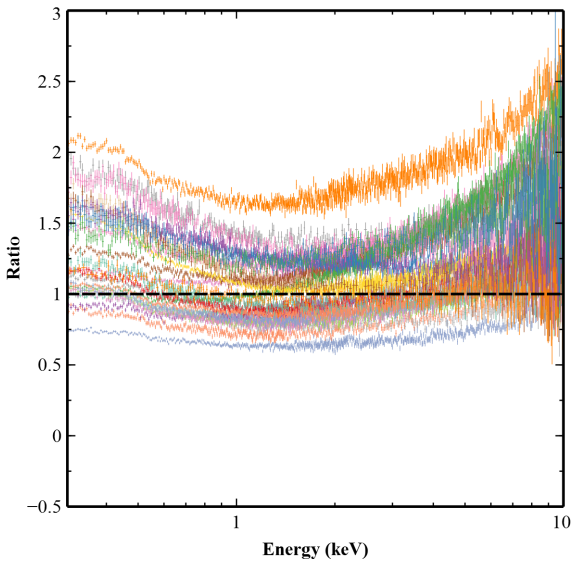

Figure 4.1: Ratio of 27 spectra of 3C 273 spanning from 2001 to 2016 to an average power law fit. The source is highly variable, both in the shape of the spectrum and the total flux. This is typical of blazars.

total variability that they are responsible for, expressed as a percentage in the plot. In every case, the remaining components showed no discernible shape and were not significant compared to the first three, and so are not plotted.

The results were very similar for every object, independent of the number or duration of observations, the time between observations, or the brightness of the object. In every case, the first principal component is uniformly above zero, meaning that all energy bands varied in a correlated manner. Compared to models generated in Parker et al (2015), this indicates a change in the overall normalization of the spectrum.

The second component in each object shows an anti-correlation between flux changes in the low and high energy bands. This is consistent with a pivoting of

