Data Modeling In Scientific Images Using Simulated Annealing

J.R. Parker
Department of Computer Science
University of Calgary
2500 University Drive NW
Calgary, Alberta, CANADA

G. Groisman
Facultad de Econamia y Administracion
Universidad Nacional del Comahue
Argentina

Abstract

One way to measure objects in some classes of scientific image is to model the objects by functions and then make the measurements on the functions. This works especially well for images that contain information in the form of relationships between grey level pixels. Here, the use of a Moffat function as a data model is explored, as is the use of simulated annealing to fit many instances of this function to the data in the image. Two examples are presented: stellar photometry, a natural application for the Moffat function, and reading DNA sequencing gels.

1. Introduction

Scientific images can be broadly classified into two groups: images in which position, lines, and shape are the key elements (EG graphs, chart recordings) which can be referred to as **line images,** and images in which grey levels or colors and their relationships are the most important items (EG astronomical images, DNA gels) which will be referred to as **blot images.** Since the information content of a blot image is contained in the grey level values and relationships of the pixels in the image, it is generally true that mathematical functions can be used to model the objects in such an image. One specific example of a blot image is the astronomical stellar image, in which bright spots appear on a non-uniform 'dark' background (although astronomers often use the negative image). Researchers are often interested in the brightness and color of stars, especially when studying star clusters in general and globular clusters in particular. A problem is that stars are often too close together to obtain an accurate measurement of brightness. Moreover, there may be many thousands of stars in an image, a fact that makes computer assisted analysis very important.

Another type of blot image results from chromatography, in which chemical compounds can be separated from each other by diffusing them through some medium (paper, gas, or gels). A DNA sequencing gel is a specific example of this process, and a gel image has some properties similar to those of astronomical images. A problem with blot images is that the source of the image often dictates the structure of the data to be extracted. For example, with DNA gels the brightness of the bars is not that important, since it is the location of the bars that is used to sequence the bases. On the other hand, when doing stellar photometry the positions of the stars is not relevant.

The work appearing in this paper involves modeling the data objects in a blot image, so that the relevant scientific information may be extracted more easily. Mathematical functions, different for each image type, will be fit to the image objects, and the properties of the objects will be inferred from the parameters of the functions. For example, measuring the brightness of a stars is equivalent to fitting a function to the data and measuring the height of the fitted function. When star images overlap, many stars would have to be fitted simultaneously. Optimizing this fit requires of a method able to explore the variable space involved and not get caught in local minima, since the best result is that corresponding to the global minimum. Because of this normal downhill methods may not yield the best results. It is proposed here to apply the techniques of *simulated annealing*, a biased random walk method, to the optimization of this fit. Simulated annealing mimics the way a slowly cooling metal reaches its optimum crystal structure, by employing the same decision making mechanism dictated by quantum mechanics.

The two major examples to be explored are crowded field stellar photometry and automatic reading of DNA sequencing gels. Section 2 of this paper describes the simulated annealing method in detail and outlines methods for improving performance in this type of data fitting application. Section 3 describes the problems encountered in computer assisted photometry of crowded star fields and shows how the simulated annealing method has been used for this application. Section 4 shows another example, that of decoding DNA sequencing gels, and again shows how simulated annealing can be used to fit a function to the gel bars. Finally, Section 5 summarizes the work and gives some hints about future efforts.

2. Simulated Annealing

Many methods have been devised to optimize the value of a function in one of more parameters. These methods usually employ a figure of merit that determines how good the optimization is; then this value is optimized by changing the parameters repeatedly until the ’best’ result is achieved. The most straight forward approach is to choose new values of the parameters by changing them in the direction that reduces the value of the figure of merit. Though this would work fine for functions with a single minimum, it has the unpleasant property of trapping the optimization process in local minima. To improve the chances of locating the one global minimum these methods are usually run several times from several different starting points and the best result obtained is taken as the global minimum.

In 1953, Metropolis [1] proposed a different approach to the optimization problem. The idea was to observe how nature optimizes the placement of atoms in a lattice when a metal sample is slowly cooled, a method known as *annealing*. When this happens the atoms find the lowest energy configuration possible, yielding a very regular structure. On the other hand, if the sample is cooled quickly (quenching) then the result is a very irregular structure far from the optimum energy state possible.

To appreciate the role the temperature plays it is necessary to think of the process in action, as described by quantum mechanics. At any given temperature T the probability of a system of having energy E is given by

$$Prob(E) / e^{\frac{-E}{kT}}$$

where k is the Boltzmann constant: this is the so-called Boltzmann probability distribution, and shows that the probability of the system jumping into a state of higher energy is finite. The system sometimes goes uphill, unlike the steepest descent methods, which allows the system to climb out of a local minimum. As the temperature gets smaller the chance of an uphill climb diminishes as well.

Simulated annealing can be seen as an algorithm that constantly tries to transform the current configuration of parameter values randomly into a new one whose value are near the previous one. This process can be simulated easily with a computer algorithm. First, a way is needed to express the probability P of the system changing energy by an amount ΔE defined as $E_2 - E_1$, the difference between the new energy value and the old one. This change is negative for a downhill change (E_2, the new energy is smaller than the old one) and positive otherwise. The probability of accepting the change is now defined as:

$$P(\Delta E) = \begin{cases} 1 & \text{if } \Delta E \leq 0 \\ e^{-\Delta E / kT} & \text{if } \Delta E > 0 \end{cases}$$

This is because it is always desireable to take downhill steps ($\Delta E \leq 0$) but the uphill ones are only accepted based on the Boltzmann probability distribution. This is known as the Metropolis criterion. The process of deciding based on this probability can be simulated in a computer [3], by comparing the value of a random number between 0 and 1 with the value of $P(\Delta E)$.

$$Random \leq P(\Delta E) \rightarrow accept$$

$$Random > P(\Delta E) \rightarrow reject$$

Finally a decision has to be made about how the temperature will be lowered: this is known as the *annealing schedule* [4,5,6]. The schedule can be as simple a process as reducing the temperature by a constant fraction every so many parameter adjustments or as complex as relating the temperature to the current error values. The choice of annealing schedule is crucial to acceptable performance of the algorithm.

