PUBLICATIONS OF THE ASTRONOMICAL SOCIETY OF THE PACIFIC
102: 932–948, August 1990

ON THE GROWTH-CURVE METHOD FOR CALIBRATING STELLAR PHOTOMETRY WITH CCDs

PETER B. STETSON*

Dominion Astrophysical Observatory, Herzberg Institute of Astrophysics, National Research Council of Canada
5071 West Saanich Road, Victoria, BC V8X 4M6, Canada

Received 1990 January 22, revised 1990 April 30

ABSTRACT

Stellar photometry with CCDs usually relies on the construction of aperture growth curves for its ultimate calibration. In the past this has been a tedious chore requiring a great deal of human intervention, mostly to select data suitable for defining empirical growth curves from data sets containing some corrupt values. A computer program has been written which incorporates a priori knowledge of the typical morphology of stellar profiles and is capable of taking a synoptic overview of all the aperture growth curves from an entire night or observing run. The program is thus enabled to make its own judgments as to the reliability of individual data points and to draw physically reasonable growth curves without human supervision, even for individual frames with insufficient or badly contaminated data. The program runs quickly and independently and produces results not noticeably inferior to those obtained by traditional hand-and-eye methods.

Key words: data-handling techniques–photometry: general

1. Introduction

Charge-coupled devices (CCDs) have revolutionized the field of star-cluster photometry. Their high quantum efficiency and generally excellent linearity and stability represent major advances over older photographic and electron-tube technology. Furthermore, their two-dimensional format offers more than the simple multiplex advantage that comes from their ability to record many stars in a single integration: Profile-fitting techniques permit precise photometry to be obtained in fields too crowded for work with the older technologies. Crowded-field stellar photometry is now routinely carried out with sophisticated, heavily automated computer software packages such as ROMAFOT (Buonanno *et al.* 1983; Buonanno and Iannicola 1989), STARMAN (Penny and Dickens 1986; see also Penny, *STARMAN: A User's Guide*), WOLF (Lupton and Gunn 1986), DAOPHOT (Stetson 1987; see also Stetson, *DAOPHOT User's Manual*), and DoPHOT (Mateo and Schechter 1989).

The profile-fitting techniques used by these computer programs depend on intensity scaling to define the relative magnitudes of stars contained within a given digital image. First, a model two-dimensional brightness profile resembling star images in the frame is somehow concocted and stored in the computer. Then for each star an intensity scaling factor is computed such that, when the model profile is multiplied by this factor, the intensity values actually observed in the detector pixels containing the image of the star are best[1] reproduced. In practice, this scaling factor is determined primarily by the innermost few pixels of the image of a star where the signal-to-noise ratio of the observed intensities is highest, the fractional contribution of the diffuse sky brightness is lowest, and the chance of contamination by other astronomical objects is least. Usually, "best" estimates of the centroid position of the star and of the local sky brightness are determined at the same time.

The magnitudes returned by such profile fits are relative—they represent the brightness ratios (⇒ the magnitude differences) between stars recorded in a given frame and the model stellar profile for that frame. To relate these differential magnitudes to an *external* photometric system—i.e., as defined by some set of standard stars—one must establish a fundamental magnitude zero point for the frame, which will be determined by the total number of photons actually detected from a given star per unit time. To the extent that the model profile may differ from one exposure to another because of changes in the seeing, telescope tracking errors, or defocusing, this zero point cannot be determined by a simple relative scaling of the model profiles of the different frames. Thus, except in the unusual and fortunate circumstance of having previously established photometric standards in each of the

*Visiting Astronomer, Cerro Tololo Inter-American Observatory, National Optical Astronomy Observatories, operated by the Association of Universities for Research in Astronomy, Inc., under contract with the National Science Foundation.

[1]Usually in a least-squares sense.

program fields, it is necessary to determine the absolute number of photons contained within the model profile of each frame. Only then may the inherent stability of the detector's sensitivity be used to define the magnitude zero point of one frame with respect to another. From this point on, the reduction of the program observations to a standard system—viz., by direct observation of standard stars and derivation of transformation and extinction corrections—follows the methods traditionally used in photomultiplier photometry (e.g., Hardie 1962; Harris, FitzGerald, and Reed 1981; see also Stetson and Harris 1988, §IV).

Determining the number of photons contained within the model stellar profile of a given frame is not necessarily straightforward. Some computer programs (e.g., ROMAFOT, STARMAN, DoPHOT) encode the model star image as an analytic function, with parameters which can be adjusted by fits to actual observed star images so as to optimize the simulation. Once these parameters have been determined, the analytic function could—in principle—be integrated numerically from minus infinity to plus infinity in x and y to determine how much total flux it represents. However, for practical reasons the analytic formulae typically contain only a few adjustable parameters, and these must be determined so as to optimize the profile match near the *center* of the stellar image, which will ultimately dominate the profile fits. There is no guarantee that the same analytic function will adequately mimic the extended wings of the profile. These wings usually span many arc seconds so, even though their surface brightness is low, they can still contain a non-negligible fraction of the total flux of the star.

A few computer programs (e.g., WOLF, DAOPHOT) represent the observed stellar profile—either in whole (WOLF) or in part (DAOPHOT)—by a digital array of brightness values actually observed for one or more bright stars in the frame; the profile is then evaluated by numerical interpolation within this data array. In principle, the tabulated brightness values could be summed to yield the total number of detected photons for the star, but once again high precision in the empirical profile is needed only for the central, bright portion of the star image which will dominate the actual profile fits and will not necessarily be achieved in the extended wings.

As a result, most profile-fitting photometry relies on two additional reduction steps, synthetic aperture photometry followed by a growth-curve analysis, to measure total integrated magnitudes for one or more stars in each program image. The mean difference between these "true" instrumental magnitudes and the relative profile-fitting magnitudes for the same stars establishes an absolute zero point for the profile-fitting results, which allows a valid comparison to be made between program-star data and observations of photometric standards in different CCD frames. Active practitioners of CCD photometry have long been familiar with this process (see, e.g., Da Costa, Ortolani, and Mould 1982, §II(a); Rich, Da Costa, and Mould 1984, §II; Stetson, VandenBerg, and McClure 1985, §II; Smith *et al.* 1986, §II, steps 9 and 10), and—while I do not recall ever seeing the statement made in print—many CCD photometrists believe it may be the single most crucial step limiting the absolute accuracy of the final photometry (e.g., G. S. Da Costa, private communication; R. D. McClure, private communication; M. Bolte, private communication). However, a comprehensive justification of the aperture growth-curve technique has only recently appeared in the refereed literature (Howell 1989); the reader is directed particularly to this paper for a thorough discussion, but I will provide a brief summary here.

2. Summary of the Aperture Growth-Curve Method

The principle of synthetic aperture photometry is simple: By adding up the observed flux in pixels within some (comparatively large) radial distance of the centroid of a star and subtracting from that total the contribution expected from the diffuse sky brightness, one obtains an estimate of the total flux from the star alone. To the extent that no other objects are contained within the aperture and the sky-brightness estimate is correct, and provided that the "aperture" radius is sufficiently large that seeing, tracking, and focus errors do not affect the fraction of the star's flux which falls *outside* the aperture, one then has a basis for comparing stellar fluxes among digital frames with different star-image profiles. The computer-defined "aperture" within which the pixel values are summed in the two-dimensional data array is a direct analog to the physical aperture placed at the telescope focus in a standard photometer.

The need for a growth-curve analysis is more subtle and rests on signal-to-noise arguments. As discussed briefly by, for instance, Da Costa *et al.* (1982) and Stetson *et al.* (1985), and recently in much more detail by Howell (1989), a larger synthetic aperture obviously contains more stellar flux (= signal) than a smaller one. However, the *rate* at which the total signal grows with increasing aperture size declines as the wings of the stellar profile tail off toward zero intensity. On the other hand the noise of the measurement grows rapidly with increasing aperture radius. Such *random* errors as the total readout noise contained within the measuring aperture, the Poisson shot noise in the diffuse-sky contribution, and flat-field errors on small spatial scales all grow as the square root of the number of pixels—that is, linearly with aperture radius. Local *systematic* errors caused by a misestimated sky brightness or the potential for contamination by unrelated astronomical objects grow linearly as the aperture area—that is, as the square of the radius. As a result, the signal-to-noise ratio of the flux measurement achieves a maximum value at some intermediate aperture radius (e.g., Howell 1989, Fig. 6). The aperture with the maximum signal-to-noise ratio will not necessarily be large enough to contain a seeing- and tracking-independent fraction of the stellar flux. Furthermore, the aperture with the maximum signal-to-noise ratio will be different for stars of different apparent magnitude.

The simplest way to improve the accuracy of the total flux measurement is to photometer each star through two concentric apertures, one small and the other large. The mean difference, in magnitudes, between the measurements through the two apertures is determined from as many bright, isolated stars in a given frame as possible. Then this mean difference is added as a correction to the small-aperture measurements for all the stars in the frame (e.g., Da Costa *et al.* 1982). In general, it is neither practical nor necessary that the larger aperture contain the total flux of the star. Provided it is large enough (typically of order 10″) that the fraction of light missed be independent of seeing, guiding, and focusing errors, and as long as the *same* large aperture be used for all program and standard-star frames, the constant missing flux will be absorbed into the zero point of the derived transformation equations.

A more powerful approach is to measure as many stars as possible through a series of *several* concentric apertures of increasing radius and to calculate the observed magnitude differences between successive apertures for each star. These are plotted as a function of radius, and a smooth curve is sketched through them to yield the average "growth curve"[2] of the frame. The average magnitude differences between successive apertures are then read from this curve and summed from the outside inward to yield cumulative corrections from each of the smaller apertures to the system of the largest. This multiple-aperture technique offers several advantages over the simpler two-aperture method.

(1) Some stars may not be measurable in the largest aperture: It may extend beyond the edge of the frame; it may include a companion star or a charged-particle event which contributes some unknown amount of flux, or an image blemish which obliterates some flux; or—for fainter stars—the signal-to-noise ratio of the measurement in the largest aperture may simply be too poor to be useful. On the other hand, such stars may still be measurable through several of the smaller apertures. Determining the mean magnitude difference between each successive pair of apertures from *all* the stars which can be reliably measured through *those two* apertures, and then summing the mean differences from the outside in, yields a more precise determination of the total correction from the smallest aperture to the largest than would be obtained if only those stars were used which could be measured in the largest aperture.