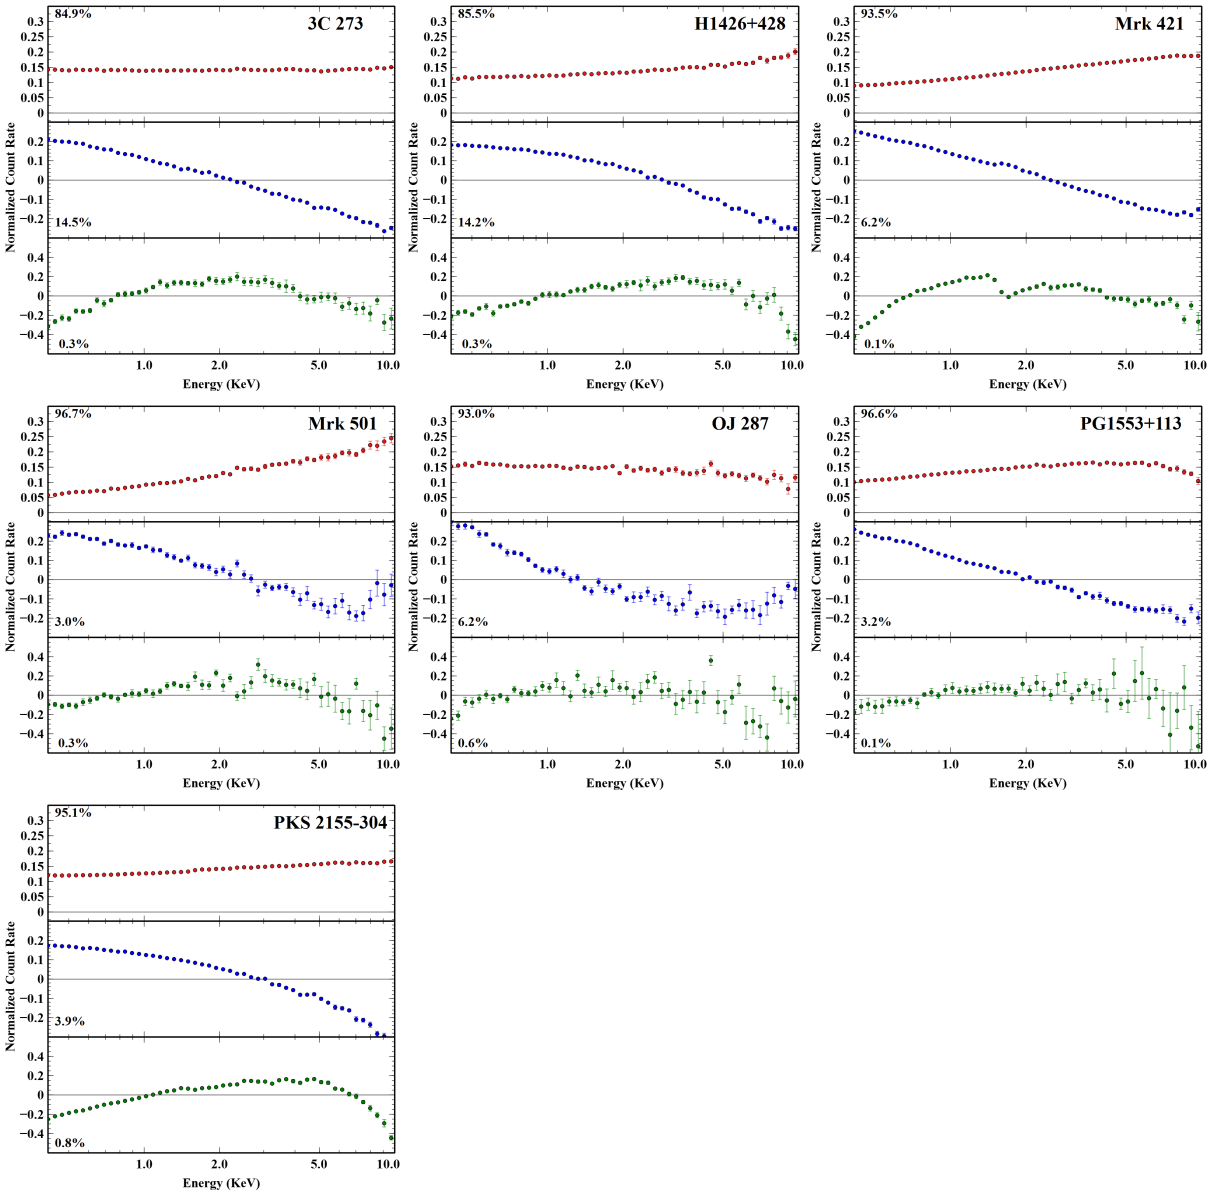

Figure 4.2: Long-term PCAs for each of the seven objects that had at least four separate observations at different epochs. All show similar results. The first component is uniformly above zero and mostly flat, indicating changes in normalization. The second component shows an anti-correlation between low and high energies, consistent with changes in the photon index of a power law. The third component is shaped like an arch, and has no obvious physical explanation. Instead, it is likely a mathematical artefact of the PCA process caused by a breakdown of the linearity assumption (see Section 6). Components beyond the third were not significant in any object. These results indicate that long-term variability in blazars is dominated by changes in a power law model, varying both in shape and normalization.

the spectrum brought about by changes in the photon index of a power law model.

The third component has an arch-like shape, meaning energy changes in the low and high energies were correlated with each other and anti-correlated with changes in the energy band between them. This shape has no obvious physical explanation, and is probably a mathematical artefact of the PCA process. This is investigated further in Section 6.

These three components are consistent with a power law model varying in both normalization and Γ. The primary component is always due to changes in normalization for our sample. This accounts for most of the variability (> 84 per cent) for our objects. The second component accounts for much less of the variability (3-15 per cent) and is attributed to pivoting of the power law. For our sample of blazars, the long-term X-ray variability over years appears to be dominated by changes in the brightness of the source, and less so by changes in the shape of the spectrum.

These power law components are to be expected of blazars, whose spectra are dominated by the effects of the jet, a highly-variable feature with a prominent power law shape. Note that this does not guarantee that a lone power law is sufficient to model any of these objects, only that the power law component is responsible for the vast majority of the variability. Although these long-term PCAs are not identical, many of the differences between them can be explained without introducing a more complicated model by changing the degree to which Γ and normalization are assumed

to be varying (see Section 6).

Additional comments on each individual PCA are presented in Appendix A.

Chapter 5

Short-term (single-observation) variability

This analysis was performed on observations with at least 40ks of good time. These observations were split into 10ks parts, which comprised the input spectra for the PCA. The results show X-ray variability over timescales as short as a few hours. In many cases, only one component was significant, whereas there were always three significant components in the long-term analysis.

There were several observations that showed no discernible shape in any component, indicating that the object was either constant at that time, or varying on timescales longer than the observation itself. These are not plotted.

The remaining short-term PCAs can be divided into two groups: those with a shape similar to those seen in the long-term analysis, presented in Figure 5.1, and those showing shapes unique to the short-term analysis, shown in Figure 5.2.

For the first case, shapes corresponding to normalization and Γ are seen, but their shapes are not as well-defined as those in the long-term cases. This could be caused by the smaller data set, but could also be an indicator of weaker variations. It is also

