1 What Are Active Galactic Nuclei? And Why Does Anyone Care?

1.1 What Makes Them Interesting?

Active galactic nuclei (hereafter abbreviated "AGNs") are among the most spectacular objects in the sky. They produce prodigious luminosities (in some cases apparently as much as 10^4 times the luminosity of a typical galaxy) in tiny volumes (probably $\ll 1$ pc^3). This radiation can emerge over an extraordinarily broad range of frequencies: in at least one case the luminosity per logarithmic frequency interval, that is, the luminosity per "band," is roughly constant (to within factors of several) across thirteen orders of magnitude in frequency! Their line spectra are almost as remarkable as their continua: in the optical and UV, they often display emission (and occasionally absorption) lines whose total flux is several percent to tens of percent of the continuum flux, and whose widths suggest velocities ranging up to $\sim 10^4$ km s^{-1}.

Although in most cases we cannot as yet make images in which AGNs are resolved, in certain objects it is possible to do so at radio frequencies. In those cases, one often sees variable structure with apparent speeds in the sky plane of $\sim 10c$!

Active galaxies are also noteworthy in displaying very strong cosmological evolution. The most luminous active galaxies were a thousand times more numerous at redshift 2.5 than they are today. Because their high luminosities and distinctive spectra make them relatively easy to pick out, they are disproportionately represented in our tally of known high redshift sources. The fact that the luminosity of AGNs is such a strong function of redshift suggests strongly that there is something special about youthful galaxies that promotes the creation of active nuclei.

Another interesting connection between AGNs and the evolution of the Universe comes from their use as light sources for studying intervening gas clouds, galaxies, and clusters of galaxies. Our knowledge of the intergalactic medium comes almost exclusively from the study of absorption lines in the spectra of distant quasars. On the other hand, systems that are transparent, but possess moderately deep gravitational wells, like galaxies and clusters of galaxies, can form gravitational lenses. The distorted—and sometimes multiple—images of distant AGNs formed by these lenses can be used to infer the nature of the gravitational potential—and therefore the mass—of the lenses.

To the degree that we do understand AGNs physically (a rather controversial point), their basic nature involves events of considerable physical drama. Most workers in the field now believe that the power for AGNs comes from accretion onto massive black holes. If this is true, their most basic properties depend on some of the most exotic physics we know: strong-field relativistic gravity.

This whole stew—exotic physics, photons detectable with virtually every sort of astronomical instrument, and a deep connection to cosmology—has made AGNs the focus of a significant fraction of the world astronomical community's attention for the past 30 years.

1.2 What Exactly Are We Talking About? The Most Salient Properties of AGNs

In order to start any discussion of AGNs, we must first define what we mean by the term *active galactic nuclei*. Unfortunately, although the black hole model has achieved widespread acceptance, it is not yet completely confirmed; moreover, and this is part of why its confirmation remains incomplete, direct signatures of accreting massive black holes are much harder to see than a variety of more indirect signals. Consequently, the only clear way to define AGN is operationally; all we can do is to list the observable phenomena we use to find them. This procedure is, admittedly, circular, but such a "bootstrap" approach is all that is possible until we achieve a more fundamental understanding.

It turns out that AGNs can be found in many ways, but not all AGNs have every single property. Thus, there is no single defining list of qualities to look for. A better way to think about the situation is to imagine a "menu" of phenomena from which they choose, with some items more popular than others. Note that in evaluating "popularity," one's terms of definition must be carefully chosen; in samples selected on the basis of a certain property, that property will always be very popular, whereas in samples selected otherwise it may be only a specialized taste.

Table 1.1 illustrates this point by listing some of the salient observational signatures of AGN, along with brief comments on how often these

properties are found, and a few concise caveats. Each of these features is hardly ever seen in normal galaxies. In the following subsections we will discuss these signatures at somewhat greater length.

1.2.1 Very small angular size

The first property is certainly the most visually striking. When the AGN is near enough that a host galaxy of reasonable surface brightness can be seen, in optical images the nucleus often appears to be a bright point whose flux can often rival, or even exceed, the flux from all the rest of the host galaxy (e.g., NGC 1566, a nearby AGN, is shown in Plate 1).

However, this simple picture is a bit misleading. Our ability to see both a bright point and its surrounding host is a function of the luminosity contrast between the nucleus and its host galaxy, a quantity that varies both from case to case and as a function of wavelength. If the luminosity of the nucleus is too small relative to the host, it will not stand out, of course. Conversely, when the luminosity of the nucleus is much greater than that of its host, light from the nucleus can overwhelm any light from the host.

This latter effect is exacerbated by the kinematics of our expanding Universe. For a start, at $z \simeq 1$ the angular size of an average galaxy is only $\sim 1''$, comparable to the seeing disk even at a good ground-based observatory. The problems don't end there. The typical luminosity of the AGNs we can find at $z \simeq 2$ is $\sim$ 100 times greater than at the present epoch, so the luminosity contrast between the nucleus and its host is very large. In addition, although the nucleus remains effectively a point, the bolometric surface brightness of the host falls as $(1 + z)^{-4}$ (see Appendix F). Finally, if the AGN has a redshift greater than about unity, the light we observe in the visible band was ultraviolet in the rest frame, and most galaxies are comparatively dim in the UV. All these effects combined make hosts of high-redshift AGNs very difficult to see (§13.1.1).

The picture we see also depends strongly on wavelength. As we will discuss at slightly greater length in the next section, many AGNs have a much greater ratio of X-ray luminosity to optical than does any normal galaxy. For this reason, their X-ray images are essentially pure points. On the other hand, radio emission, more often than not, extends over a sizable region, frequently much larger than a galaxy (§§ 1.2.7, 9.1.2).

In the long run we can hope to obtain images of even the smallest AGN structures, but the requirements are daunting. As we shall see later, there are AGN pieces spread over a very wide dynamic range of radial scale.