(2) The radius of the aperture providing the maximum signal-to-noise ratio for a brightness measurement depends on the apparent magnitude of the star. Having measured each star through a number of apertures of differing radius, one can then choose for each star the aperture which maximizes the signal-to-noise ratio of that star and correct *that* measurement to the system of the largest aperture.

(3) In plotting for each star the magnitude difference between successive apertures as a function of radius, one may see whether and at what radius the personal growth curve of each star begins to diverge from the mean growth curve for its frame. This will draw attention to previously unrecognized companions or defects and will help in deciding which is the largest reliable aperture for the star.

This is the process employed by Rich *et al.* (1984), Stetson *et al.* (1985), Smith *et al.* (1986), and others, as summarized by Howell (1989).

This empirical aperture-growth-curve method still has some problems, however. First, it is extremely tedious. Even with some computer assistance it can take many minutes per frame to compute the magnitude differences between successive apertures, plot them as a function of radius, identify questionable data, produce a smooth curve through the rest, read the mean corrections from the curve, add up the cumulative corrections for the various apertures, and apply them to the raw observations. To carry out this process for an observing run of some hundreds of frames could take weeks. The vast amount of effort involved encourages the taking of short-cuts, such as using only the brightest few stars per frame rather than all stars which could contribute usefully to the mean growth curve, and giving all stars equal weight in sketching the mean curves, rather than weighting by the estimated standard error of each aperture measurement; it discourages repeating the process to check for computational mistakes or to experiment with different sets of apertures. Second, the method requires that every frame of interest contain some bright stars that can be reliably measured out to the radius of the largest aperture. This is not always the case. Underexposed frames may have no stars bright enough for high signal-to-noise measurements in a large aperture, yet may still be of scientific interest; even in well-exposed frames the bright stars may have companions, or may lie near defects or the edge of the frame; the diffuse sky brightness may be spatially variable on the scale of the largest aperture (an underlying galaxy or emission nebulosity, for instance), and hence the precision needed in correcting for its contribution to the largest aperture may not be achievable from the information available. Methods for dealing with faint companion stars and backgrounds which vary on large spatial scales have been described in the past: The fitted model profiles may be used to subtract companion stars

[2]The earliest published use of the phrase in this context which is known to me is in Rich *et al.* (1984).

from the frame (e.g., Stetson *et al.* 1985; Stetson and Harris 1988; see also Penny and Dickens 1986, §3.3, for a different approach), and robust filtering techniques can remove *smooth* (i.e., on spatial scales much larger than a stellar image) background variations (e.g., Stetson and Harris 1988). However, the other possible complications just mentioned can still render a classical growth-curve analysis unfeasible for some frames.

What is needed is a computer program which (a) is completely automatic, enabling the computer to evaluate the growth curve of a frame and apply it to the raw measurements, deciding for itself which apertures are to be used for which stars without human guidance; and (b) can process the data from an observing run in a time comparable to that required for other CCD reduction stages such as flat fielding (i.e., much less than a day), so this step will no longer be a serious bottleneck. To be successful in achieving these goals, the program should (c) have a priori expectations of reasonable growth-curve morphologies and (d) be able to take a synoptic view of all the growth curves for all the frames from a given night or observing run. The program must have this "experience" with representative stellar profiles to be able to judge reliably which observed aperture differences are likely to be erroneous—a feature which is particularly important for frames having empirical growth curves derived from only one or two stars, so that purely internal consistency checks are impossible. Furthermore, under the assumption that stellar profiles which are similar at small radii will also be similar at large radii, the computer can use this experience to extrapolate growth curves to the largest radius even for frames which contain *no* stars with reliable measurements in the larger apertures. The next section describes a working prototype of such a program. A semi-interactive version of this method was employed by Stetson and Harris (1988); on the basis of that experience a completely automated version has now been written and is currently being employed in other photometric investigations.

3. General Principles

The fundamental assumption underlying the present technique is that it is possible to define a one-parameter family of aperture growth curves for a given night or observing run, in which the azimuthally averaged outer part of the stellar profile is asymptotically independent of seeing, guiding, and focusing errors. This assumption is suggested by the study of King (1971), who provided an empirical description of the internal structure of star images. According to King, much of the flux of a star is contained within a small, bright "core", in which surface brightness is well approximated by a Gaussian function of radius; this Gaussian form is almost certainly imposed by the statistical effects of seeing and guiding. The extended outer "aureole" of the image falls off as an inverse power of radius with an exponent close to 2 and is probably caused by scattering in the atmosphere and in the telescope optics. (However, the surface brightness cannot really fall as slowly as r^{-2} because integration over all radii would imply a divergent total flux.) Between the Gaussian inner core and the outer power-law aureole, King found a transition region which is reasonably well matched by an exponential function of radius, whose physical cause is not intuitively obvious.

For the present work, then, my specific assumptions are that (1) to a reasonably good approximation a stellar profile may be represented by a sum of a Gaussian function, an exponential, and a third function which resembles an inverse power law with exponent $\gtrsim 2$ at large radii; (2) the radial scale lengths of the Gaussian and exponential components of the profile will differ from frame to frame due to differences in the seeing conditions, guiding history, and degree of telescope defocusing; and (3) the relative fractions of stellar flux contained in the seeing-dependent and seeing-independent parts of the stellar profile will either (a) be constant properties of a given night or observing run or (b) depend solely on the air mass of the observation.

The fundamental equation which I use to describe the general stellar profile is

$$\begin{aligned} I(r,X_i;R_i,A,B,C,D,E) = {} & (B+E\cdot X_i)\cdot M(r;A) \\ & +(1-B-E\cdot X_i)\cdot[C\cdot G(r;R_i) \\ & +(1-C)\cdot H(r;D\cdot R_i)] \ , \end{aligned} \tag{1}$$

where r is a radial distance measured (in pixels) from the center of the concentric apertures which are in turn assumed to be concentric with the star image; X_i is the airmass of the ith data frame (known a priori); R_i is the seeing- (guiding-, defocusing-) related radial scale parameter for the ith data frame; and M, G, and H are Moffat, Gaussian, and exponential functions, respectively:

$$M(r;A) = \frac{A-1}{\pi}(1+r^2)^{-A} \ ,$$

$$G(r;R_i) = \frac{1}{2\pi R_i^2}\exp\left(-r^2/2R_i^2\right) \ ,$$

$$H(r;D\cdot R_i) = \frac{1}{2\pi(D\cdot R_i)^2}\exp\left[(-r/(D\cdot R_i))\right] \ .$$

In the reduction of the data from a given night or observing run, a separate value of the image-radius parameter R_i is computed for each image, but the values of the other fitting parameters $A,\ldots,E$ are required to be the same for all frames. Some specific points to note: (1) The parameter A affects the asymptotic power-law slope of the outer, seeing-independent part of the profile, and for a physically reasonable profile we must have $A > 1$; (2) the fraction of the total flux of the star which is contained in the aureole is given by $B + E \cdot X_i$ (the parameter E allowing the relative strength of the wings to depend linearly on airmass); (3) the parameter C defines the relative importance of the Gaussian and exponential contributions to the seeing-dependent part of the profile; (4) D permits the Gaussian and exponential components to have different—though linearly related—seeing-imposed scale lengths. Finally, (5) a Moffat function is used rather than a simple power law to prevent an unphysical divergence of the profile at $r = 0$. The Moffat function, unlike a power law, does have a characteristic scale length at small radii. However, this has arbitrarily been set to 1 pixel, which is invariant and is enough smaller than the smallest aperture commonly used ($\gtrsim$3 pixels, in my own work) that in the radius domain of interest this implicit scale length is unimportant. In my experience to date the parameters D and E are comparatively unimportant; they could be fixed at 0.9 and 0, respectively, as I will discuss below. Nevertheless, I include them as free parameters for greater completeness and in case future data sets turn out to require the additional flexibility.

Equation (1) is applied only in a differential sense. That is, given a set of instrumental magnitude measurements $m_{i,j,k}$ through a set of apertures $k = 1,\ldots,n$ with radii r_k (I assume throughout that $r_k > r_{k-1}$) for *stars* $j = 1,\ldots,N_i$ in frame i, we work with the observed magnitude *differences*

$$\delta_{i,j,k} = m_{i,j,k} - m_{i,j,k-1} \ , \qquad k = 2,\ldots,n \ .$$

These values of δ are fitted by a robust least-squares technique to equations of the form

$$\delta_{i,j,k} \doteq -2.5\log\left[\frac{\int_0^{r_k} I(r,X_i;R_i,A,B,C,D,E)(2\pi r)dr}{\int_0^{r_{k-1}} I(r,X_i;R_i,A,B,C,D,E)(2\pi r)dr}\right] . \quad (2)$$

Through no coincidence at all, the three functions which I chose to represent the components of the stellar profile all have analytic integrals,

$$\int_0^{r_k} M(r;A)\,(2\pi r)dr = 1 - (1+r_k^2)^{-A} \ ,$$

$$\int_0^{r_k} G(r;R_i)\,(2\pi r)dr = 1 - \exp(-r_k^2/2R_i^2) \ , \quad (3)$$

$$\int_0^{r_k} H(r;D\cdot R_i)\,(2\pi r)dr$$

$$= 1 - [1+r_k/(D\cdot R_i)]\exp[-r_k/(D\cdot R_i)] \ ,$$

and have been so normalized as to have unit total volume when integrated to infinity.