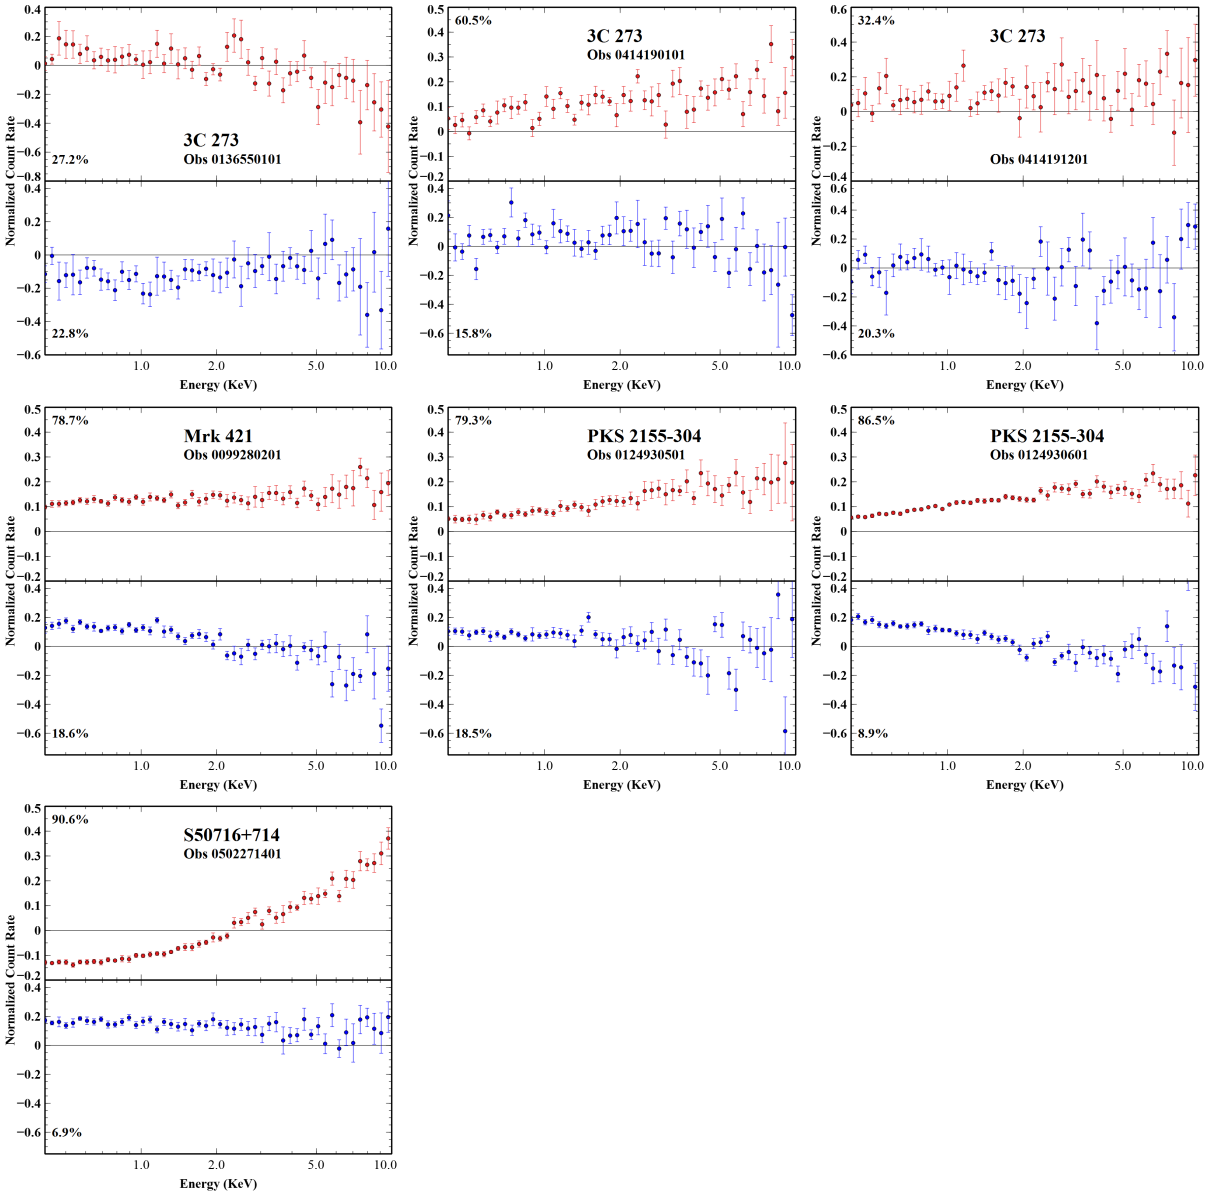

Figure 5.1: Short-term PCAs that display components similar to those seen in the long-term analysis, corresponding to changes in the normalization and photon index of a power law. The results are not as pronounced as they are in the long-term case, which could be due to smaller variations, or simply a result of lower signal-to-noise due to having less data.