The figure of merit often employed to determine the goodness of the fit is the χ^2 value, which gets smaller as the fit improves. What remains to be decided are the values of k and T. As we shall see later these are related to the function being optimized and can easily be decided based on information available during the fitting process.

An example, introduced by Bohachevsky [3], can be used to illustrate these concepts. The function to be minimized is:

$$f(x,y) = x^2 + 2y^2 - 0.3cos(3\pi x) - 0.4cos(4\pi y) + 0.7$$

As can be seen in Figure 1 this function has many local minima and one well defined global minimum at (0,0). Any downhill method could get caught in one of these local valleys and never reach (0,0). To test simulated annealing we start at (1,1) and let the algorithm proceed.

To illustrate the main points of simulated annealing a very simple implementation is used. The function to be fitted is known, and to detect the global minimum (or any minimum for that matter) the step size in x and y should be smaller that about 0.2 since this is about the size of each minimum. The program then begins at (1,1) and takes a step in the x direction of size $Rand\ step_x$, where *Rand* is a random number between 0 and 1, and $step_x$ is the maximum step size allowed in the x direction. It then computes the value of the function at this new point and applies the Metropolis criterion to decide whether to move there or not. Then next step is taken in the y direction in the same way, and the whole process is repeated until the global minimum is located. Knowing when this happens is also made simpler by previous knowledge of the function being optimized.

Figure 1b shows the path followed by the algorithm, where the points indicate the steps that the algorithm followed. For this simple test the temperature was lowered by 0.95 every 30 steps and the value of k was set to 1. The plot clearly illustrates why simulated annealing is considered a biased random walk.

For every pixel in the region being fitted it is necessary to locate all those objects near enough to

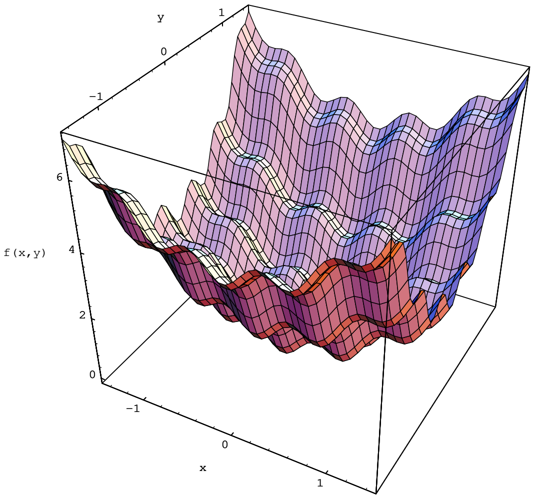

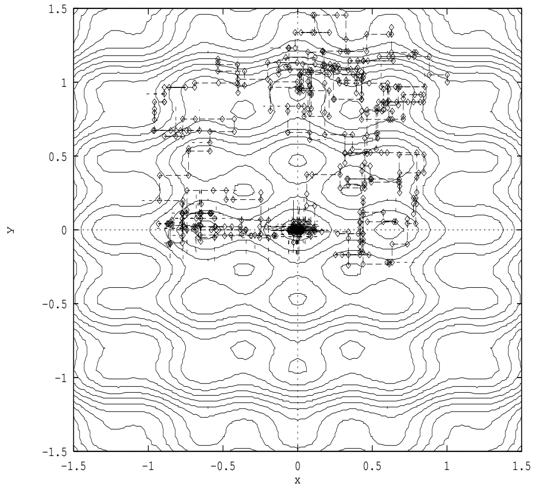

Figure 1 - Simulated annealing test result. (a) Plot of the function.
(b) Steps taken by the simulated annealing method while locating the minimum.

contribute some overlap. This is called *grouping*, and is done from a list of all located objects and their positions. One object is selected and then those which lie close enough are located and removed from the list. The process is recursively repeated for all these objects until no other overlapping object can be found. At this point a new object is selected from the list and the whole process repeated until no more objects remain in the list. The resulting grouping can be seen in Figure 2 for a typical star image. These groups of objects must have functions fit to them all simultaneously.

As we have mentioned, the χ^2 value for the fit can be used as a measure of how well the functions fit this particular group; in this case it is clear that the problem is to minimize χ^2. A better idea is to normalize the χ^2 function by dividing by the expected pixel uncertainty, giving a value that approaches one (1) as the fit improves. This is known as the *reduced* χ^2, which will be denoted by χ^2_υ. If the final value of the χ^2_υ is much larger than 1 we can conclude that either the fit is bad or the noise was underestimated. If it is much smaller than 1 then the fit is too good to be true, so the errors were overestimated.