This growth-curve methodology is complementary to the profile-fitting techniques mentioned in Section 1. In profile-fitting photometry one strives to produce an accurate model of the bright, inner portions of a star image for the best possible *relative* intensity scaling. A crude estimate of the tail of the profile is acceptable because the surface brightness in the wings of a star image is low enough that even a poor estimate is sufficient to correct for the narrow sector of the wing of one star which overlies the core of another when their profiles are being fit. In evaluating the total flux of a star, however, the surface brightness in the wings must be multiplied by $2\pi r\Delta r$ as it is integrated over azimuth. Thus, any significant error in the inferred profile at large radii could make a major difference to the derived total flux. For precise and accurate *absolute* brightness measurements we must, therefore, obtain the best possible representation of the *outer* part of the profile. Conversely, for the growth-curve analysis the detailed structure of the star image inside the smallest aperture is immaterial: The fact that the exponential function has an unphysical cusp at the origin and the fact that an arbitrary radial scale length has been imposed on the Moffat function are unimportant, because the only way the profile inside the innermost aperture appears in the analysis is as the integral $\int_0^{r_1} I(r)\,(2\pi r)dr$, which is merely a constant contributing to the numerator and denominator of all equations (2).

In order to improve the signal-to-noise ratio of the low-surface-brightness wings, the two-dimensional data are compressed to a single, radial dimension by summation in the azimuthal direction. Even then, the assumption that these stellar profiles can be represented by a single one-parameter family of curves differing only in a seeing-imposed radial scale-length parameter will not be entirely valid. In particular, the effects of tracking vagaries, telescope aberrations, decentering, and defocusing will be different and will impose their own characteristic structures on the stellar profile. However, it may be expected that these differences will be reduced in importance by the azimuthal summing, and furthermore they should be most important in the smallest apertures; we may hope that toward larger radii the assumption of a one-parameter family will be asymptotically correct, at least to within the precision of the observations.

This analytic approach to growth-curve analysis offers one further advantage over the standard computation of empirical, mean aperture-to-aperture corrections: One may now relax the requirement that the outermost aperture contain a seeing-independent fraction of the stellar flux. The analytic approximation to the outer part of a stellar profile which is defined in equations (1) and (2) offers a means of extrapolating the integral of that profile to radii larger than the largest aperture, even if the total flux contained within the largest aperture is still affected to some extent by seeing. By integrating the analytic model for a given frame to some radius much larger than that of the largest aperture, one obtains a numerical correction for the lost flux. The error in the estimate of the total flux is now not the integral of the true stellar profile beyond the largest aperture but the integral at large radii of the difference between the true stellar profile and its analytic approximation.

The rest is conceptually straightforward. One simply writes a computer program which reads in the concentric aperture photometry for all the frames of a given observing run, computes the magnitude differences $\delta_{i,j,k}$, and fits them using equations (1) and (2) to obtain values for the R_i and $A,\ldots,E$ (ignoring "obviously spurious" data in the process). Then for the ith frame the model correction from the kth aperture to a hypothetical aperture of somewhat larger radius $r_T > r_n$ is given by

$$\hat{\Delta}_{i,k} \equiv \sum_{k'=k+1}^{k_n} \hat{\delta}_{i,k'} - 2.5 \log \left[\frac{\int_0^{r_T} I(r,X_i;R_i,A,B,C,D,E)(2\pi r)dr}{\int_0^{r_n} I(r,X_i;R_i,A,B,C,D,E)(2\pi r)dr} \right] ,$$

where the individual points on the model growth curve, $\hat{\delta}_{i,k}$, are obtained by substituting the final parameter values into equations (1)–(3). There are still some problems of detail to be considered, however, and I will describe my current working solutions for them in the next section.

4. Details

(1) Present-day computers are neither large enough nor fast enough to use all the stars that a modern telescope and CCD can record in a given observing run in the solution of equation (2); some selection of stars must still be made. My prototype implementation of the synoptic, analytic growth-curve method described above (a FORTRAN program which I call DAOGROW) works directly from data files generated by the PHOTOMETRY routine in the computer program DAOPHOT (Stetson 1987). Along with instrumental magnitudes measured through up to 12 concentric apertures of differing radii, PHOTOMETRY produces an estimated standard error of each aperture magnitude based on knowledge of the readnoise and gain of the CCD and the observed diffuse sky brightness in the vicinity of the star. DAOGROW allows the user to specify a maximum value for this standard error: When the files produced by PHOTOMETRY are read into DAOGROW, any aperture measurement whose standard errors are larger than this value will not be used. Furthermore, since these standard errors increase with aperture radius (see above), DAOGROW will ignore any star whose standard error in the next-to-smallest aperture exceeds this value. (Obviously, two valid aperture measurements are the minimum requirement for at least one valid magnitude difference.) In addition, DAOGROW imposes a maximum on the number of stars from any given frame which will be used in the reductions; at present this number is 40, but it is easily changed. Thus, if before DAOGROW is executed the files containing the concentric aperture photometry are sorted by increasing apparent magnitude in the next-to-smallest aperture (e.g., with the SORT routine in DAOPHOT), for each frame DAOGROW will use the brightest 40 stars having at least one reliable magnitude difference; even for those stars it will only use data out to the radius where the standard error reaches the specified maximum.

This selection procedure could impose a bias on the results, arising from the fact that magnitude errors are relative: $\sigma(magnitude) \sim \sigma(brightness)/brightness$. As a result, since it is $\sigma(brightness)$ which is actually computed from first principles, if an aperture measurement is accidentally too bright its $\sigma(magnitude)$ will be correspondingly too small, and the measurement is more likely to be included. This bias can be fought by adopting a limiting σ which is comparatively large, say 0.10 mag: Since bright, well-exposed stars typically have combined readout and Poisson errors ~ 0.001 mag, a star near the cutoff—where the bias is imposed—would have a natural weight of order 10,000 times less than the bright stars that will dominate the solution. Furthermore, the weight-fudging scheme described in item (5) below will also preferentially reduce the influence of observations whose estimated σ's are too small for their true accuracy. In any case, to the extent that this bias is likely to be similar for program and standard frames, it may be ignored.

(2) For crowded frames, of order 20 to 40 bright stars in each field are chosen by visual inspection, and the profile-fitting and star-subtraction routines of DAOPHOT are used to remove surrounding companion stars before the concentric aperture photometry is carried out for the hand-picked ones. This step is the current bottleneck and is not usually necessary in standard-star frames and other sparse fields.

(3) The magnitudes observed for a given star through a set of concentric apertures are not statistically independent: The flux measured through one aperture will also be a major fraction of the flux measured through the next larger one and, furthermore, the same sky estimate is used for each. As the standard error of a given magnitude difference (needed for properly weighting the data), I therefore simply use the standard error of the magnitude from the larger of the two apertures. Thus, if

$$\delta_{i,j,k} \equiv m_{i,j,k} - m_{i,j,k-1} ,$$

then

$$\sigma(\delta_{i,j,k}) \equiv \sigma(m_{i,j,k}) \equiv \sigma_{i,j,k} .$$

A more correct weighting scheme could be devised, but I doubt that it would make much difference to the results.

(4) Convergence of the solution to equations (2) could be tricky to achieve. First, the equations are extremely nonlinear in the fitting parameters, $R_j,A,\ldots,E$. Second, even for bright stars, and even if they are handpicked, measurements of some stars through some apertures will be corrupted by unrecognized companions or image defects. For these reasons, cookbook-style weighted least squares cannot be relied on to give the optimum growth-curve solution. Instead, (a) the computer must be given

some reasonable guidance on likely answers, and (b) it must be given some capability to recognize and disregard erroneous data.

In regard to (a), as in all nonlinear least-squares problems, starting values must be provided for the fitting parameters. In all data which I have played with to date—from Kitt Peak National Observatory (4-m prime focus and No. 1 0.9-m Cassegrain), Cerro Tololo Inter-American Observatory (4-m prime focus and 0.9-m Cassegrain), and the Canada-France-Hawaii Telescope (prime and Cassegrain foci)—the following have turned out to be reasonable first guesses; exponent in the denominator of the Moffat function, $A \sim 1.2$; fraction of light contained in the seeing-dependent component which is in the Gaussian (as opposed to the exponential) function, $C \sim 0.5$; ratio of the exponential scale length to the Gaussian scale length, $D \sim 0.9$; and airmass dependence of the fraction of light in the Moffat-function aureole, $E \sim 0.00$ airmass^{-1} (which suggests that scattering in the optics typically dominates over scattering in the atmosphere). In the event of reasonable evidence supporting other parameter values, these starting guesses may be easily changed, but I expect that the values given are within the capture radii of most reasonable solutions.

At the beginning of a reduction, equation (2) is solved with the values of $A,\ldots,E$ fixed at the starting values, and solutions are performed only for the scale lengths of the individual frames, R_i. When these solutions have converged, the other parameters are freed one by one starting with A. With each iteration of the nonlinear solution, the standard deviation of the fitting residuals is computed, and when this stops decreasing (i.e., provisional convergence has been achieved) an additional parameter is freed, until final convergence has been achieved with all parameters free at once.

Finally, to obtain physically reasonable solutions even from poor data, it is necessary to impose the following limits on the parameter values:

$$\begin{aligned} A &> 1 \ , \\ 0 \le B &\le 1 \ , \\ 0 \le C &\le 1 \ , \\ D &> 0 \ . \end{aligned} \tag{4}$$

When a parameter approaches one of these limits, the maximum correction that any one iteration may make to the current value of the parameter is reduced, so that the parameter is prevented from passing beyond the bound of physical reasonableness (although it still *can* approach arbitrarily close to that bound). It might be cynically expected that during the reductions one or more of the parameters would tend toward these externally imposed limits and then stick there. This turns out not to be the case. In fact, with one exception to be discussed below, a physically reasonable final convergence is always achieved with all the parameters somewhere in the middles of their allowed ranges. The important thing seems to be to prevent the parameters from overstepping these bounds during the early iterations, while the program is thrashing around looking for an approximate solution and trying to decide which observations are spurious (see item (5) below). Once the reduction has survived this early period of uncertainty, the parameters invariably settle down to final values not far from the starting values I listed above.