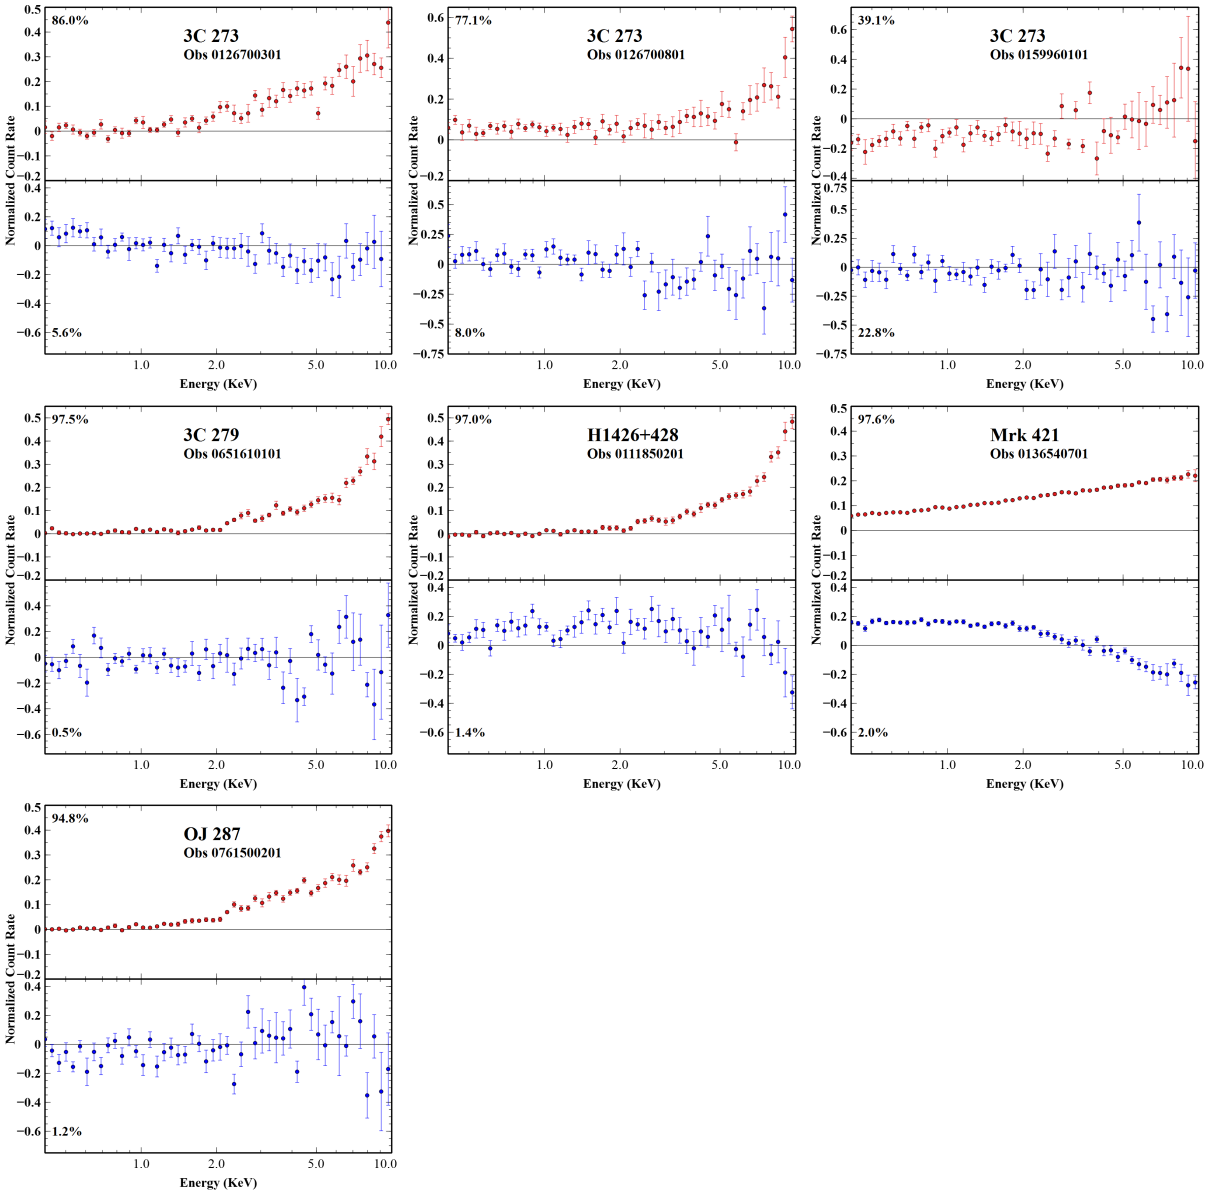

Figure 5.2: Short-term PCAs with shapes unlike those seen in the long-term analysis.

worth noting that the third principal component seen in the long-term observations does not appear with any significance in the short-term, even in observations with both a normalization and a pivoting component.

Overall, these PCAs are indicative of the same sorts of changes seen over long timescales, and the long-term results could be seen as the summation of many years worth of these short-term changes.

The second group exhibits components, that differ from the straightforward pivoting power law interpretation. Many of them look like broken power laws, showing no significant variability below some low energy (usually around 2 keV) and then increasing variability with increasing energy. Some of those show an additional upward curve, displaying an almost exponential shape. Shapes like these are not seen in the long-term PCAs, which means they are insignificant over long timescales. It is unclear what, if anything, these components correspond to physically, and simulations have been unable to replicate the upward-curving shape. To ensure that these shapes were not influenced by background effects, some short-term PCAs were performed again with background subtraction turned off. This did not significantly effect the results beyond introducing more noise, indicating that these shapes are part of the source spectrum.