An improved algorithm for simulated annealing optimization consists of a repeated series of adjustments to the parameters of the functions being fit; there would normally be one adjustment per parameter being fit (not all parameters will always be involved in this process). After a fixed number of cycles have been performed the algorithm consults a record of how many adjustments were accepted for every parameter. The intention here is to keep the ratio of acceptance to rejections to about 1/2. To enforce this the

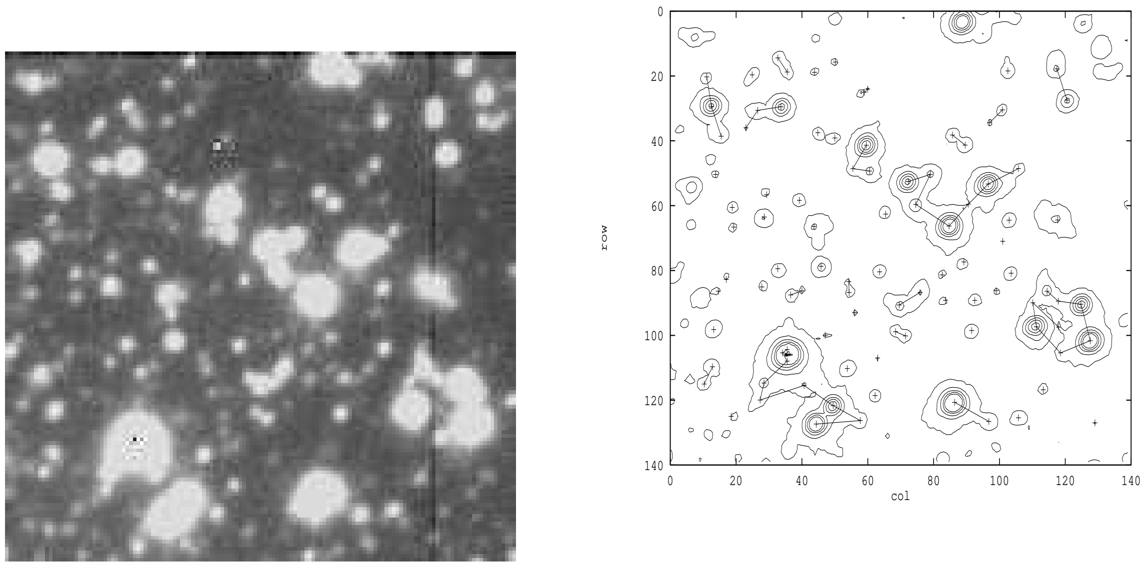

Figure 2 - Groups of objects, taken from a star image. (a) Original image.
(b) Object groups, that must be fit simultaneously.

step sizes are adjusted and the whole process is repeated. Several sets of step adjustments are performed, after which the temperature is reduced. At this point the whole process is restarted, but with the parameters set to the optimum value found in the previous iterations. This as a whole is repeated until the stopping criterion is achieved.

It is necessary to establish a way of deciding when the process must be stopped. For real data a fit is rarely exact and hence one cannot expect the value of χ^2_υ to reach 1. A more reasonable way of detecting a solution is to look at how the values of χ^2_υ progress. When the fractional change between the previous and the current values and the fractional change between the optimum and the current value both fall below a certain value, denoted by ε, it can be said that the algorithm has converged on a solution.

3. Simulated Annealing in Stellar Photometry

One common activity in astronomy is to measure the brightness of stars [8,9]. If the image of a star as seen through a telescope were just a dot of light, one would be tempted to just measure one pixel in our image and consider that to be the brightness of the star. Unfortunately, the light has been spread over a region of the image according to the transfer function of the telescope. Instead of just measuring one pixel what should be done is to integrate all this light back into one value. This is equivalent to undoing the spreading of the transfer function.

If no other stars are present a simple way of doing this integration is simply to select a region of the image and add up the light in all those pixels. This is known as *synthetic aperture* photometry [10] because the region simulates how the aperture of a single detector would have measured the signal. A second approach to photometry uses the transfer function of the device. Here we use a function to account for the fact that the star was a point source that was spread into the resulting image; it is common to refer to this function as the point spread function (PSF). If we can come up with an analytic expression for the PSF then given the image it should be possible to adjust this expression until the function best fits the data. We can now compute the brightness of the star by integrating this function.

When deciding on a function to approximate the PSF of a telescope the best choice is fairly obvious. Moffat [11] studied the way in which an image forms in a telescope, and he showed that a simple function with 2 free parameters was enough to account for most of the distortions. The function suggested has come to be known as the Moffat function and is of the form

$$I(r) = I_0(1 + (\frac{r}{\rho})^2)^{-\beta}$$

where I_0 is the intensity of the star, $I(r)$ is the intensity observed at a distance r from the center of the image and ρ and β are shape parameters. If the PSF is space invariant then so are the shape parameters. For example, given a value for peak intensity and full width at half maximum, only one possible shape exists for both the Gaussian and Lorentzian curves. For the Moffat curve there is an infinite combination of values of ρ and β that still give the same full width at half maximum, as can be seen if Figure 3. It is this flexibility that makes this function a more suitable choice.

Of course, to achieve a high degree of accuracy all noise sources must be identified and their values computed. Two such sources have been accounted for in the simulated annealing photometry system: the background, or *sky*, and the usual uncertainty associated with a counting process, which will be a Poisson noise. Using these, the uncertainty of any one pixel can be expressed as the sum of two sources of error:

$$\sigma_{pixel}^2 = \sigma_{sky}^2 + \frac{DN}{g}$$

where σ_{sky} represents the measured standard deviation of the background, and DN/g represents the Poisson noise measured in DN units. The g represents the gain of the CCD used to capture the image.

The photometry system was tested in two ways: first on an extensive suite of simulated images, and

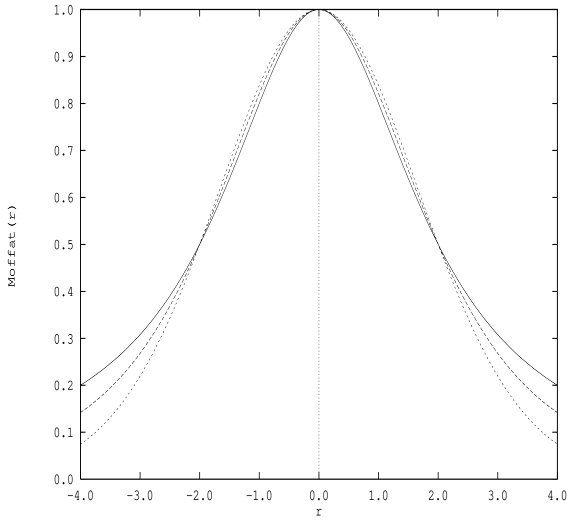

Figure 3 - Moffat's function

then on real images for which good photometry has been published. Testing on synthetic images was done to explore the ability of the algorithm to solve the different problems crowded-field photometry imposes in a controlled environment. Several sets of test data were created and different types of noise added to evaluate performance under conditions that went from ideal to realistic.