Table 1.1: The Menu

Property	Popularity	Comments and Exceptions
Very small angular size	Many	Wavelength-dependent
Galactic (or greater) luminosity	Many	Lower luminosity is hard to find; obscuration and beaming may mislead
Broad-band continuum	Most	Often $dL/d\log\nu \simeq const.$ from IR to X-rays; sometimes to γ-rays
Strong emission lines	Most	Sometimes very broad, sometimes not
Variable	Most	Modest amplitude; short wavelengths stronger, faster than long
Weakly polarized	Most	$\sim 1\%$ linear; a minority much stronger
Radio emission	Minority	Sometimes, but not always, extended on enormous scales
Strongly variable and polarized	Small minority	Correlated with bright radio and high-energy γ-rays; in some cases emission lines absent

Equivalent elements in different AGNs tend to have roughly the same effective temperature (i.e., luminosity emitted per unit area $F = \sigma T_{\rm eff}^4$, where σ is the Stefan-Boltzmann constant) because the effective temperature roughly characterizes the nature of the structural element. The effective temperature is also a rough guide to the typical wavelength emitted by that part of the AGN ($\lambda \sim ch/kT_{\rm eff}$) if the radiation mechanism is roughly thermal, but nonthermal mechanisms can easily make the typical wavelength far shorter.

Even at the heart of the beast, $T_{\rm eff}$ is rarely much more than $\sim 10^5$ K (§§7.1, 7.3.3); discounting radio emission regions, at the outer edge of an AGN proper, it may fall to a few hundred K (§12.4.4). We can then roughly predict that in an AGN whose observed flux is $F_{\rm obs}$, the angular size of the region whose effective temperature is $T_{\rm eff}$ is

$$\theta = 87(1+z)^2 \left(\frac{F_{\rm obs}}{10^{-11}\ \mathrm{erg\ cm^{-2}\ s^{-1}}}\right)^{1/2} \left(\frac{T_{\rm eff}}{10^3\ \mathrm{K}}\right)^{-2} \quad \text{microarcsecond}, \tag{1.1}$$

where the fiducial value of $F_{\rm obs}$ has been chosen at a level corresponding to roughly the tenth brightest AGN in the sky. Thus, the angular size of any particular AGN structure scales $\propto F_{\rm obs}^{1/2}$, so that the apparently brightest AGNs are also the largest (in angular terms)—and even they are very small indeed.

1.2.2 High luminosity

We know of AGNs with luminosities all the way from $\sim 10^{42}$ to $\sim 10^{48}$ erg s^{-1}. To put this in perspective, the characteristic luminosity of the field galaxy distribution (L_*) is $\sim 10^{44}$ erg s^{-1}. In other words, we see AGNs whose power output ranges from as little as 1% of a typical galaxy to $\sim 10^4$ times as great. However, one must be careful in interpreting these luminosities. On the one hand, we cannot easily detect active nuclei much weaker than the host galaxy, and so there may be a large population of "mini-AGNs" that are as yet unknown. On the other hand, there is also reason to think (see Chap. 12) that in many AGNs the active nucleus is obscured by extremely thick dust extinction, so that we can be grossly misled as to its true luminosity if our only measure is the power output in optical or ultraviolet light. Relativistic beaming (§§9.3.4, 12.2.1) can also substantially distort the angular distribution of light from AGNs, and, of course, there is a strong selection effect in favor of observing those whose radiating material is moving toward us. Thus, in objects where beaming is likely to be significant, it is important to distinguish between the luminosity inferred assuming isotropic radiation and the true luminosity. Obscuration, or beaming directed away from us, can, of course, so weaken the light we see from an AGN that we may not even recognize it.

1.2.3 Broad-band continuum emission

It is best to begin this discussion with a digression about an issue of notation. In most fields of astronomy, the property of light that is measured is the specific flux (F_ν or F_λ), the rate at which energy arrives per unit area per unit frequency ν or per unit wavelength λ. The standard unit for F_ν is the Jansky, 10^{-23} erg cm^{-2} s^{-1} Hz^{-1}. In high-energy (i.e., X-ray and γ-ray) astronomy, where photon-counting devices prevail, the customary measured quantity is N_ϵ $(= F_\nu/(h\epsilon))$, the rate at which photons arrive per unit area per unit energy ϵ. However, when we speak of a "band" of the electromagnetic spectrum, whether it is radio, infrared, visible, or X-ray, we generally mean a span in the logarithm of the wavelength. The term *infrared*, for example, generally refers to a range of $\sim 10^2$ in wavelength, from ~ 1 to ~ 100 μ. Therefore, for describing which band is most important in terms of energetics, the most convenient quantity is $\nu F_\nu = dF/d\log\nu = dF/d\log\lambda = \lambda F_\lambda$, the energy flux per logarithmic bandwidth. In most of this book (starting with fig. 1.1), this will be the favored form for the presentation of spectra.

To understand what is meant by "broad-band" continuum radiation, one should first contrast the spectra of ordinary galaxies (see fig. 1.1). To a first approximation, galaxies are piles of stars. A good zeroth-order approximation to a stellar spectrum is that it is a blackbody, so the great majority of a star's luminosity comes out within a factor of three in frequency. The total span of stellar surface temperatures is only about a factor of ten, and in any particular galaxy, the stellar mix is usually such that a limited temperature range dominates the total power output. Thus, a typical galaxy emits nearly all its power within no more than one decade of frequency, and usually rather less. The only possible modification to this picture is due to interstellar dust. In many spiral galaxies the dust extinction is great enough that a significant fraction of the optical and ultraviolet light is absorbed by cool dust grains and reradiated in the far-infrared. Because there is a wide range in extinctions, the relative size of this secondary peak in the infrared varies substantially from one spiral galaxy to the next. In ellipticals, on the other hand (as shown in fig. 1.1), there is generally little dust, and hence at most weak infrared emission.