(5) At the same time, to deal with complication (b) mentioned at the beginning of item (4), I employ a more sophisticated version of the blunder-rejecting scheme described briefly by Stetson (1987; §III.D.2.d) and which I hope to justify in somewhat more detail in Stetson (in preparation). Specifically, after the first pass through the data all the normal (viz., $1/\sigma^2$) weights of all observations are multiplied by a fudge factor which depends on the size of the residual, $v_{i,j,k} \equiv \delta_{i,j,k} - \hat{\delta}_{i,k}$ (where the definition of the model-predicted value, $\hat{\delta}_{i,k}$, is given at the end of the preceding section):

$$w_{i,j,k} = \sigma_{i,j,k}^{-2} \times F_{i,j,k} \ , \tag{5}$$

where

$$F(v_{i,j,k}/\sigma_{i,j,k};a,b) = \left[1 + \left| \frac{v_{i,j,k}}{a \cdot \sigma_{i,j,k}} \right|^b \right]^{-1} .$$

F has the following properties: $0 \le F \le 1$; $F(v/\sigma) = F(-v/\sigma)$; $F \to 1$ as $v/\sigma \to 0$; and $F \to 0$ as $v/\sigma \to \pm\infty$. In fact, under the assumption that when a stellar measurement through a particular aperture is dubious, all measurements of that star through larger apertures will also be dubious, to a given magnitude difference $\delta_{i,j,k}$ I actually apply the minimum of the F values computed for that star from that aperture and all smaller ones:

$$F_{i,j,k} \equiv \mathrm{Min}\,[F(v_{i,j,k'}/\sigma_{i,j,k'};a,b)], \qquad 2 \le k' \le k \ .$$

Note that it is only the weights that are fudged, and not the observations themselves. Throughout the reductions the quantity a is fixed at a value of 2 (easily changed), so after the first pass any observation which departs from the current estimate of the mean growth curve of its frame by 2σ is given half-weight, with correspondingly lower weight for larger discrepancies. At the beginning the quantity b is set to unity. It may be shown (Stetson, in preparation) that this has the effect of driving the iterated solution of equation (2) toward the mean of those observations with errors $\ll 2\sigma$ and toward the median of those observations with errors $\gg 2\sigma$. Thus, even if the initial guesses are far from the correct parameter values, nevertheless the derived growth curves will be pulled toward the center of the bulk of the observations. When the reduction has reached the point where the fitting parameters A, B, and C are all free (see item (4) above), the value of b is changed to 2. This further reduces the weight of the most discrepant observations and restores the fudge factor for observations with residuals $< 2\sigma$ to a value closer to unity; this drives the solution toward a broadly defined weighted mode of the observations. Finally, after a provisional convergence has been achieved with all parameters free, the value of b is changed to 3. The weights of the extreme outliers are still further reduced, the concordant observations receive weights still closer to their natural values, and the solution is driven toward a more narrowly defined weighted mode of the observations. Extensive experience shows that this procedure is extremely effective in leading the computer to arrive at about the same solution that an experienced astronomer would have sketched by hand through data containing some corrupted values.

(6) As I mentioned in Section 4 above it is unreasonable to expect that a single one-parameter family of growth curves will accurately model everything that can go wrong with a stellar profile. In particular, defocusing will not broaden the stellar core in the same way as poor seeing, and the effects of erratic guiding will be quite different still. Nevertheless, one may with some reason expect that these departures from the nominal model growth curve will be most noticeable in the smallest apertures. Fortunately, it is precisely here that the raw observations are best: A given frame will contain many more stars which can be effectively measured through small apertures than through large ones, and those measurements will contain the smallest random and systematic errors. Therefore, in parallel with the *model* growth curve for the ith frame (the $\hat{\delta}_{i,k}$ given by the robust least-squares solution of equations (1) and (2)), one may also generate the *empirical* growth curve, $\bar{\delta}_{i,k}$, given by a straight robust average of the $\delta_{i,j,k}$ values actually observed for stars in the frame, without any assumptions as to analytic form. (The empirical growth curve, of course, is exactly what we CCD photometrists have been using right along: Rich *et al.* 1984; Stetson *et al.* 1985; Smith *et al.* 1986; Howell 1989; etc.) In the computerized calculation of these $\bar{\delta}_{i,k}$ I use a weight-fudging scheme similar to that described in paragraph (5) above: successive fudged means with $a = 2$ and $b = 1$, 2, and 3 are taken to drive the average difference first toward the weighted median and then toward the weighted mode of the raw magnitude differences observed with a given aperture pair for all stars in a given frame. In this instance the residuals $v_{i,j,k}$ are computed with respect to the current estimate of the empirical growth curve, not the model growth curve. The *adopted* growth curve for a frame, δ^*, is then some compromise between the empirical mean growth curve and the analytic model growth curve, with the compromise heavily weighted toward the empirical curve at small radii—where the raw data are good but the actual profile may depart from the one-parameter family—and toward the model growth curve at large radii—where the raw data for any given frame may be poor or nonexistent but the assumption of a one-parameter family of curves should not be too bad. The current scheme for effecting this compromise is simply to take a weighted mean of the empirical and model growth curves at all radii. The weight of each point on the empirical curve is just the sum of the fudged weights of the individual raw differences going into the average (eq. (5)). The weight to attribute to a point on the analytic model curve, however, is much harder to quantify. It boils down to the question of just how much the mean growth curve of one frame can be expected to depart *legitimately* from the one-parameter family of the observing run (or night) and includes such considerations as the difference between seeing which is constant at some value during the exposure, and seeing which is variable about a mean value (since the sum of several Gaussians is not Gaussian), and the different effects of seeing, guiding, and defocusing. Furthermore, it can even be imagined that the growth curve for any particular star may legitimately differ from the mean growth curve for its own frame as a result, for instance, of optical aberrations or decentering. However, some experimentation suggests that

$$\sigma(\hat{\delta}_{i,k}) \times 0.1 \cdot \hat{\delta}_{i,k}$$

is not too bad an estimate of these additional uncertainties. (The numerical value of the multiplier can easily be changed if data warrant it.) On the basis of these ideas, then, the adopted growth curve for the frame is given by

$$\delta^*_{i,k} = \frac{\left(\sum_j w_{i,j,k}\right)\bar{\delta}_{i,k} + \left(\frac{1}{0.1\cdot\hat{\delta}_{i,k}}\right)^2 \hat{\delta}_{i,k}}{\left(\sum_j w_{i,j,k}\right) + \left(\frac{1}{0.1\cdot\hat{\delta}_{i,k}}\right)^2} ,$$

and the uncertainty associated with the adopted value for each point on the adopted curve is given by

$$\sigma^2(\delta^*_{i,k}) = \frac{1}{\left(\sum_j w_{i,j,k}\right) + \left(\frac{1}{0.1\cdot\hat{\delta}_{i,k}}\right)^2} .$$

The adopted *cumulative* growth curve is calculated by summing the adopted differential growth curve from the outside in, with an outer tail given by extrapolation of the analytic model:

$$\Delta^*_{i,k} = \sum_{k'=k+1}^{n} \delta^*_{i,k'} - 2.5 \log\left[\frac{\int_0^{r_T} I(r,X_i;R_i,A,B,C,D,E)(2\pi r)dr}{\int_0^{r_n} I(r,X_i;R_i,A,B,C,D,E)(2\pi r)dr}\right] ,$$

$$\sigma^2(\Delta^*_{i,k}) = \sum_{k'=k+1}^{n} \sigma^2(\delta^*_{i,k'}) .$$

I have neglected the error of the extrapolation from r_n to

r_T because it is only the error in the *seeing-dependent* part of this extrapolation that matters: Any error in the constant part of the tail will be absorbed into the zero point of the transformation. Unless the aperture radii are very poorly chosen, the error in the seeing-dependent part of the outer tail will be much smaller than the other uncertainties in the analysis. (I note at this point that it is probably not a good idea to set $r_T = \infty$, because if the least-squares solution of equation (2) should result in $A \approx 1$, then the outer part of the profile approaches an inverse-square law, the integral of which diverges when taken to infinity. A slight error in the parameter E could then introduce a large airmass-dependent error into the extrapolation. This would be absorbed into the extinction coefficient for the night but, nevertheless, should be avoided if possible. I have adopted $r_T = 2r_n$.)

(7) Now that we have complete cumulative growth curves for all apertures and all frames, each aperture magnitude measurement may be corrected to the system of the (hypothetical) aperture of radius r_T:

$$m^*_{i,j,k} = m_{i,j,k} + \Delta^*_{i,k} \ .$$

In estimating the uncertainty of this corrected magnitude, we must include the standard error of the raw magnitude measurement itself as well as the standard error of the adopted growth-curve correction: normally this would be

$$\sigma^2(m^*_{i,j,k}) = \sigma^2(m_{i,j,k}) + \sigma^2(\Delta^*_{i,k}) \ .$$

In choosing *which* aperture k to use for a given star, one could then exploit the fact that the error of the raw magnitude measurement $[\sigma(m_{i,j,k})]$ grows rapidly toward larger aperture radii, while the error of the correction $[\sigma(\Delta^*_{i,j,k})]$ grows rapidly toward smaller radii. Thus, the combined error for a given star will have a minimum in some—usually intermediate—aperture. However, at this point, I must remind the reader that there is one final complication: Some aperture measurements for any given star may be corrupted by companion stars, image defects, or a poor sky determination. I have already introduced a mechanism for recognizing and dealing with the suspicion of such contamination, namely the weighting fudge factors F discussed in paragraphs (5) and (6) above. I recycle that fudge factor here by setting the uncertainty of the corrected magnitude equal to

$$\sigma^2(m^*_{i,j,k}) = \frac{\sigma^2(m_{i,j,k})}{F_{i,j,k}} + \sigma^2(\Delta^*_{i,k}) \ .$$

If the individual growth curve for a particular star begins to diverge from the adopted growth curve for its frame due to whatever reason, the program begins to lose confidence in those measurements: For that aperture where the divergence first becomes apparent (that is, the residual becomes recognizably large in comparison with its standard error) and for all larger apertures, the F factor of the star decreases, the estimated errors of the corrected magnitudes increase, and $\sigma(m^*_{i,j,k})$ now achieves its minimum value in some smaller aperture, where the contamination is not yet significant. The corrected magnitude based on *this* aperture is adopted as the total magnitude of the star, $m^*_{i,j}$, and the minimum σ itself becomes $\sigma(m^*_{i,j})$.

Adopting the corrected magnitude from that aperture where $\sigma(m^*_{i,j,k})$ is minimized differs slightly from the approach recommended by Howell (1989), who suggests that the weighted average of the corrected magnitudes $m^*_{i,j,k}$ for a given star through several consecutive apertures k be adopted as the total magnitude for the star. Since the measurements through the various concentric apertures are highly correlated, it is not clear that Howell's approach adds significant additional information and, in fact, it may be more sensitive to contamination than the approach suggested here.