In a test image a star is represented by a two dimensional peak. A pair of x and y coordinates determine the center of the star, and a value of I in DN its brightness. The PSF of the image is controlled through the shape of the star function. Two types of functions were used: Gaussians with σ=1 and Moffat's with ρ=2.0 and β=3.0. The images were subjected to three types of noise, and in all cases the stars were located and fit to within a small error. For example, test stars at (3,3) and (6.0,6.3) having peak intensities of 2000.0 and 1000.0 respectively were fit with functions centered at (2.98,3.01) and (6.0,6.2) with peaks of 2065 and 987 respectively.

Another type of test uses a sequence of two star images where the two stars are moved progressively closer to each other. They correspond to a separation of 3, 2.8, 2.5, 2.1, 1.8 and 1.4 pixels between the centers. Only noise of the Poisson type was tried with these images. As seen before, if the method performs well under this type of noise it can be expected to do as well or better under less noisy situations.

The parameters used to create these blended images are:

Image	I	X	Y	Image	I	X	Y
1	2000	3.0	3.0	4	2000	3.0	3.0
	1000	6.0	6.0		1000	5.1	5.1
2	2000	3.0	3.0	5	2000	3.0	3.0
	1000	5.8	5.8		1000	4.8	4.8
3	2000	3.0	3.0	6	2000	3.0	3.0
	1000	5.5	5.5		1000	4.4	4.4

Even at a distance of only 1.4 pixels, annealing does a good job of estimating both the center coordinates and intensities of the stars:

Table 2: Fits to pairs of Moffat stars at varying distances + Poisson noise

Image	χ	ρ	B	x	y	I
1	1.3402	2.01(2)	3.11(7)	3.00(4)	3.01(8)	2062.(105)
				6.02(5)	5.9(1)	1022.(75)
2	1.2188	2.02(3)	3.07(3)	3.00(5)	3.0(1)	2051.(87)
				5.8(1)	5.75(6)	972.(68)
3	1.1473	2.03(1)	3.14(5)	3.00(3)	3.02(3)	2016.(53)
				5.50(8)	5.46(2)	1029.(88)
4	1.3312	2.23(2)	3.41(7)	2.9(1)	2.99(5)	1928.(49)
				5.09(7)	5.08(9)	920.(59)
5	1.3333	2.19(4)	3.38(3)	2.97(2)	3.00(7)	1978.(129)
				4.84(9)	4.7(2)	941.(89)
6	1.2627	2.13(3)	3.26(5)	2.98(4)	2.96(2)	1950.(119)
				4.33(6)	4.37(4)	1058.(47)

Tests were also performed to verify the ability to fit stars by iterating when a star is missed in a first pass, and later added to the fitting list. For example, an image was created with three stars in it. Two were large and separated enough to make sure they would be detected in the first pass, and a third smaller star was placed between these two so it would not be possible to see it in a first pass. The parameters for this image can be seen in Figure 4, which also shows both the image and a three dimensional representation of the data. The original image clearly seems to contain only two stars. When only two stars are fitted the median filtered residuals look like those in Figure 4b. Based on this a third star is added and all three fitted again. The resulting median filtered residuals can be seen in Figure 4c. Notice that in the first attempt the lack of the third star pushes the centers of the two fitted stars closer, to compensate for the missing star in

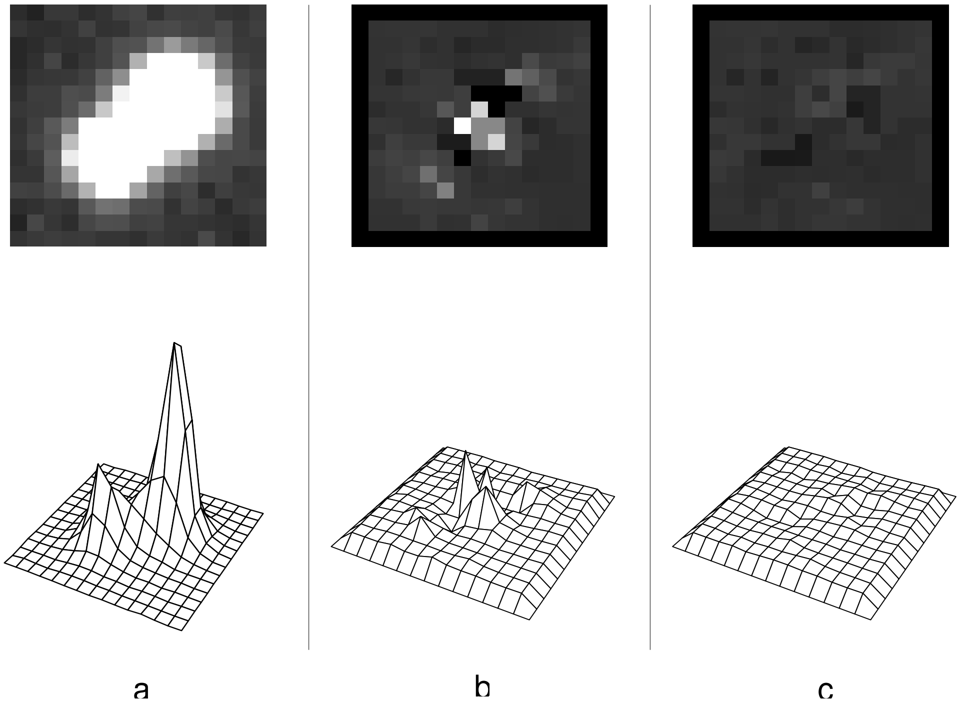

Figure 4 - 3 star residuals. (a) The 3 star image (b) fit of two stars to this image and residuals (c) Fit of three stars and residuals.

the center. Their brightness are also larger than expected to account for the extra DN. The residuals from the fit clearly show the poorly fitted data where that star should be. When a third star is added all parameters converge to proper values.