Most (but not all) AGN continuum spectra look spectacularly different from normal galaxy spectra. A particularly well observed example,

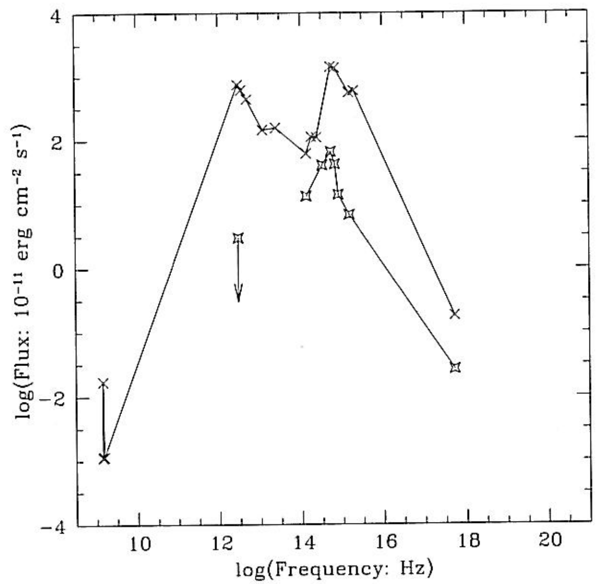

Fig. 1.1 The broad-band spectra of two typical galaxies, M 101 (Hubble type Sc: plot symbol $\times$) and NGC 4168 (Hubble type E2: plot symbol a four-pointed star). The star with an arrow is an upper limit. Connecting lines are meant only to guide the eye. In both cases, the intrinsic radiation is confined to a narrow range of optical frequencies, with only minor amounts of power emerging in the radio or X-ray bands. In the case of M 101, however, almost half the intrinsic power is absorbed by dust and reemitted at ~ 100 μ. The spike at 1.4 GHz in the spectrum of M 101 is the 21 cm HI line. At all frequencies but the near-infrared, the fluxes are integrated over the entire galaxy from maps; in that band the apertures used cover most, but not quite all, of the galaxy and therefore give fluxes that are slightly too low. Data sources are: White and Becker 1992, Rice et al. 1988, Fabbiano et al. 1992, and NED.

NGC 4151, is shown in figure 1.2. In terms of νF_ν, NGC 4151, like most AGNs, has a spectrum that is flat (to within factors of several) all the way from the mid-infrared to the hardest X-rays observed (the highest observed energy is anywhere from a few to a few hundred keV, depending on the object). Compared to normal galaxies, the fraction of the bolometric luminosity radiated in the radio band is generally an order of magnitude greater, but in some cases it is several orders of magnitude larger still; the fraction of the power that emerges in X-rays is three to four orders of magnitude larger in AGNs than in normal galaxies. NGC 4151 is so well observed because it is very nearby ($z = 0.003$), but its luminosity is relatively low,

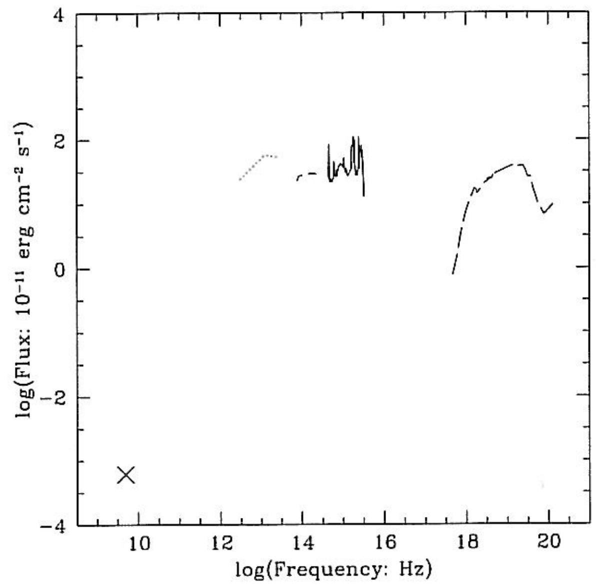

Fig. 1.2 The broad-band continuum spectrum of one of the nearest AGNs, NGC 4151, compiled from nonsimultaneous (!) data. Its flux per logarithmic bandwidth has comparable peaks in the infrared, the optical/ultraviolet, and the hard X-ray bands. Far less luminosity is radiated in radio frequencies. The dip in the soft X-ray region is due to intervening absorption, possibly in the AGN itself; the bumps in the optical/ultraviolet are smoothed versions of the strong emission lines seen in this object. The data are taken from: Ulvestad, Wilson, and Sramek (1981), the ×; Edelson and Malkan (1986), the dotted line; Kriss et al. (1995), the broken line in the optical band; the *HST* archive, the solid line; and Zdziarski, private communication, the dashed line in the X-ray band. In all cases but the four infrared points at 12 μ, 25 μ, 60 μ, and 100 μ, the apertures used were small enough that the AGN dominates the flux.

merely comparable to its host galaxy. Composite spectra illustrating what AGNs look like at high redshift and high luminosity are shown in figure 1.3.

All told, the range over which νF_ν is roughly flat is more than a factor of 10^5 in frequency, producing a spectrum far broader than any normal galaxy's. Although there are weak local maxima, it is clearly inappropriate to speak of any one frequency band dominating the output. These local maxima may, however, be signaling to us that the primary emission mechanisms change as functions of frequency (Chaps. 7, 8, 9, 12). In the case of one particular subclass of AGN, the flux in photons as hard as 1 GeV is at least as great as that in lower frequency bands, and there are a few