5. Examples

During three photometric nights in June 1988, I acquired 203 CCD frames with RCA5 on the Cerro Tololo 4-m telescope. Once concentric-aperture photometry had been derived for these frames, DAOGROW required 25 minutes on a VAX 11/780 to compute the 203 synoptic growth curves and the corrected total stellar magnitudes—a task that would have taken me probably one or two weeks in the old days and most likely would not have been done as well even then.

Of course, that 25 minutes represents only one part of the total reduction chore. Since I have automatic routines to set up the star-finding and aperture-photometry tasks for batch processing, for the standard-star frames and other uncrowded fields the only astronomer effort required to obtain the concentric-aperture photometry was that involved in typing the image file names into the computer, viz., a few seconds per frame; the star finding and aperture photometry then ran overnight and were ready the next morning. For the crowded (in this case, globular-cluster) program frames, on the other hand, much more time was needed to define the point-spread functions (PSFs) that are required for the profile-fitting photometry. Working at top speed I can presently define of order 20 to 25 PSFs per normal working day; with the interruptions common at any workplace, rather less than this is more typical (although I hope for increased throughput some day, with new software and a "private" workstation). Thus, one astronomer week per night of observing is a reasonable order-of-magnitude estimate for the time spent in this task, which has always been a, or *the*, major bottleneck in DAOPHOT reductions. However, once this chore has been completed, no significant *additional* astronomer time needs to be spent in choosing stars for the aperture-growth-curve analysis: It is simple to arrange it so that those stars which were chosen to define the PSF (typically 20–30 per frame) are subjected to new concentric-aperture photometry after all other stars have been subtracted from the frame, and automatic software to perform this task has been written.

Clearly, 25 minutes of CPU time is little enough that the computation and application of growth curves should no longer be a major contributing factor to the effort of reducing our CCD photometry. It is perhaps not fast enough for real-time photometry in the dome (possible for uncrowded fields only), but the philosophy behind DAOGROW is inappropriate to that application anyway since the synoptic growth curves should be based on all the frames from a given night or observing run which are, of course, not available until the night or run is over.

The random errors introduced by the procedure are hard to evaluate quantitatively. When repeat observations for particular stars are intercompared (many fields were observed a dozen or more times during the run), the observed scatter in the individual corrected magnitudes is found to exceed the raw read and shot noise by amounts in the range 0.000–0.005 mag (rms)—that is, these are the standard deviations which must be added in quadrature with the $\sigma(m^*_{i,j})$ defined above to reproduce the observed scatter. This estimate of 0.000—0.005 mag includes all sources of observational scatter besides readout noise and Poisson photon errors including, e.g., flat-fielding errors, wispy clouds and other extinction variations, and shutter-timing errors, in addition to any scatter introduced by inadequacies in the growth-curve analysis. This is as good as or better than the results seen in previous studies (e.g., Stetson *et al.* 1985; Smith *et al.* 1986; McClure *et al.* 1987; Hesser *et al.* 1987) where classical growth-curve analyses were performed manually and is comparable to the results achieved by Stetson and Harris (1988), who employed a process similar to that described here but with much more manual intervention. M. Bolte (private communication) estimates that DAOGROW reduces the scatter in his photometry by as much as a factor of 2 in comparison with his own hand-sketched growth curves. I therefore conclude on the basis of the currently available evidence that his growth-curve methodology produces results no poorer than those obtainable by more tedious methods.

Figure 1 illustrates the analytic-model differential growth curves for the best- and poorest-seeing frames of the three nights and also for three intermediate values of R. The actual computed values of $\hat{\delta}_{i,k}$ are plotted as open circles, and these are connected by smooth curves to make the separate growth curves clearer. The radii of the 12 apertures formed a geometric sequence from $r_1 = 3.0$ to $r_{12} = 20.0$ pixels, with

$$r_k = \sqrt[11]{20/3} \cdot r_{k-1}, \qquad k = 2,\ldots 12 \ .$$

Note that the growth curves in Figure 1 still have a

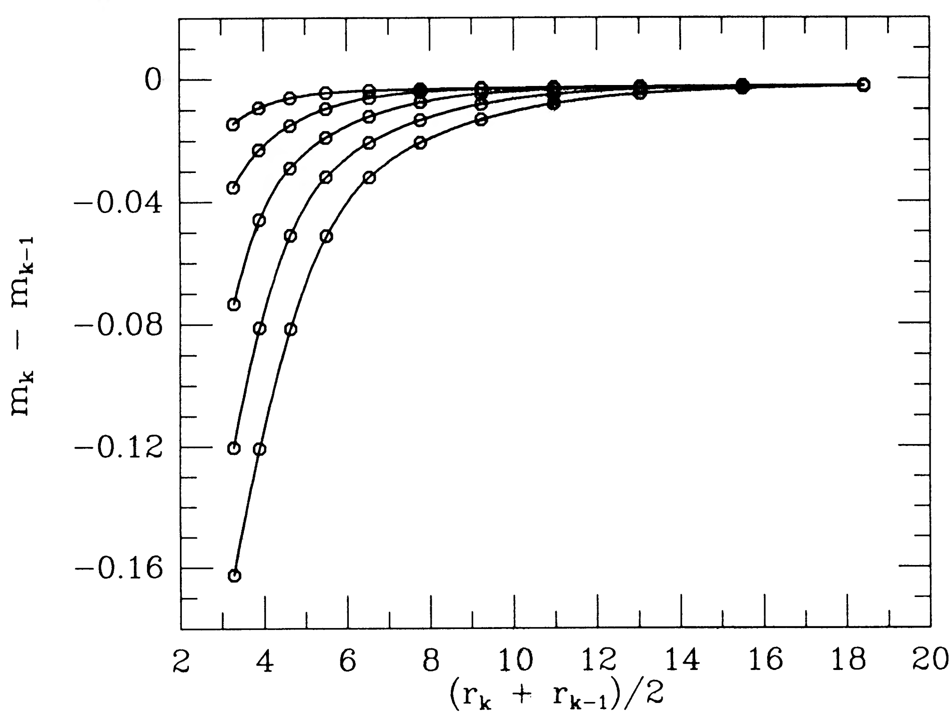

FIG. 1–Five selected analytic model growth curves from an observing run in 1988 at the prime focus of the Cerro Tololo 4-m telescope. Shown are the model curves for the individual frames with the best and poorest seeing of the entire three-night run and curves for three intermediate seeing values. The points represent the model-predicted magnitude differences for aperture magnitudes from consecutive concentric apertures, plotted as a function of the mean radius of the two apertures. Solid lines connect the magnitude differences corresponding to each individual growth curve.

slightly nonzero value even at the largest radius. That is because of the large radial extent of the outer tail of the profile: The geometric increase in the aperture radii partly compensates for the power-law falloff to produce a perceptible increment of flux even in the largest apertures. The effects of seeing—at least on these three nights—appear to have faded away by the radius $r_{11} = 16.8$ pixels $\approx 10\farcs1$, since some frame-to-frame differences are still barely perceptible in the model $\hat{\delta}_{11} = \langle m_{11} - m_{10} \rangle$ values but not in $\hat{\delta}_{12} = \langle m_{12} - m_{11} \rangle$.

Figure 2(a) shows as open circles the raw observed magnitude differences for the frame with the best seeing; note that some artificial scatter has been introduced in the abscissae of these points to reduce overlap. The model growth curve is shown as short dashes, the empirical growth curve is shown as long dashes, and the adopted growth curve is shown as a solid line; it may be seen that all three curves coincide at small radii, while at large radii the empirical curve becomes poorly determined and wanders off, while the adopted curve remains faithful to the smooth analytic model. This frame has a seeing parameter $R = 0.650$ pixel, and the total magnitude correction from r_1 to r_T (arbitrarily defined as $2r_n$) is -0.0625 ± 0.0027 mag; that is, 94.4% of the total flux of a star within 40 pixels = 24″ is inside 3 pixels = $1\farcs8$, and the growth-curve correction contributes < 0.003 mag estimated uncertainty to the total flux of a star as extrapolated from the measurement in the smallest aperture. For comparison, Figure 2(b) shows the raw data and the three curves for the poorest-seeing frame, which has $R = 1.856$ pixels. Here the analytic model does not appear to match the raw data perfectly at small radii ($r \lesssim 8$ pixels); accordingly, the adopted curve follows the empirical one at small radii, transferring its loyalty to the analytic model only for the largest apertures ($r \gtrsim 13$ pixels). According to the adopted curve, the total correction from the three-pixel aperture to $r_T = 40$ pixels is -0.4942 ± 0.0220 magnitude: Less than two-thirds of the light of a star falls within the $1\farcs8$ aperture in this case, and a total magnitude estimated on the basis of the smallest aperture would be uncertain by at least 2% after application of the growth-curve correction. However, the correction from aperture 5, with $r_5 = 6.0$ pixels = $3\farcs6$, to r_T is -0.0981 ± 0.0051; hence, some 91% of the flux is within a radius of 6″, and the program (and the astronomer) would be more inclined to prefer corrected magnitudes from these larger apertures where the uncertainty in the correction is smaller.