We also note that some objects fall in both groups. For example, 3C 273 sometimes exhibits rapid variability consistent with the long-term (yearly) variations, while also

having epochs where the PCA spectral shape is unique.

Additional comments on each observation's PCA are presented in Appendix B.

Chapter 6

MODELS

The model-independent nature of PCA is useful, but it has its drawbacks. One weakness is that it is often unclear what the resultant principal components correspond to physically. Simulations can help with this. In this section, PCA is performed on a set of fake spectra generated according to a given model, with each spectrum varying the model parameters randomly within a certain range. Performing PCA on a known model allows for comparisons to the real data to be made and can identify how shapes are associated with physical parameters. As one would expect, simulations of a power law varying in both normalization and Γ can closely reproduce the results of the long-term PCAs presented in Section 4 (Parker et al 2015). Furthermore, many of the differences between the various PCAs can be reproduced by changing the ways in which the model parameters vary in relation to each other.

Figure 6.1 shows the results of PCAs performed on three sets of 100 simulated spectra conforming to a power law model varying in both normalization and Γ. In the first PCA, Γ varied randomly by up to 10 per cent, while normalization varied randomly by up to a factor of four. As seen in the long-term PCAs, there are three significant components: a flat component representing changes in normalization, a pivoting component representing the changing slope of the power law, and an arch-

shaped third component. These are the archetypal power law results that explain most of the shape in the long-term PCAs.

The second simulation put more emphasis on Γ, allowing it to vary by up to 25 per cent. This induces an upward slope on the first component, which is seen in several of our objects (H1426+428, Mrk 421, Mrk 501, and PKS 2155-304). Increasing the amount Γ varies increases this slope of the first principal component. The final simulation in Figure 6.1 varied the parameters as in the first simulation, except they now varied in a correlated manner, rather than independently. Changes were correlated such that normalization increased or decreased as Γ increased or decreased, resulting in a softening of the spectra as their brightness increased. This induces a negative slope in the first component, as well as weakening the second component slightly at high energies, and suppressing the third component entirely. All three of these effects are seen to some degree in OJ 287.

While simulations can do a good job of reproducing PCA shapes, each component's fractional contribution to the total variability is less easy to simulate. In simulations of a varying power law, the first component generally accounts for >90 per cent of the variability, whereas the third component is responsible for only a tiny fraction, even compared to the results from real data. Because of the presence of noise in the real data, and the fact that the noise contributes significantly to the total variability, it is difficult to reproduce the correct share of the total variability for each component.

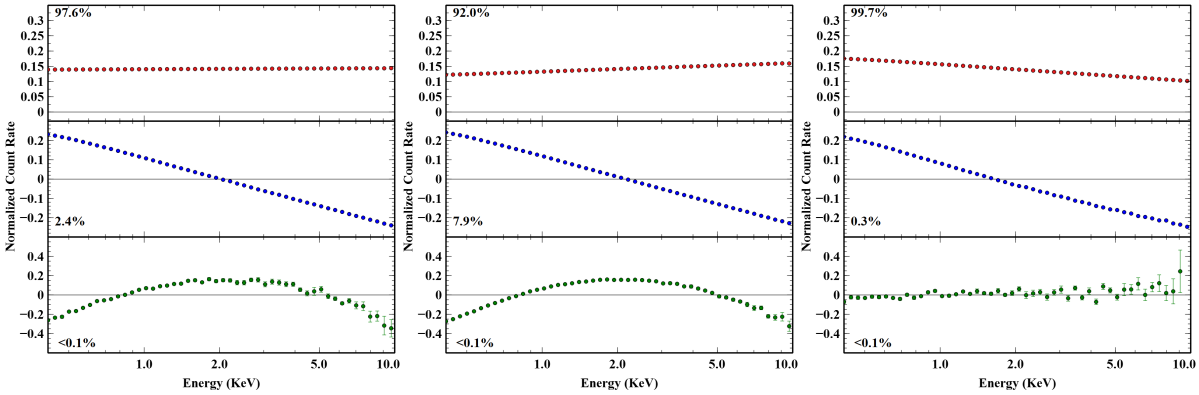

Figure 6.1: Results of PCA performed on 100 simulated power law spectra varying in normalization and Γ. Left: Γ varied randomly by up to 10 per cent, normalization varied by up to a factor of four. Middle: Γ varied by up to 25 per cent, while normalization still varied by up to a factor of four. Right: Same as in the left panel, except the variation in both parameters were correlated; normalization increased or decreased as Γ increased or decreased, leading to a softening of the spectra with brightness.