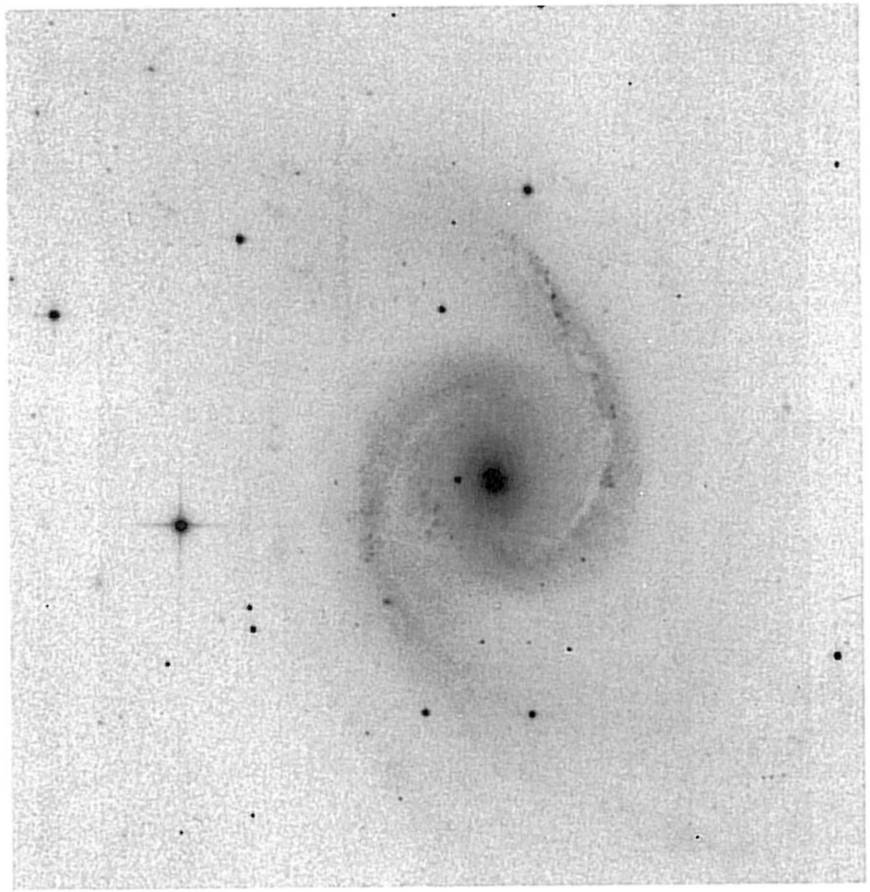

Plate 1 NGC 1566, a nearby active galaxy. It appears to be a fairly normal spiral galaxy, but for the extremely bright nucleus (image courtesy of Z. Tsvetanov).

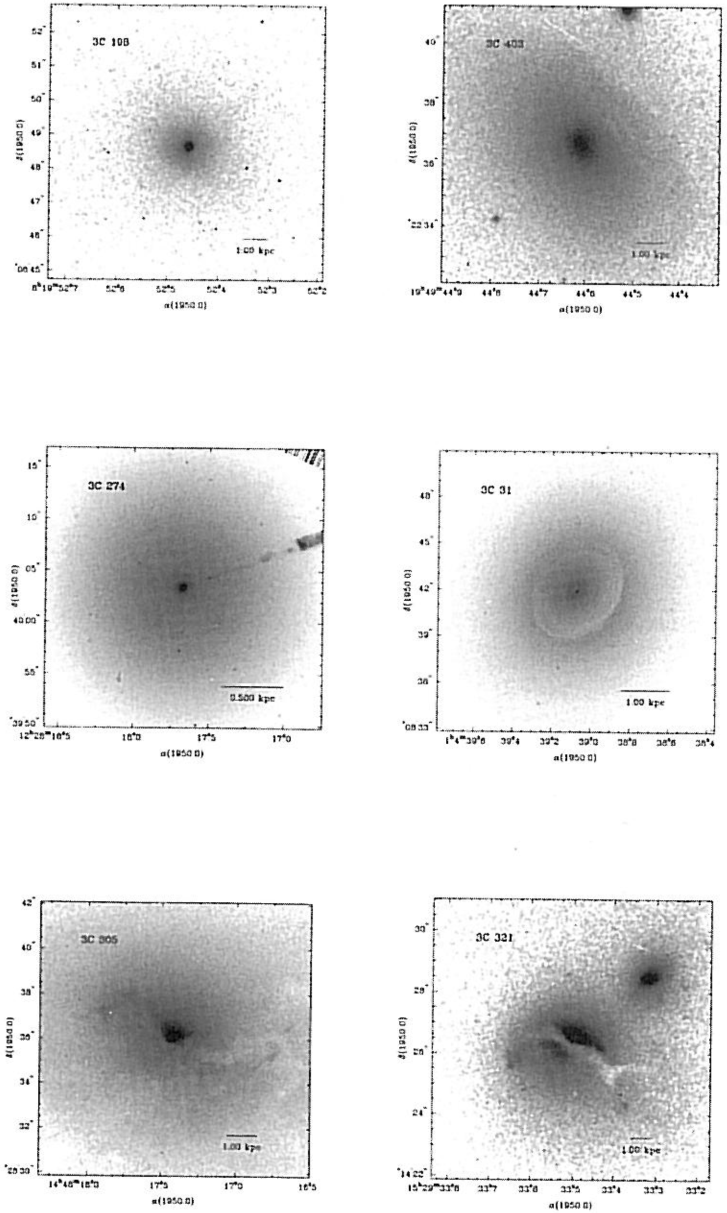

Plate 10 A sampler of R-band images of the hosts of low-redshift 3C radio galaxies, taken from Martel et al. (1997). From upper-left to lower-right, they are 3C 198 (z = 0.082), 3C 403 (z = 0.059), 3C 274 (z = 0.004), 3C 31 (z = 0.017), 3C 305 (z = 0.041), and 3C 321 (z = 0.096). The distance scale given for each is computed assuming $H_o = 75$ km s^{-1} Mpc^{-1}. Most of these are obviously elliptical galaxies, but several just as unmistakably show unusual features. 3C 274 is also known as M 87; its jet (apparent all the way from radio frequencies through X-rays) is clearly visible. 3C 305 has been disturbed in some fashion. 3C 321, for which a near-UV and polarimetric image is shown in Plate 8, has a dark dust lane across its center, and a nearby companion.