While synthetic test images provide control situations were the performance of the algorithm can be tested, the final test is the application to real data. For this purpose several CCD images of the globular cluster NGC6397, and a number of others, were obtained and reduced.

In astronomical work the independent values obtained for each star in the frame are rarely considered individually. Usually two frames are available for reduction, one taken with a V filter (meaning visual, a yellowish color) and another one with a B (blue) filter. Once the stars are measured in both frames it is possible to establish their color from the two sets of results. A normal way of describing the color is by expressing the difference between these values, for example B-V.

From stellar evolution theory it is known that there is a strict relationship between the age, brightness and color of any star. In the case of a globular cluster all stars are expected to have very similar ages, so their magnitudes and colors are strictly correlated. A plot of these two values is called a Color-

Magnitude diagram. A pair of B and V observations of the cluster NGC6397 were reduced using the methods presented here.

First the frames were reduced selecting a few calibration stars and were fitted with $\varepsilon = 0.1$. This meant a fast reduction, but the quality could be expected to be marginal. The resulting diagram can be seen in Figure 5(a). Though the main sequence is clearly visible, the spread of the data points is too large. This is mostly due to a large value of ε which translates into the annealing process stopping too soon. A second run was performed with a value $\varepsilon= 0.05$. this increased the execution time significantly (about a factor of 4) but gave much better results as can be seen in Figure 9(b). The more strict parameters forced the fitting process to continue until a better fit (hence a lower value for χ^2_υ) was found.

Since the selection of calibration stars can affect the results it was decided to go back to the beginning and select new calibration more carefully. These were then fitted for all parameters, shape parameters included. The resulting fits were examined and averages for ρ and β computed. New fits were performed with these new values and the smaller value of ε. The results can be seen in Figure 5(c).

Data of a previously published C-M diagram for NGC6397 was also obtained from papers by Alcaino[12,13,14,15]. This made it possible to compare the final C-M diagram with published data. Alcaino's photometry had been calibrated by observing some stars of known brightness. To make both diagrams match several stars were identified in Alcaino's work and in the present work and the correlation computed. The resulting values and fit can be seen in Figure 6(a). Using the resulting fit the measured magnitudes were converted to the same scale Alcaino used. The two sets of Color-Magnitudes are shown in Figures 6(b) and (c). Notice that the published work was able to cover a much wider range of

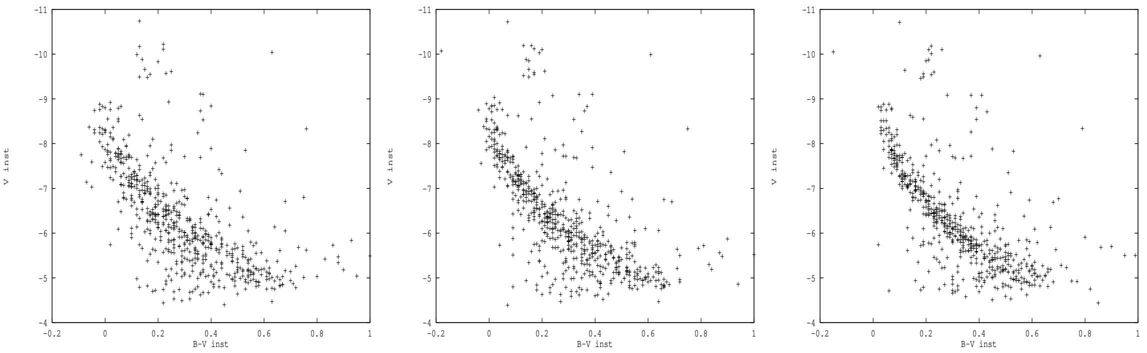

Figure 5 - (a)Color Magnitude diagram of NGC6397, ε=0.1, (b) C-M diagram with ε=0.05, (c) C-M diagram with new calibration stars.

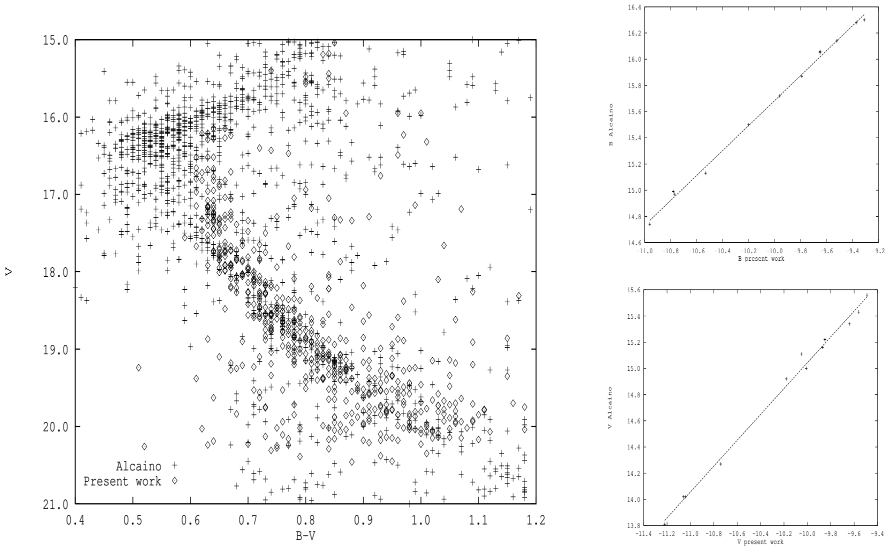

Figure 6 - (a) Both C-M diagrams, Alcaino[12] and present work.
(b,c) Calibration of our observations of NGC 6397.

magnitudes by combining several types of observations, both photographic and with CCDs. The results from the present work on the other hand come from a single pair of images and are much more restricted in range. The overlap region shows how the agreement between the two methods is remarkable.