Notably, the third principal component still appears in simulated PCAs, even though we can be absolutely certain that the model can be described by only two components. This indicates that the third principal component does not correspond to any sort of model parameter, but rather is created as a by product of the PCA process. One of the assumptions made during PCA is *linearity*, meaning that the new basis vectors are a linear combination of the old ones, and that any correlations among the original data set are linear. However, this is not entirely true for most spectra, even those conforming to a simple power law model. A power law changing its slope, for example, can not be described linearly. A linear approximation of a power law will always undershoot the model at both low and high energies, and overshoot it in the middle, no matter the slope of the power law. The PCA process sees this as a problem, and fixes it by creating a new component with just the right shape to make up for the places that the linear approximation fails.

Since the second (pivoting) component is the one responsible for describing the changing shape of the power law, this third component should be strongest whenever the second component is strongest (in either direction). Strongest, in this case, refers to the normalizations of the principal components. If this explanation is true, a plot of the second component against the first should display a V-shape.

Figure 6.2 shows the normalizations of the second and third components plotted against each other for PCA performed on 1000 simulated power law spectra that varied as in the second simulation in Figure 6.1. The results show that the third component grows in strength as the second component increases in either direction, as expected. This is how the PCA code accounts for its inability to describe a changing power law using a linear function. It also explains why the third simulation in Figure 6.1 shows a suppressed third component: by varying Γ and normalization together, we have introduced linear correlation in the data, therefore reducing the need for a corrective third component. Even though this third component is not physical, it does appear in real data and thus can still be used as an indicator of a changing power law.

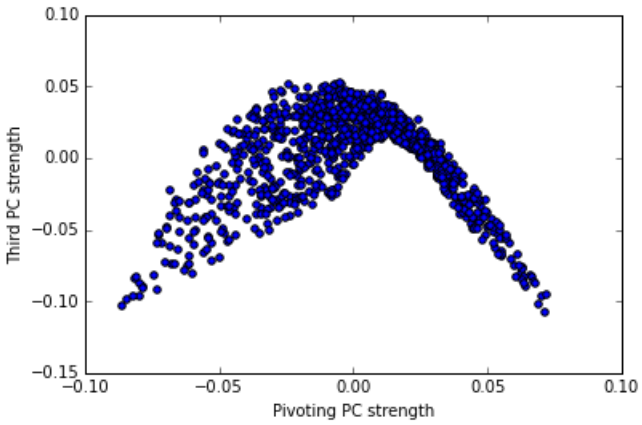

Figure 6.2: Normalizations of second and third principal components plotted against each other for PCA of 1000 simulated power law spectra. The third component is strongest wherever the second component is strong in either direction. In other words, wherever a linear approximation would differ significantly from a power law, the third component accounts for the difference.

Chapter 7

DISCUSSION

The PCA of blazar X-ray variability over years indicates the variation arising from changes in a power law component. The primary principal component for all sources in our sample indicated that changes in brightness (normalization) is the dominant factor, responsible for >84 per cent of the variability in each source. The secondary effect (component two) was the changes in power law shape Γ, which accounted for up to 15 per cent of the variability.

All of the sources in our sample showed a significant third component in the long-term that is not obviously associated with a physical spectral parameter. In Section 6 we demonstrate that this arching principal component is not a physical effect at all, but rather a mathematical artefact of the PCA process caused by a lack of linearity within the data. This mathematical factor becomes much less prominent if the changes in normalization and photon index are correlated.

Therefore, taken at face value, the long-term variability in our sample of blazars can be described by random variations of the power law brightness and photon index, or perhaps correlated variations between the parameters with some time delay.

Although the results for each object were similar, the differences between the various long-term PCAs can tell us a surprising amount. In particular, the slope of the first component is stronger in objects where Γ varies across a wider range. This slope is angled away from zero when Γ and normalization vary independently of each other, and towards zero when they vary together. This can be used to distinguish between emission mechanisms in blazars. Blazars are known to emit in the X-ray through either synchrotron emission, the inverse Compton effect, or some combination of both according to their luminosity (Donato, 2001), and correlations between Γ and flux are indicative of inverse Compton emission (Fatima, 2017). A decrease in the first principal component with energy can therefore be used as an indicator of Comptonization.

The degree to which a power law shape dominates our results becomes obvious when compared to similar analyses of radio-quiet objects, such as many of the objects in Parker et al (2015). Radio-quiet objects display much more complicated components that can include prominent features corresponding to emission lines, absorption edges, blackbodies, and so on. While this makes PCA a necessary identifying model components in radio-quiet objects, it also means that it is harder to identify what each principal component corresponds to. In particular, even with just the two parameters of a power law model, a third, non-physical component is required to complete the PCA. In more complex objects with competing spectral models, it may be harder to pick out the useful results from the mathematical artefacts caused by the non-linearity of the data set.

The short-term PCAs are more complex and interesting when compared to the long-term results. Long-term variability is simple the sum of many smaller variations, and yet the same does not always seem to be true of the PCAs. Most observations show only a single principal component, perhaps due to limited signal-to-noise.