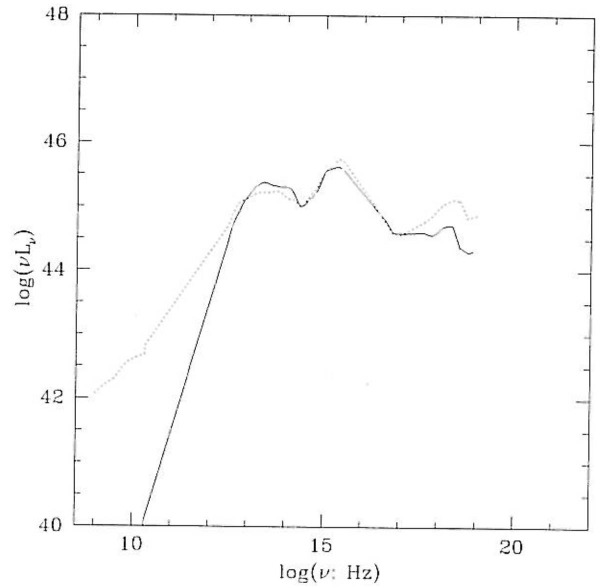

Fig. 1.3 Composite spectra for two different samples of high-redshift, high-luminosity AGNs, one (solid curve) relatively weak in the radio band, and the other relatively strong (dotted curve). The straight line segments in the far-IR/sub-mm band and the EUV are both interpolations across spectral regions where no real observations exist. Although the curves end at photon energies of a few tens of keV, this is not because AGNs are weak there (see, e.g., figs. 1.2 and 1.4).; it is because good data in the hard X-ray band and beyond exist for comparatively few objects. As these curves are composites, the luminosity scale (given in units of erg s^{-1}) should not be taken literally. Data are from Elvis et al. (1994).

examples of this subclass (e.g., Markarian 421; see fig. 1.4) with as much flux at 1 TeV as anywhere else in the electromagnetic spectrum. However, as always, there are exceptions. There are also AGNs in which the bulk of the light arrives at Earth within a frequency span of only one decade in the infrared (NGC 1068, another very nearby AGN, is a good example of this variety; see fig. 1.5).

1.2.4 Emission lines

AGN emission lines have received a great deal of attention for two reasons. First, they are often very prominent (equivalent widths are often ~ 100 Å). This makes AGN spectra stand in great contrast to the spectra of most stars and galaxies, where lines are generally relatively weak and

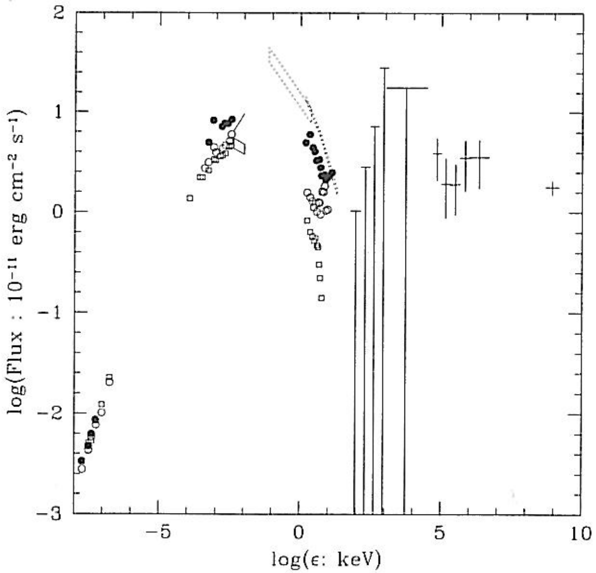

Fig. 1.4 Markarian 421 is a BL Lac object (see table 1.2). Its luminosity is spread almost equally per logarithmic frequency interval over 13 decades in frequency! Where different levels are shown at the same energy, the contrast denotes variability. The data for this plot were collected by Zdziarski and Krolik (1993).

predominantly in absorption. Figure 1.6 presents a composite spectrum that dramatically illustrates just how prominent the emission lines are. Second, because we know a great deal about atomic physics, it is easy and productive to study them (Chap. 10).

The emission lines that we see are remarkably stereotyped from one object to the next. When we have the wavelength coverage to look for them, if there are any lines at all, we almost always see Lyα, the Balmer lines, the CIV 1549 doublet, [OIII] 5007, and several others that are generally weaker. The Fe Kα X-ray line near 6.4 keV is also frequently seen.

However, there is an interesting split in the line width distribution. In some objects, many of the lines have broad wings extending out several thousand km s^{-1} from line center, whereas in others the lines are never broader than a few hundred km s^{-1}. Interestingly, the permitted and semi-forbidden lines are seen in both of these classes; in fact, when the broad wings appear, there is often a narrow core as well. The forbidden lines, on the other hand, are only seen with narrow profiles. In another

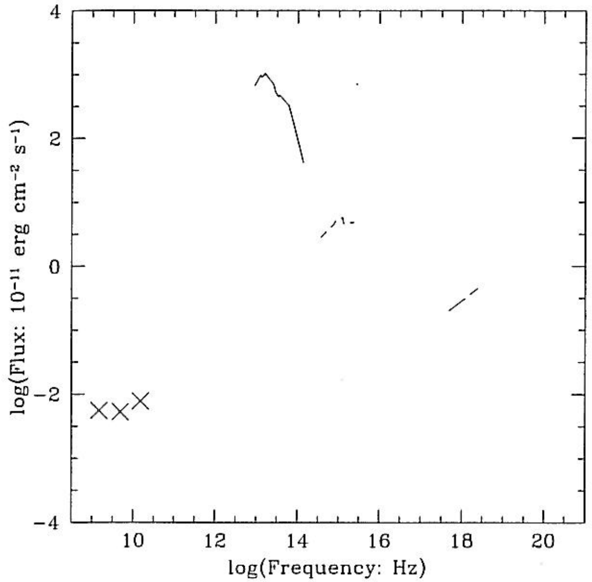

Fig. 1.5 NGC 1068 is another very nearby AGN. Unlike NGC 4151, its spectrum has a strong peak in the infrared, most likely due to very optically thick dust obscuration that reradiates the nucleus's luminosity in that band (§12.4.4). The data are from: Wilson and Ulvestad (1982), the ×'s; Rieke and Low (1975), the solid line; and Pier et al. (1994), the dashed lines. They have been analyzed in such a way as to show only nuclear radiation.

interesting correlation, those objects with only narrow lines are often quite weak from the near-infrared through the X-ray band; most of their light is generally emitted in the mid-infrared (as in the example shown in fig. 1.5).