4. DNA Sequencing Gels

DNA sequencing involves the determination of the structure of a segment of a DNA molecule in terms of the order of occurrence of the *bases* (guanine = G, adenine = A, thymine = T or cytosine = C) of which it is composed. The base sequence codes for the construction of a protein, and all proteins in a cell have a code that defines them; thus it is not uncommon for a molecular biologist to be interested in decoding DNA sequences.

These sequences are determined using the technique of electrophoresis, in which the bases are marked with radioactive ’tags’ and separated electrically on a thin agarose gel. When a photographic film is placed on this gel and exposed for a few days an image is produced in which radioactive regions correspond to dark bars on a lighter background. Each band represents a base in the original DNA sample; the position of the base in the DNA is a function of the distance of the dark band from the top of the image. Each base is restricted to a fixed set of columns on the image, one column for each base type, as seen in

Figure 7.

Reading these gels involves locating the bars from the top down and recording the bases in their order of occurrence. However, the bars interfere with each other in places, much like the problem seen in crowded star fields. The use of functions to model the bars and the use of simulated annealing to do the fit seems an obvious enough choice given the success of these methods on the similar photometry problem.

The problem will be reduced to a one dimensional fit by extracting data along the center of one of the columns. This gives a one dimensional curve with dips at the location of each dark band. These dips shall be referred to as inverted peaks, or just peaks from here on. This curve will be fit with a combination of a flat sloping background

$$A + Bx$$

where A is the average value of the background and B the slope; and a set of Moffat functions in one dimension, given by

$$I_o = \left(1 + \left(\frac{x - x_o}{\rho}\right)^2\right)^{-\beta}$$

where I_o is the height of a peak located at x_o. Since the bars are darker, the values of I_o can be expected to be negative. The values of ρ and β (the shape parameters) will be kept the same for all peaks in one column. The full model will contain n bars, and the error will be set to one for all data points.

In the case of the DNA gels it is easy to estimate the background. The empty space between the columns can be used to calibrate the flat background fit knowing that contamination there is minimal. The location of the bands can then be estimated by searching for the peaks in the curve. Once all the larger peaks have been identified the column is fitted. There is no need to fit groups separately as the number of bars present is not very large, and a one dimensional fit requires less computation.

First the full set is run and the program allowed to adjust only the position and depth of the peaks, the background and the shape parameters of the Moffat function are held constant. A second run is made later with the position and intensities obtained from here, but now all parameters are adjusted.

To explore the problems involved two sets of columns were reduced. The first contained clearly differentiable bands that nevertheless overlapped to some degree. The second was a set with regions of highly blended bands. Both can be seen in Figure 7. The left side of Figure 7 shows the set that was first reduced. The central 5 columns of data along each bar was averaged for fitting. From these averages the

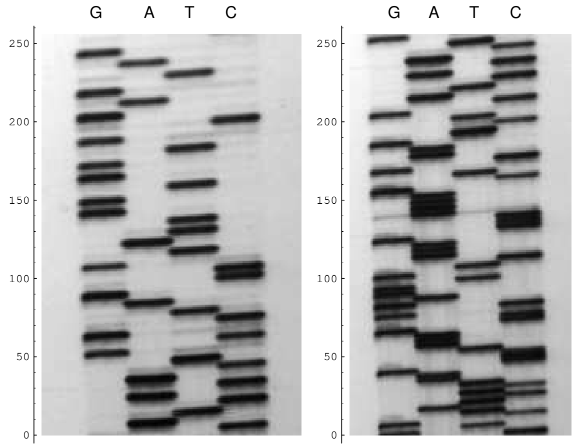

Figure 7 - Test sets of DNA sequencing gels, reduced. The left is more sparse with well defined bands and little overlap. The right one shows more crowding of bands in certain regions.

location of the center of the peaks were estimated. With this estimate a first run was made, The purpose of which was only to refine the positions of the peaks. Figure 8 shows this process at three different stages.

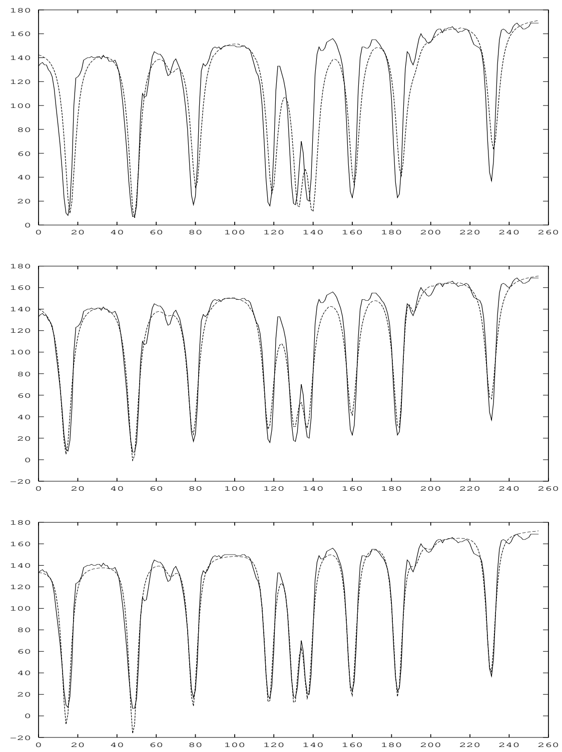

Figure 8 - Iterative refinement to the fits of column T. The original approximation can be seen the top, better values for the positions are obtained after the first run as seen in the middle. Finally shapes and background are also fit in a second run as shown in the bottom plot.