Raw data and growth curves for some other frames are shown in Figures 3–5. These have been selected to illustrate particular points and should be regarded as pathological rather than typical. Figure 3 is from a frame of a globular-cluster field with such a short exposure time that, even though there are many stars in the frame, nearly all of them are quite faint. As a consequence the outer parts of the growth curves are poorly defined. Nevertheless, there is a sufficient number of stars whose inner profiles are reasonably distinct that the program can use its experience of the growth curves from similar frames to select the physically reasonable data points and produce a suitable growth curve out to large radii. Figure 4 shows data from a long-exposure globular-cluster frame. The field is so crowded that every star has many neighbors. Although these neighbors have been subtracted to the extent possible, the subtractions introduce additional noise into the cleaned frame. In some pixels the neighbor stars have been oversubtracted, leaving behind a blemish which causes the truncation in the growth curve of the star; in other pixels there has been undersubtraction, giving the effect of excess flux at large radii. Such a frame could hardly be touched with a traditional one-frame-at-a-time empirical growth-curve analysis. As in Figure 3, however, DAOGROW can use the concordant observations at small radii to select the appropriate member of the one-parameter family of model growth curves, based on its experience of less-crowded frames with similar seeing, and use it to produce the full set of cumulative corrections. Finally, Figure 5 shows the most extreme example I have ever seen of a frame whose empirical growth curve does not resemble any member of the one-parameter family of model curves of its observing run. This frame contains several Landolt (1973) standards, most of which are both well exposed and completely uncrowded. Faced with these data, DAOGROW adopts the empirical growth curve out to a radius ~ 6 pixels. Here, where the raw data become noticeably fewer and noisier, it begins to transfer its allegiance from the empirical curve to that model curve which best—albeit poorly—matches the overall profile. From $r \sim 10$ pixels ~ 6″ outward DAOGROW relies on the model curve alone.

The fourth night of my observing run at Cerro Tololo had much poorer and highly variable seeing, and it clouded up toward the end of the night. The seeing was so poor that in some frames a 3-pixel ($1\farcs8$) aperture contained ~ 10% of the total flux. When the innermost apertures reach this far into the stellar core, the current version of DAOGROW has revealed itself to be not entirely trustworthy: It tends to play the Gaussian, exponential, and Moffat functions off against each other, occasionally driving a parameter or two up against the limits in a desperate attempt to match details of the inner core profile. Therefore, I reduced the data of this night with a different set of apertures. Here I used $r_1 = 5$ pixels, $r_{12} = 20$ pixels, and a different progression of intermediate radii:

$$r_k = r_{k-1} + (k-1) \cdot \delta_r \ , \qquad \text{where} \qquad \delta_r = \frac{15}{66} \ .$$

This produced the family of analytic model growth curves shown in Figure 6, where again I have shown the model

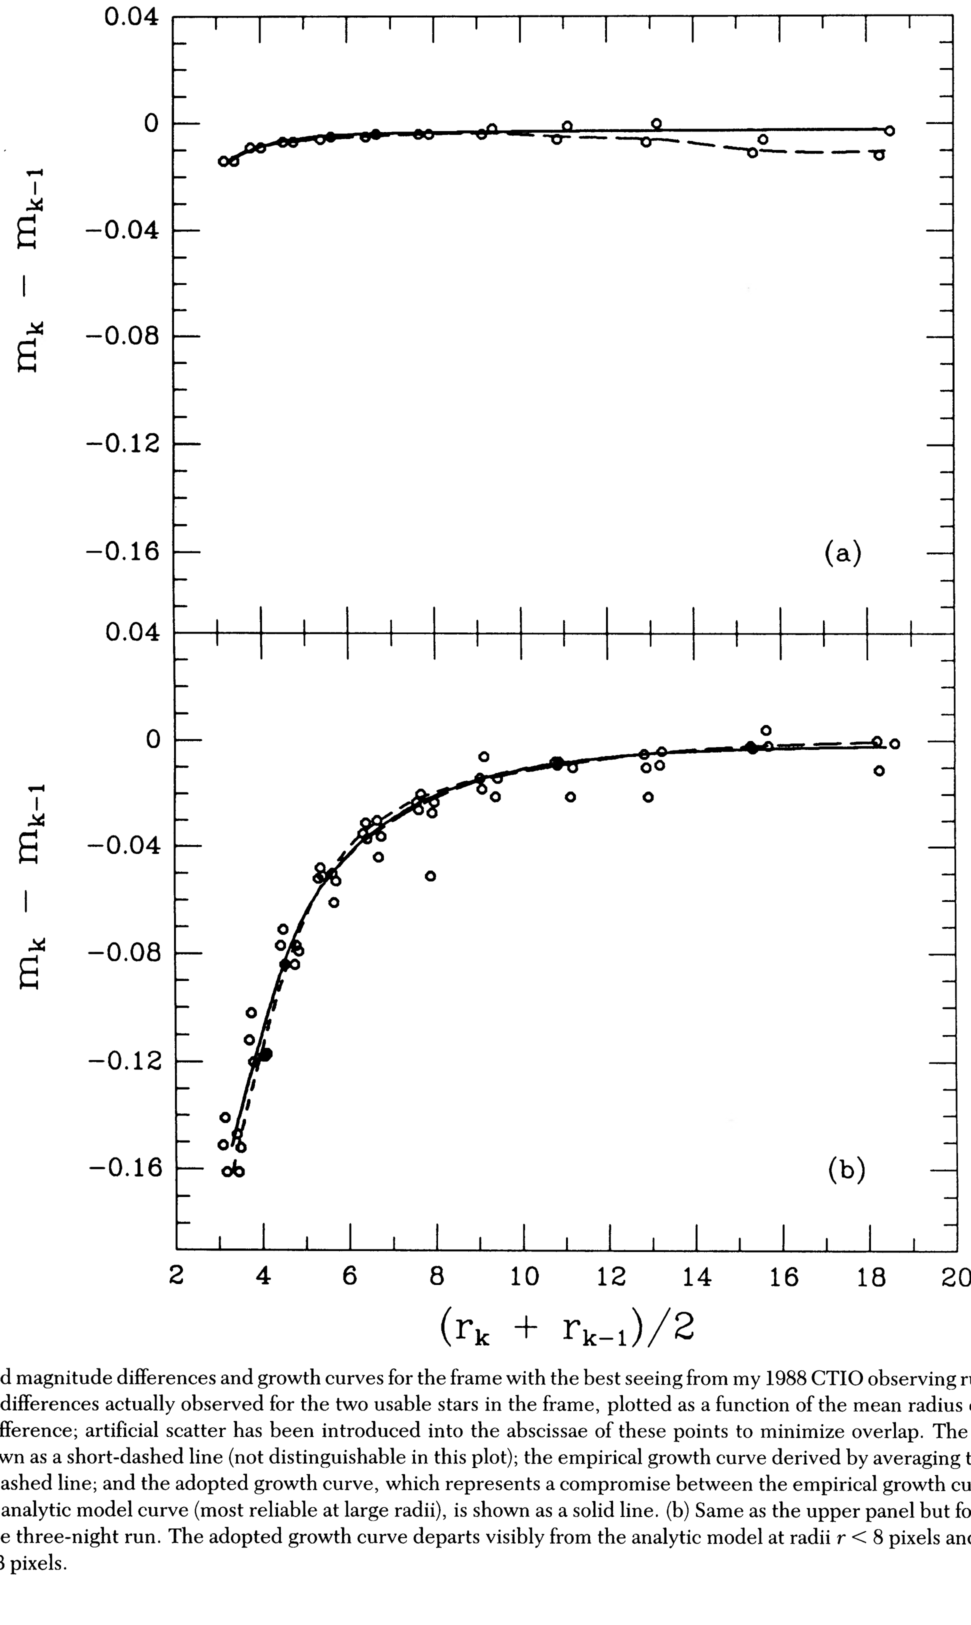

FIG. 2–(a) Observed magnitude differences and growth curves for the frame with the best seeing from my 1988 CTIO observing run. Points represent the raw magnitude differences actually observed for the two usable stars in the frame, plotted as a function of the mean radius of the two apertures involved in each difference; artificial scatter has been introduced into the abscissae of these points to minimize overlap. The best-fitting analytic growth curve is shown as a short-dashed line (not distinguishable in this plot); the empirical growth curve derived by averaging the raw observations is shown as a long-dashed line; and the adopted growth curve, which represents a compromise between the empirical growth curve (most reliable at small radii) and the analytic model curve (most reliable at large radii), is shown as a solid line. (b) Same as the upper panel but for the frame with the poorest seeing of the three-night run. The adopted growth curve departs visibly from the analytic model at radii $r < 8$ pixels and from the empirical curve at radii $r > 13$ pixels.

curves for the frame with the best seeing (R = 1.529 pixels, 90.5% of the flux within 3."1), the frame with the poorest seeing (R = 2.962 pixels, 92.0% of the flux within 5."9), and for three intermediate values of R. The raw data for the best- and poorest-seeing frames are reproduced in Figure 7. In Figure 6 some seeing dependence is still perceptible out at least to aperture 11 (r_{11} = 17.50 pixels ≈ 10."5), since a range of values is seen for $\hat{\delta}_{12} = \langle m_{12} -$

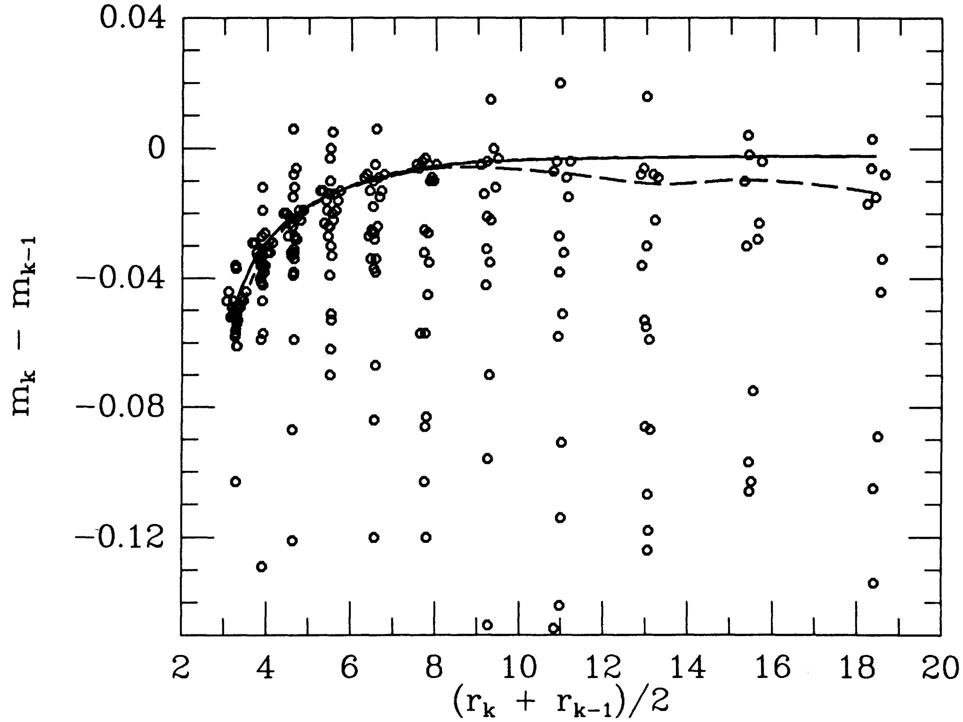

FIG. 3–Observed magnitude differences and growth curves for a shallow exposure of a crowded globular-cluster field. Symbols and line types have the same meaning as in Figure 2.