1.2.5 Variability

It is often loosely remarked that variability is a hallmark of AGNs. This is only partially true. In the optical band, most AGNs, unlike normal galaxies, can be seen to vary, but the typical amplitude over human timescales (e.g., a few years) is often only 10% or so. Incomplete evidence suggests that the variability amplitude on the most easily observed timescales increases toward shorter wavelengths, with factors of two often seen in the X-rays. Figure 1.7 shows two cases in point, once again the frequently observed AGN NGC 4151, and the almost equally popular NGC 5548.

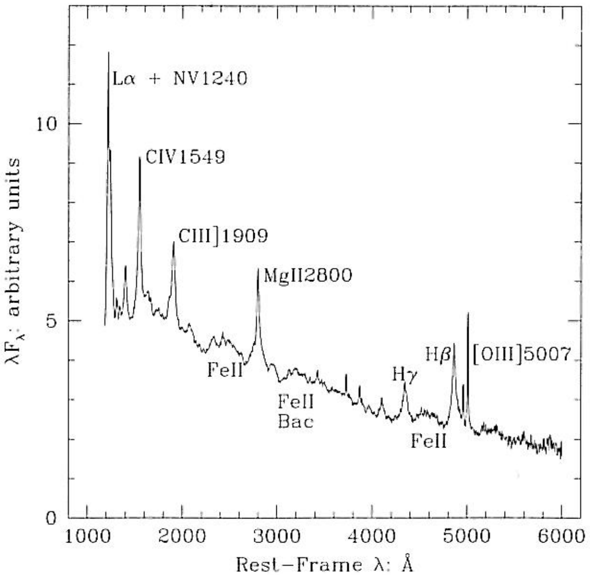

Fig. 1.6 A composite optical/ultraviolet AGN spectrum as compiled by Francis et al. (1991) from a large quasar survey. The very large contrast between emission lines and continuum, and the lines' substantial breadth, are both immediately apparent. Some of the more prominent emission lines are labeled; "FeII" is a shorthand for blends of many FeII multiplets, and "Bac" refers to the Balmer continuum.

It cannot be emphasized too strongly (fig. 1.7 underlines this point) that any statement about variability is strongly dependent on the timescale in question. A well known theorem in Fourier analysis states that the variance of a function of time $F(t)$ sampled at intervals Δt over a duration $t_{\rm tot}$ is equal to the integral of the power spectrum over the range of frequencies to which such measurements are sensitive:

$$\mathrm{Var}(F) = \int_0^{t_{\rm tot}} dt\,[F(t) - \langle F\rangle]^2 = \int_{1/t_{\rm tot}}^{1/(2\Delta t)} df\,|\hat{F}(f)|^2. \qquad (1.2)$$

The form of this relation makes it very clear that the measured variance depends on the sampling unless it is good enough that essentially all the variability power is contained within the observed frequency range. Unlike stars, whose variability is often dominated by periodic components (consider eclipsing binaries or Cepheid variables), AGNs for the most part vary with no special timescales; that is, their Fourier spectra are broad-band, just as their photon spectra are. In consequence, the amplitude of variability for AGNs is a slippery thing to measure.

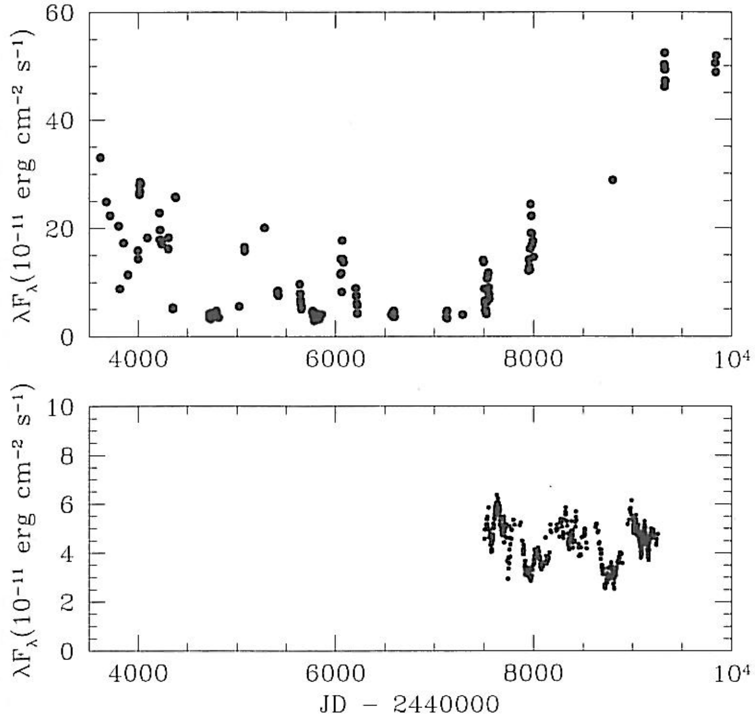

Fig. 1.7 (Top) A long-term UV (1455 Å) lightcurve for NGC 4151 (from Ulrich et al. 1997); (bottom) a somewhat shorter optical (5100 Å) lightcurve for NGC 5548 (data from Peterson et al. 1991, 1992, and 1994). In the UV, fluctuations of factors of several are common and can occur on timescales ranging from weeks to years. In the optical band, the fluctuations tend to be rather smaller.

A small subset of AGNs vary much more strongly, even in the optical band. In some cases, fluctuations of a factor of two have been seen from night to night, and cumulative changes of factors of 100 have occurred over year timescales (see, e.g., fig. 1.8). Strong variability is also very strongly correlated with three other properties: strong polarization, compact radio structure, and strong high-energy γ-ray emission.