The top plot shows the data being fitted in solid line, and the fit in dashed line. The middle plot shows the fit after the first run of annealing was completed. Only the positions and depths of the peaks are fitted in this pass. The location of the peaks has been improved significantly. The shape of the peaks can be improved since it can be seen that the fit fails to model properly the full shape of the data. The background is also likely to be wrong at this stage.

To fix this a new annealing run is performed. The new positions and depths of the peaks are used as starting values, and this time all parameters are adjusted. By the end of this second round the fit has improved significantly, as can be seen in the bottom plot of Figure 8. In looking at the final fit, we note that the position of each band can be computed to an accuracy in the order of half a pixel. The same process was followed with all four columns, and the resulting fits can be seen in Figure 9. After all bands were fitted, their positions had to be rotated to allow for the correct sequencing. The angle of the columns to the edge of the image was measured to be 1.8 degrees, so all four columns had to be rotated as a whole by that amount. Once this was done it was possible to sequence all four columns by sorting the peak positions. It can be said that no conflict arises when the position of any two adjacent bars is outside each other's uncertainty. This is the case for all bands in this gel where the average uncertainty in position is in the order of 1/10 of a pixel.

The resulting sequence can now be compared to accepted values. For the section presented above:

```
CATCACACTGCGCTAGCCGTATTGGTGGTGCGAGTAG
catcacactgcgctagccgtattggtggtgcgagtag
```

where the top row represents the values found and the one on the bottom the accepted ones. Agreement is complete. The next attempt was made with a more crowded gel. Here the overlap is more serious which makes the detection of the bars more difficult. Based on the apparent width of the bands a guess was made about their positions. These values were then used as the starting guesses for a reduction sequence as explained above. The resulting fits can be seen in Figure 10. Though the closest two bands are less than a pixel away, the uncertainty for these bands is in the order of less than 1/10 of a pixel.

The resulting sequence can now be compared to accepted values:

```
GTGCTATCTCT AGCC TAA GCCGCGAGGTGTCAAAGCCCAAAAGCTGCAAGTCTGCATCACACTG
```

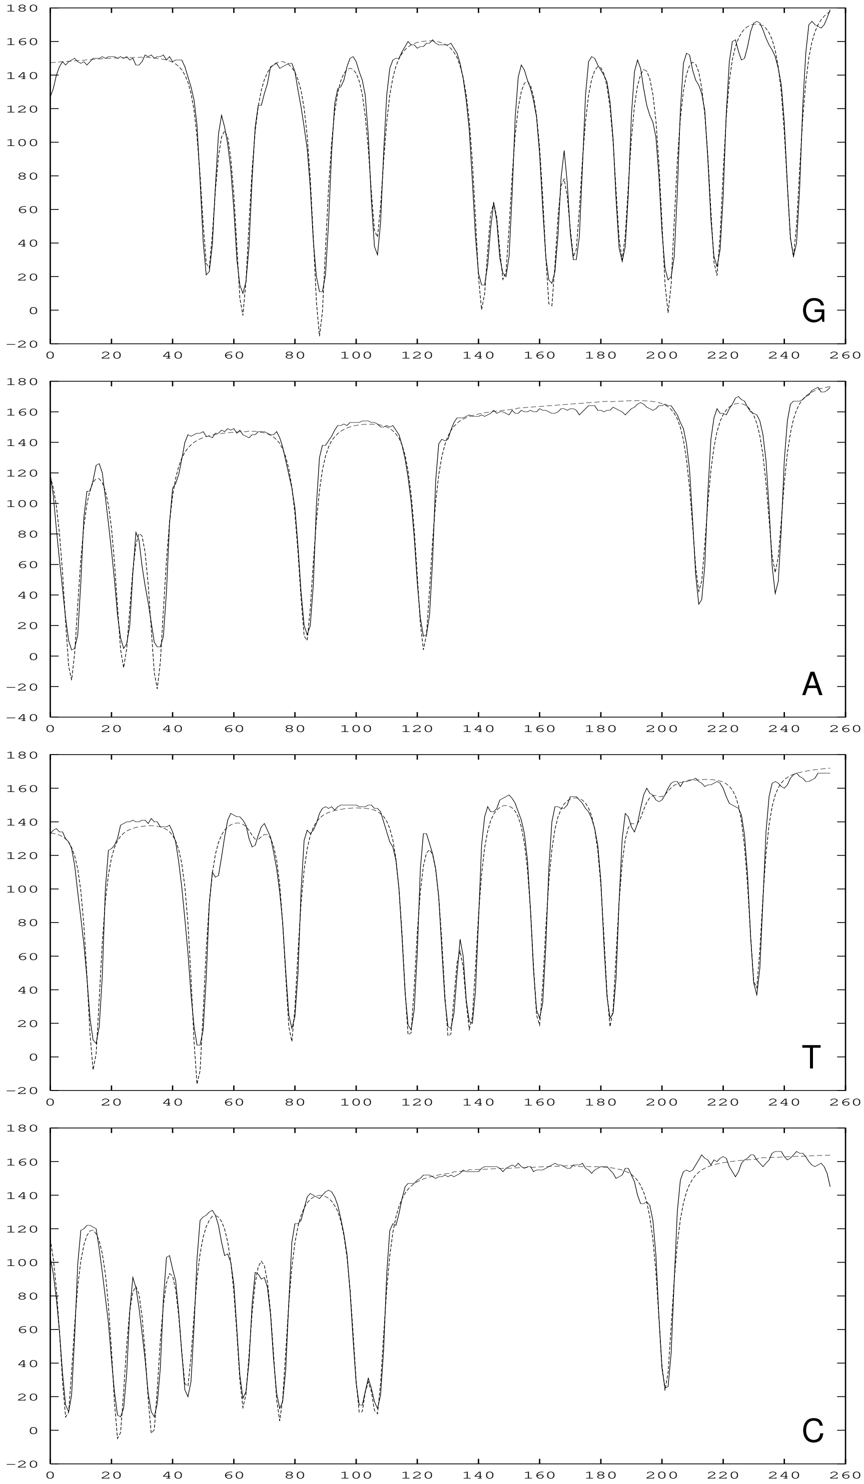

Figure 9 - Fits to the sparse data set.