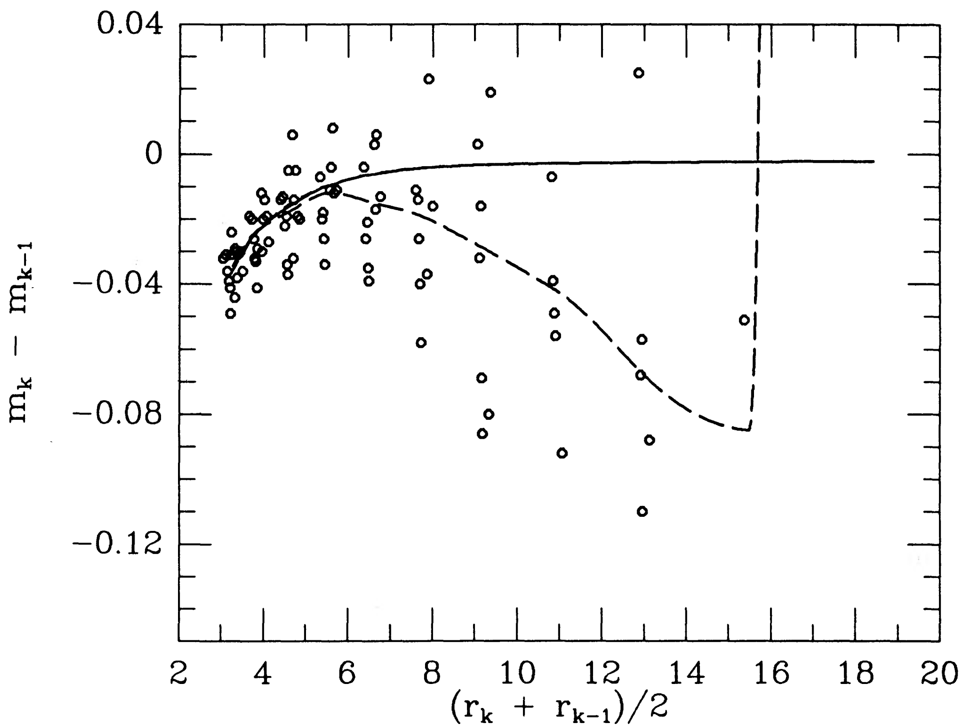

FIG. 4–Observed magnitude differences and growth curves for a deep exposure of a crowded globular-cluster field, where before performing the concentric-aperture photometry for individual bright stars the neighbors were subtracted out to the extent possible using results from model-profile fits. Symbols and line types have the same meaning as in Figure 2.

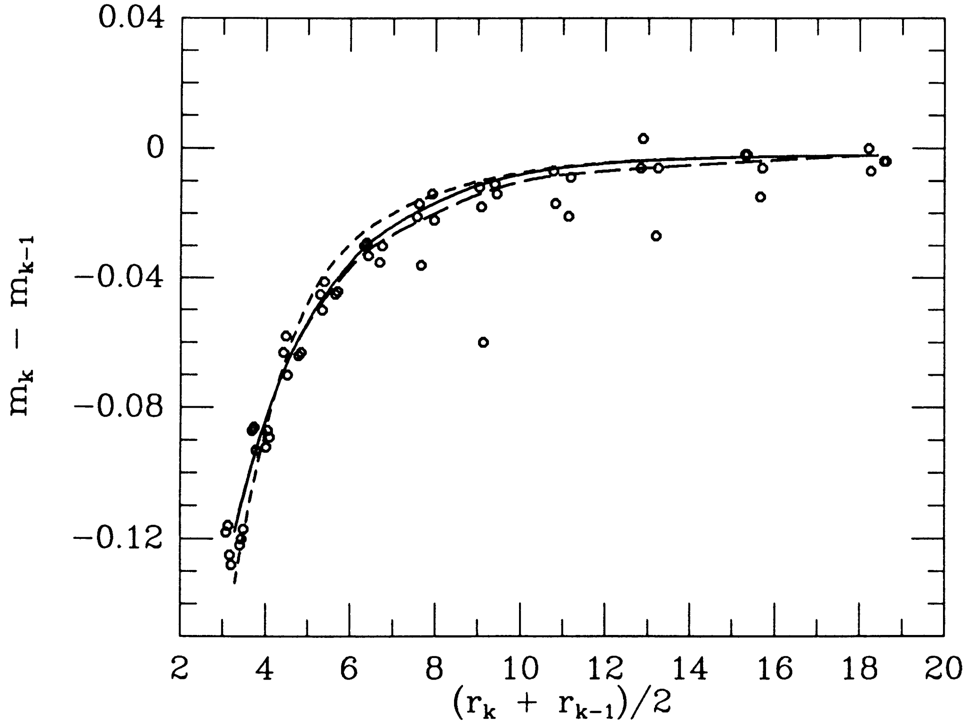

FIG. 5–Observed magnitude differences and growth curves for a frame where, although there are several bright, uncrowded stars, the empirical growth curve is not well matched by any member of the one-parameter family of analytic curves which match the remainder of the frames from the observing run. The departure of the best analytic model (short-dashed curve) from the actual data is quite pronounced at small radii. Accordingly, the adopted growth curve follows the empirical curve out to a comparatively large radius, making a full transition to the analytic model curve only for $r \gtrsim$ 10 pixels = 6″.

$m_{11}\rangle$. With such an odd progression of aperture radii and only hand-drawn graphs and empirical mean magnitude differences, it might be tricky to extrapolate a correction for the seeing-dependent part of the profile beyond the largest aperture; more likely the frames would have to be rephotometered with larger apertures to get out into the invariant part of the profile. The signal-to-noise ratio would be very poor in these larger apertures and the correct curves correspondingly difficult to draw. With the help of best-fitting analytic models, the extrapolation is straightforward.

Practical experience thus suggests that it is worthwhile to give some thought to the choice of aperture radii. DAOGROW seems to work well as long as the innermost aperture contains not much less than a third of the total flux of a star. However, even then the utility of the results can be affected by the way in which the aperture radii are distributed between the minimum and maximum values. The use of a constant difference between successive apertures, $r_k = r_{k-1} + \Delta r$, seems to be a particularly poor choice, since it gives too little spatial resolution at small radii, where there is the most leverage on the seeing parameter (R_i) of a frame. Conversely, a linear progression of aperture radii makes the annuli too narrow at large distances, where the surface brightness is low, the signal-to-noise ratio of the magnitude differences is very poor, many stars cannot be measured at all, and things are not changing very fast anyway. Finally, for the faintest stars the signal-to-noise ratio is maximized in small apertures and is a strong function of aperture radius. When the smallest apertures are few and far between, the freedom to choose the very best aperture for such a star is diminished. Even the scheme of aperture radii following a geometric progression (as in Figs. 1–5), which I have been using for several years, tends to put the apertures too far apart near the center and too close together near the edge. A scheme more like that which I used for the data of the fourth night as mentioned above seems to be preferable. Arbitrary patterns of apertures are difficult to interpret with traditional hand-and-eye techniques, because it is hard to know a priori what the shape of the hand-drawn curves should be. DAOGROW, with its precise knowledge of the actual aperture radii and its (I. R. King's) prior experience with the typical morphology of a stellar profile, has no problem with an arbitrary progression of apertures.

In Section 3 above I mentioned that I found the parameters D and E in the formulation of the analytic model star profile to be unimportant. To illustrate these points I rereduced the data for the three photometric

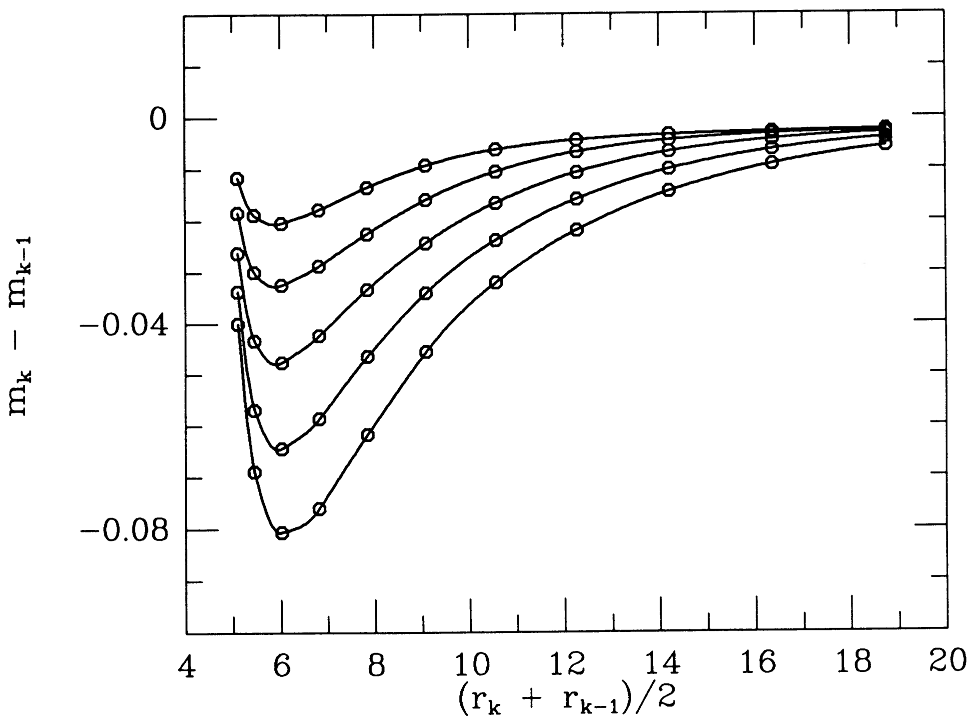

FIG. 6–Five selected analytic model growth curves from one night in 1988 at the prime focus of the Cerro Tololo 4-m telescope. This night was characterized by seeing which was much poorer and more highly variable than the preceding three nights (cf. Figs. 1 and 2). Shown are the model curves for the individual frames with the best and poorest seeing of the night and curves for three intermediate seeing values. The points represent the model-predicted differences between aperture magnitudes from consecutive concentric apertures, plotted as a function of the mean radius of the two apertures. Solid lines connect the magnitude differences corresponding to each individual growth curve. The morphology of these curves differs from that in Figs. 1–5 because a different rule was used to set the spacing between consecutive apertures.

nights with the parameter D fixed at 0.9, once more with D free and E fixed at 0.0, and once again with both D and E fixed. The weighted standard deviations of the residuals of the raw observations about the best solutions, averaged over all stars and all apertures, were: five parameters free, 0.002639 mag; D fixed, 0.002608 mag; E fixed, 0.002638 mag; D and E fixed, 0.002604 mag. In addition, when data for various observing runs were reduced with E completely free, it seemed as likely to come out negative as positive, conflicting with the common-sense expectation that atmospheric scattering should increase toward higher airmasses. Accordingly, I conclude that the formation of the aureole is probably dominated by the optics, not the atmosphere, and the airmass dependence can probably be neglected.