1.2.6 Polarization

Most stars are intrinsically unpolarized, but the light we receive is generally linearly polarized by $\sim 0.5\%$ due to interstellar dust transmission polarization. The same is true for most galaxies. Most AGNs are also weakly polarized, but just enough more strongly for their polarization distribution to be statistically distinguishable from that of stars: typically they are lin-

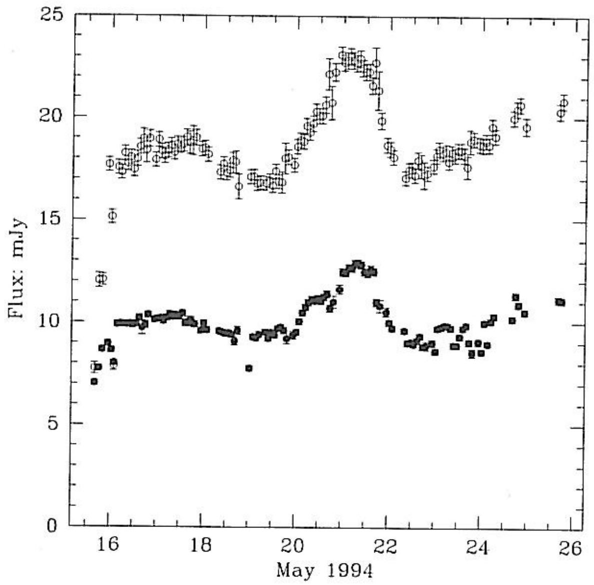

Fig. 1.8 Two ultraviolet lightcurves for PKS 2155-304. The open circles represent continuum flux at 2800 Å; the filled circles show the continuum flux at 1400 Å. In this case, changes of several tens of percent occurred within 1 day. Data courtesy of E. Pian, originally published in Pian et al. (1997).

early polarized, with fractional polarization $\simeq 0.5$–2%. However, again as for variability, a minority—including the same minority that have strong optical variability, as well as some of the AGNs with only narrow emission lines—are much more strongly polarized, often $\sim 10\%$ in linear polarization. Those objects that are strongly polarized and strongly variable in total flux are also highly variable in both the magnitude and the direction of the polarization. This variability has its bounds, however; circular polarization has never been detected, and the limits are reasonably tight.

Technical caveats are in order for discussing polarization as well. When the luminosity of the nucleus does not completely dominate the luminosity of the host galaxy, the polarization we measure depends on the aperture used, for starlight can substantially dilute the polarization of the nucleus. This effect is most noticeable in the AGNs with only narrow emission lines. In those cases, the (linear) polarization of the true nucleus can be as high as tens of percent but can only be seen after virtually all starlight has been eliminated (Chap. 12).

In addition, whether or not we detect polarization depends on the wavelength observed. Radio emission is usually linearly polarized at the few to few tens of percent level, but the observed value can be artificially suppressed by insufficient angular resolution, for the angle of polarization often varies from place to place (§9.1.4). Some AGNs whose spectra show exceedingly broad ultraviolet absorption troughs (see Chap. 11) are strongly polarized in the absorption features and more weakly polarized in the continuum. Others become strongly linearly polarized at wavelengths just shortward of the Lyman edge. Whether the X-ray emission is polarized is a question that awaits more sensitive instruments than those available in the 1990s.

1.2.7 Radio emission

Strong radio emission is the last of the distinguishing marks of AGNs. Historically, however, it was effectively the first. Some of the earliest radio astronomical observations discovered that many bright radio sources come in the form of double lobes with a galaxy located halfway between them (e.g., Cygnus A, shown in Plate 2). These radio lobes were the first evidence recognized as indicating nonstellar activity in external galaxies. After the compilation of the 3C catalog (see Chap. 2) in the late 1950s, efforts to identify its members led to Maarten Schmidt's realization in 1963 that the bright optical point source associated with one of these radio sources, 3C 273, possessed the (then) shockingly large redshift of 0.158. This realization cracked open the field of active galaxies; suddenly the unidentifiable optical point sources associated with numerous radio sources became objects with interpretable optical spectra, and known distances and luminosities.

Because radio astronomical techniques are extremely powerful, and a large fraction of all bright radio sources are AGNs, many of the *known* AGNs are strong radio emitters, and a great deal is known about the phenomenology of that emission. For example, only in the radio band is milliarcsecond imaging (by Very Long Baseline Interferometry, or VLBI) currently available, and AGN radio emission is mostly resolved on this angular scale.

However, this depth of knowledge is a bit misleading. Even in AGNs where the radio band is relatively strong, it never accounts for more than $\sim 1\%$ of the bolometric luminosity; in addition, less biased surveys have shown that the great majority of AGNs emit a much smaller fraction of their total power in the radio.

1.3 AGN Nomenclature

Because subclasses of AGNs exist that all share the same choices from this menu, numerous subvarieties have been named. Whether the taxonomical effort that created this "zoo" has helped clarify our view of this subject, or led to further confusion, is a matter of some debate. Nonetheless, to understand the conversation, one must learn the language.

Table 1.2 presents the lineaments of AGN zoology. The names themselves reveal how roundabout scientific progress can be. Some are descriptive: *radio-loud* and *radio-quiet* are fairly self-explanatory terms, and, as we shall discuss in Chapter 9, the distribution of the fraction of the bolometric luminosity that is radiated in the radio band is quite bimodal. *OVV* is an acronym for "Optically Violently Variable," a term that is equally direct: this class is marked by exceptionally rapid and large amplitude variability in the optical band. Other designations refer to the names of the first people to identify the class: Carl Seyfert pointed out the first six *Seyfert galaxies* (thirty years later they were subdivided into two principal types, according to whether their emission lines did or did not have broad wings: see §10.1.3); the "FR" in *FR1* and *FR2* stands for Fanaroff and Riley, who pointed out an interesting distinction in both luminosity and morphology among the radio galaxies; this will be discussed in detail in Chapter 9. Some class titles are deliberate coinings. *Quasar* was originally the pronounced form of "QSRS," an acronym for "quasi-stellar radio source," but over the years its usage evolved so that it now denotes nothing about the radio luminosity of an object. Today it is often used as a synonym for "generic AGN"; very low luminosity AGNs are sometimes called "micro-quasars," for example. And some have truly quirky histories. Variable stars are given names having two letters and an abbreviated form for the constellation in which they are found. The prototype of the AGN variety now called *BL Lac objects* was originally thought to be a variable star in the constellation Lacerta and was therefore given the name "BL Lac." Quirkier still is the origin of the term *blazar* (not found in the table because, 20 years after its introduction, its use is still somewhat nonstandard). This coinage is meant to unite the OVV and BL Lac classes and was invented as a joke by Ed Spiegel, the after-dinner speaker at the first conference organized on BL Lac objects. The name stuck both because the two varieties do resemble each other (compare their two lines in the table, deliberately put immediately adjacent to each other) and because its connotation of "blazing" is very appropriate to these objects whose power output varies so dramatically.