Many observations show an upward-curving shape or a broken power law shape, neither of which are seen in any of the long-term PCAs. This would seem to suggest a variability mechanism that only manifests over short timescales, and is washed out by larger changes in the long term, but it is unclear what this sort of mechanism could be. Some observations showed no significant components at all, a sign of no rapid variability. This indicates that even highly variable objects such as blazars can show moments of steadiness, or display variability mechanisms that operate on scales greater than hours.

Chapter 8

Conclusions

PCA was used to analyse the X-ray spectra of nine blazars in order to identify variability trends across several timescales. Over long timescales, variability was found to be consistent with changes in a power law model, as should be expected in a blazar. In addition to principal components corresponding to change in normalization and Γ, a third component was seen in all objects. This component has no physical explanation, and instead was found to be a relic of the PCA process created by non-linearities within the data set. Even though each PCA shares the same broad power law shape, differences in the shapes of these components can be used to predict various qualities, such as the degree to which Γ is varying and correlations between spectral hardness and flux.

Over shorter timescales, the results were more complex. Some observations contained components similar to those seen in the long-term PCAs, which over time would add up to produce the long-term variability seen in each object as one would expect. However, others showed shapes not seen in the long-term analysis, including broken power laws and a unique, curved shape with no obvious physical analogue. Most of the short-term PCAs produced only one significant component, possibly due to low signal-to-noise. The smaller number of components and less consistent results

mean that it is harder to draw useful conclusions from single-observation PCAs at the moment, although there may be interesting physics to discover in this area if variability really does differ qualitatively in the short-term.

Principal component analysis is a useful tool that can offer a new approach with which to analyse a data set. However, even with objects as simple as blazars, it should not be trusted blindly without an understanding of its underlying assumptions and limitations.

Chapter 9

References

Abdo, A. A. et al., 2011, ApJ, 736, 2

Antonuccia, R., 1993, Annual Review of Astronomy and Astrophysics, 31, 473

Blumenthal, G. R. and Gould, R. J., 1970, Reviews of Modern Physics, 42, 237

Costamante, L. and Ghisellini, G., 2002, A&A 384, 56

Donato, D. et al., 2001, A&A, 375, 3, 739

Fatima, S. and Vierdayanti, K., 2017, AIP Conference Proceedings, 1801

Francis, P. J. and Wills, B. J., 1999, in Quasars and Cosmology, ed. G. Ferland and J. Baldwin, ASP Conf. Ser. 162, 363

Gallo, L. C. et al., 2015, MNRAS, 446, 633

Grupe, D., Beuermann, K., Mannheim, K., and Thomas, H.-C., 1999, A&A 350, 805

Grupe, D., 2004, AJ, 127, 1799

Jansen, F. et al., 2001, A&A, 365, L1

Kendall M. G., 1975, Multivariate analysis. Griffin, London

Koljonen K. I. I., McCollough, M.L., Hannikainen, D.C., and Droulans, R., 2013, MNRAS, 429, 1173

Malzac, J. et al., 2006, A&A, 448, 1125

Mastichiadis, A. and Kirk, JG., 1997, A&A, 320, 19

Miller, L., Turner, T.J., Reeves, J.N., George, I.M., Kraemer, S.B., and Wingert, B., 2007, A&A, 463, 131

Turner, T.J., Miller, L., Reeves, J.N., Kraemer, S.B., 2007, A&A, 475, 121

Parker, M. L. et al., 2015, MNRAS, 447, 72

Sambruna, Rita M. et al., 1997, ApJ, 483, 774

Shlens, Jonathon, *A Tutorial on Principal Component Analysis*, 2014

Struder, L. et al., 2001, A&A, 365, L18

Urry, C. Megan and Padovani, Paolo, 1995, Publ.Astron.Soc.Pac., 107, 803

Vaughan S. and Fabian A. C., 2004, MNRAS, 348, 1415

Appendix A

Specific Long-Term PCAs

This appendix presents the results of each long-term PCA and comments on them individually. As explained in Section 4, each follows the same pattern of a normalization component, a pivoting component, and an arch-shaped non-physical component elaborated on in Section 6. These components indicate a close fit to a power law model, which is expected for blazars. Still, there are smaller differences between the objects that deserve a closer look.

A.0.1 3C 273

These principal components are the simplest and easiest to explain of any object sampled. This PCA is an excellent match to simulations of a single power law varying in both normalization and photon index, corresponding to the first and second components respectively. The third component is not physical, but rather a mathematical artefact of the PCA process. This is discussed further in Section 6. A PCA of this object appears in Parker et al (2015), but displays a slight curve in the first component. This difference is due to the difference in methods: in this work, each observation contributed only one spectrum to the long-term PCAs, whereas Parker's earlier work splits each observation into smaller sections for every PCAs. This causes their results to look like a combination of our short-term and long-term results, explaining the

additional shape in the first component.