One might also suspect that it would be worthwhile to compute growth curves separately for different photometric bandpasses, since the strength of the aureole and possibly also the form of the seeing profile could depend on wavelength and bandwidth. During the observing run at Cerro Tololo I obtained data only in the B and V photometric bandpasses. Reducing these data sets separately, I obtain standard deviations of 0.002546 mag (V) and 0.002603 mag (B). The root-mean-square average of these is 0.002575 mag, which should be compared to the value of 0.002639 obtained when the data were reduced in combination. The 2.5% reduction in the scatter seen when the data are reduced separately corresponds to about a 5% increase in weight over the combined solution, or the equivalent of 21 observations for the price of 20. This suggests that the advantage to be gained by computing separate families of growth curves for the two filters is not large. However, this experiment should certainly be repeated, especially with data for more widely separated bandpasses.

I am grateful to Michael Bolte for letting me play with some of his data and for reading a preliminary draft of this paper.

REFERENCES

Buonanno, R., and Iannicola, G. 1989, *Pub. A.S.P.*, **101**, 294.
Buonanno, R., Buscema, G., Corsi, C. E., Iannicola, G., and Fusi Pecci, F. 1983, *Astr. Ap. Suppl.*, **51**, 83.
Da Costa, G. S., Ortolani, S., and Mould, J. 1982, *Ap. J.*, **257**, 633.
Hardie, R. H. 1962, in *Stars and Stellar Systems*, Vol. **2**, *Astronomical Techniques*, ed. W. A. Hiltner (Chicago: University of Chicago Press), p. 178.
Harris, W. E., FitzGerald, M. P., and Reed, B. C. 1981, *Pub. A.S.P.*, **93**, 507.

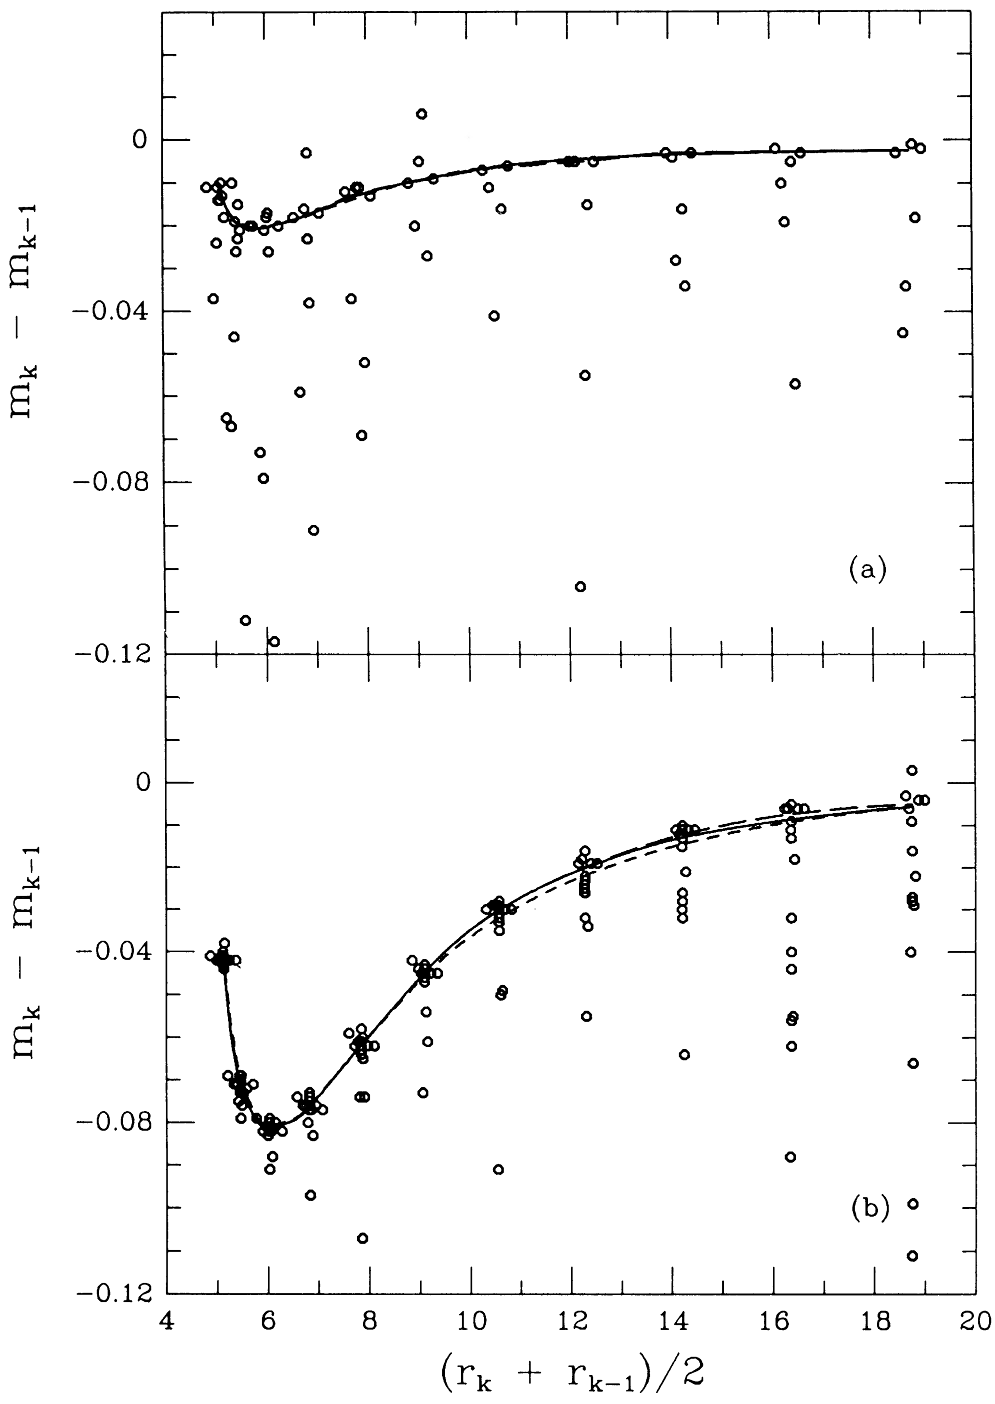

FIG. 7–(a) Observed magnitude differences and growth curves for the frame with the best seeing from the fourth night of my 1988 CTIO observing run. Points represent the magnitude differences actually observed for stars in the frame, plotted as a function of the mean radius of the two apertures involved in each difference. The best-fitting analytic growth curve is shown as a short-dashed line; the empirical growth curve derived from the raw observations is shown as a long-dashed line; and the adopted growth curve, which represents a compromise between the empirical growth curve (most reliable at small radii) and the analytic model curve (most reliable at large radii), is shown as a solid line. The three curves are virtually indistinguishable at this resolution. (b) Same as the upper panel but for the frame with the poorest seeing of the night. The individual stars are comparatively bright, well isolated, and well photometered. Since in this frame the data are precise and reasonably concordant, the adopted curve adheres to the empirical growth curve out to a large radius, making a transition to the analytic curve only around $r \sim 13$–16 pixels = 8″–10″.

Hesser, J. E., Harris, W. E., VandenBerg, D. A., Allwright, J. W. B., Shott, P., and Stetson, P. B. 1987, *Pub. A.S.P.*, **99**, 739.
Howell, S. B. 1989, *Pub. A.S.P.*, **101**, 616.
King, I. R. 1971, *Pub. A.S.P.*, **83**, 199.
Landolt, A. U. 1973, *A.J.*, **78**, 959.
Lupton, R. H., and Gunn, J. E. 1986, *A.J.*, **91**, 317.
Mateo, M., and Schechter, P. 1989, in *Proceedings of the First ESO ST-ECF Data Analysis Workshop,* ed. P. J. Grosbøl, F. Murtagh, and R. H. Warmels (Garching: ESO) (ESO Conference Workshop Proceedings No. 31).
McClure, R. D., VandenBerg, D. A., Bell, R. A., Hesser, J. E., and Stetson, P. B. 1987, *A.J.*, **93**, 1144.
Penny, A. J., and Dickens, R. J. 1986, *M.N.R.A.S.*, **220**, 845.
Rich, R. M., Da Costa, G. S., and Mould, J. R. 1984, *Ap. J.*, **286**, 517.
Smith, G. H., McClure, R. D., Stetson, P. B., Hesser, J. E., and Bell, R. A. 1986, *A.J.*, **91**, 842.
Stetson, P. B. 1987, *Pub. A.S.P.*, **99**, 191.
Stetson, P. B., and Harris, W. E. 1988, *A.J.*, **96**, 909.
Stetson, P. B., VandenBerg, D. A., and McClure, R. D. 1985, *Pub. A.S.P.*, **97**, 908.