Other classes have been defined in the past, but are now rarely used. *N galaxies* are elliptical galaxies with bright optical nuclei, generally first

noticed because of their large radio power. Most would now be called *broad line radio galaxies*. "QSO," a term still occasionally used, is an acronym for "quasi-stellar object" that has now been largely replaced by *radio-quiet quasar*. Osterbrock and his associates defined several Seyfert types intermediate between 1 and 2 (see §10.1.3 for a more detailed discussion); these classes can still sometimes be seen in the literature.

The column headings are abbreviations for items in the menu: "Point-like" refers to whether an optical point source can be seen. "Broad-band" means that there is comparable luminosity in the infrared, optical, and X-ray bands. "Broad lines" and "narrow lines" indicate the existence in the optical and ultraviolet spectra of lines several thousand km s^{-1}, or several hundred km s^{-1} in width, respectively. "Radio" means that the fraction of the luminosity emitted in the radio is relatively large, perhaps $\sim 10^{-3}$ of the bolometric luminosity (defining radio as $\nu \sim 10$ GHz). To be considered significantly variable, the members of the class should vary by an order of magnitude or more in the optical band over a timescale comparable to a human lifespan. To receive a "yes" in the "Polarization" column the optical light should be at least a few percent linearly polarized.

The astute reader will notice that these classifications divide up into groups: radio-loud versus radio-quiet, strongly variable versus all others, and narrow emission lines only versus broad emission lines as well. These groupings suggest an arrangement in a three-dimensional parameter space, illustrated in figure 1.9. The existence of these groupings has led to a great deal of effort to explain these multiple categories in terms of a smaller number of variables. These efforts are discussed in detail in Chapter 12. For now, it is best simply to take note of the connections.

This same hypothetical reader will notice that type 1 Seyfert galaxies have identical table entries to those of radio-quiet quasars. He or she might then ask why they are listed separately. That is a good question. In fact, there are a number of objects that have been classified "type 1 Seyfert" by some observers, and "radio-quiet quasar" by others. In practice, the only distinction is whether a host galaxy is visible. When it is, the AGN is called a Seyfert galaxy, whereas when none is visible, it is called a quasar. Since the distinction depends on the resolution and background level of the instrument, and does not depend primarily on the object itself, it is clearly not a very useful one. Its only objective correlate is with luminosity: Seyfert galaxies are on average two orders of magnitude less powerful than quasars. It is for this reason, of course, that the host galaxy is visible, for we can only see host galaxies when the luminosity from the AGN does not overwhelm the starlight.

Table 1.2: The AGN Bestiary

Beast	Pointlike	Broad-band	Broad Lines	Narrow Lines	Radio	Variable	Polarized
Radio-loud quasars	Yes	Yes	Yes	Yes	Yes	Some	Some
Radio-quiet quasars	Yes	Yes	Yes	Yes	Weak	Weak	Weak
Broad line radio galaxies (FR2 only)	Yes	Yes	Yes	Yes	Yes	Weak	Weak
Narrow line radio galaxies (FR1 and FR2)	No	No	No	Yes	Yes	No	No
OVV quasars	Yes	Yes	Yes	Yes	Yes	Yes	Yes
BL Lac objects	Yes	Yes	No	No	Yes	Yes	Yes
Seyferts type 1	Yes	Yes	Yes	Yes	Weak	Some	Weak
Seyferts type 2	No	Yes	No	Yes	Weak	No	Some
LINERs	No	No	No	Yes	No	No	No

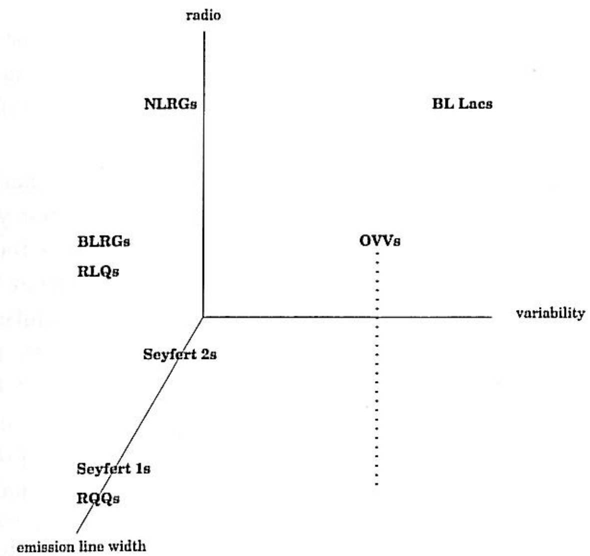

Fig. 1.9 The principal subvarieties of AGNs schematically arranged according to relative power in the radio band, emission line width, and variability. All combinations are possible except that there are no highly variable radio-quiet objects.

The last entry in the table, *LINERs*, denotes a category at the very margin of activity. As of this writing, it is still debatable whether these galaxies are truly active or not, or whether it is even a well defined category in any physical sense. Dubbed "LINER" for "Low-Ionization Nuclear Emission Region" by Tim Heckman in the early 1980s, these galaxies show strong emission lines, like Seyfert galaxies, but the relative strengths of those lines from low-ionization stages are rather greater than in Seyfert galaxies (see §10.1.3 for a quantitative definition).

It is also important to recognize that these categories may not exhaust the true list. It is entirely possible that there are varieties of AGNs that we do not yet recognize. Some people believe, for example, that the ultraluminous infrared galaxies discovered by the *Infrared Astronomical Satellite* (*IRAS*) may house AGNs (§13.6). Similarly, others suspect that there may be a class of *type 2 quasars*, by analogy with the division of the Seyfert class (§12.6.5). At present, limits on the existence of such objects are very weak. Surprises are always very much a possibility in the field of AGNs.