A.0.2 H1426+428

As with 3C 273, components corresponding to normalization and photon index can be seen clearly. Unlike 3C 273, H1426+428 shows some curvature in both major principal components, with the first increasing at high energies and the second and third suppressed at low energies. An increase in the first component can be reproduced in simulations of a power law model by increasing the degree to which Γ varies relative to normalization, as shown in Figure 6.1. Larger variations in Γ compared to normalization produce a steeper slope in the first component. A slant in the first component is seen in all of the objects, with 3C 273 being the flattest, indicating a relatively stable photon index.

The curved shape of the second component is harder to explain. A shape similar to this can be produced in simulations of a double power law model where two power laws vary together in normalization (Parker et al., 2015) but such a model does not reproduce the upward slope of the first component. Sambruna et al. (1997) find a variable warm absorber in the spectrum of this object (and of PKS 2155-304, which has a very similar PCA) but simulations of variable absorbers have also been unable to replicate the shapes of the first two principal components.

A.0.3 Mrk 421

This object's PCA is unique due to the kink found shortly before 2 keV. Mrk 421 was a highly piled-up source, and it is possible that this kink is caused by instrumental features. The silicon Kα or Kβ lines could be responsible.

A.0.4 Mrk 501

In Mrk 501, the first component again increases with energy, indicating large changes in photon index. The most notable feature is the shape of the second component, which begins rising back up around 7 keV rather than continuing downward as in the other pivoting components. The cause of this is unknown.

A.0.5 OJ 287

In OJ 287, the first principal component is decreasing with energy, rather than increasing. This can be reproduced in simulations by assuming that the variation in normalization and Γ is correlated, rather than varying them independently of one another, as shown in Figure 6.1. This can explain the shape of all three significant components in this object. Enforcing a correlation between normalization and Γ is not without precedent, and in fact seems to be the case for OJ 287 in particular. Fatima (2017) finds such a correlation, and remarks that this indicates inverse Compton emission rather than the synchrotron process.

A.0.6 PG 1553+113

PG 1553+113 shows the most similarity to Mrk 501, especially in the second component. However, the first component flattens out at high energies rather than continuing to increase, and the third component is much flatter, barely showing any shape at all.

A.0.7 PKS 2155-304

This PCA is nearly identical to that of H1426+428. Both objects are known to have warm absorbers (Sambruna 1997), but no absorption model has reproduced these principal components as of yet.

Appendix B

Specific Short-Term PCAs

Here, the results of each single-observation PCA are presented individually. They are not as similar as the long-term ones, falling into the categories discussed in Section 5.

B.0.1 3C 273

3C 273 had by far the highest number of suitable observations for this analysis, and the results are wide-ranging. Observations 012670301, 0136550101, and 0159960101 are shaped like a broken power law varying only in Γ. Observation 0414190101 looks like a change in normalization. Finally, observations 0126700801 and 0651610101 show a strange, upward-curving shape not seen in any of the long-term observations. The cause of this shape is unknown, and simulations have failed to replicate it, but it also appears in short-term PCAs of 3C 279, H1426+428, and OJ 287. This unique shape could indicate the existence of a process that drives short-term variability but has no effect in the long term, but it is hard to conceive of such a process, especially for objects with as simple a spectrum as blazars.

B.0.2 3C 279

There are not enough *XMM-Newton* observations of this object to perform a long-term analysis, but one observation long enough for short-term analysis does exist. The results are similar to those for 3C 273 observations 0126700801 and 0651610101, an upward-curving slope with no obvious explanation.

B.0.3 H1426+428

H1426+428 observation 0111850201 displays the upward curving shape seen before. There also appears to be some amount of shape to the second component, but it is of very low significance and is unlikely to be a real effect. A similar shape is seen in OJ 287 observation 0761500201.

B.0.4 Mrk 421

Unlike most of the other objects, both observations of Mrk 421 show more than one principal component. Observation 0099280201 has a normalization and a pivot component, as seen in most of the long-term observations. Observation 0136540701, on the other hand, looks like a broken power law. It is worth noting that the third (nonphysical) component seen in the long-term PCAs is not discernible in observation 0099280201. This is most likely because the second component is very weak.

B.0.5 OJ 287

The upward-sloping shape is seen again in the first principal component of observation 076500201, and, like in H1428+428, the second component seems to have a slight upward shape to it near the end as well. As with H1426+428, the second component is not significant at all, but the same semblance of a shape appearing twice is unlikely to be a coincidence. It may be an artefact of the PCA process, much like the third component seen in the long-term analyses.

B.0.6 PKS 2155-304

Observation 0124930301 indicates normalization changing alone, whereas observation 0124930601 seems to indicate changes in both normalization and Γ.

B.0.7 S50716+714

Like 3C 279, there were not enough observations to perform a long-term PCA on S50716+714. Observation 0502271401, the only one available for this analysis, shows some unexpected results. This is the only PCA where the normalization component is less significant than the pivoting one, and it is unclear why this would happen. Without information on the long-term variability of this object, it is hard to tell what separates it from the others.