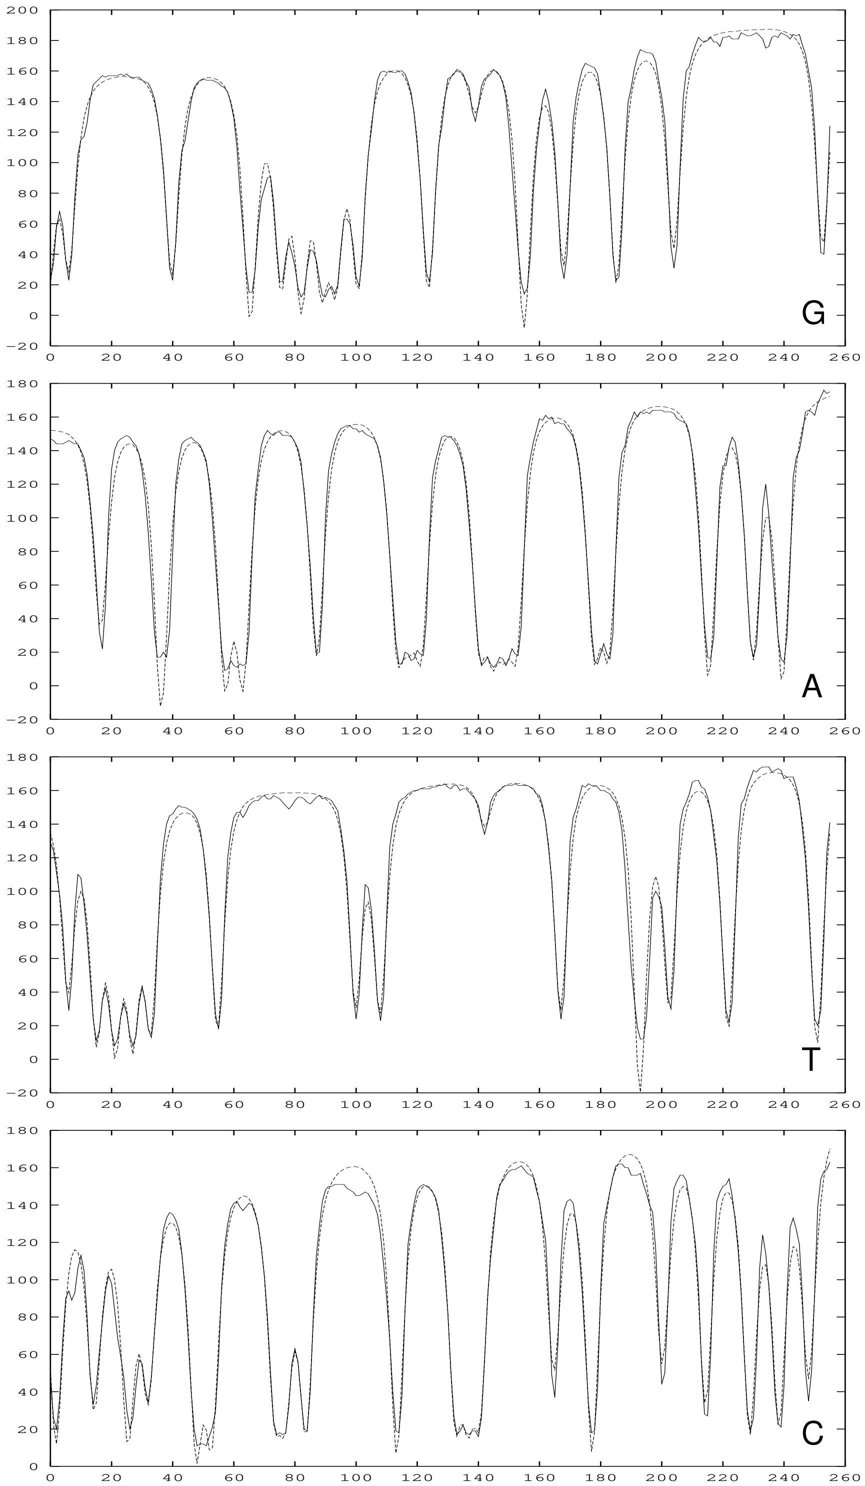

Figure 10 - Fits to all 4 columns of the crowded set.

gtgctatctctaagccctaaagccgcgaggtgtcaaagcccaaaagctgcaagtctgcatcacac

Here a couple of gaps are seen where the present method failed to locate a band. In these cases only a single data point suggests the presence of two peaks instead of one, and it could be just due to noise. Using the width of the peak (the single bar fit is narrower than the data) the existence of another bar can be inferred, and when inserted the fit improves, confirming our suspicions. Still, the fit is poor in some of these regions; it would seem the fitting process cannot quite agree whether there should be two blended peaks in this region.

The problem is mainly in the nature of the data being reduced. To make visual inspection easier the image of the gel was exposed in the order of days onto photographic film. This saturates the bands making them too dark. As can be seen in any of the plots, data inside the bands is on the order of tens of DN. When doing a functional fit this saturation hinders the results. On top of this, the non-linear nature of photographic film makes this saturated data even more difficult to model. Shorter exposures would produce fainter bands that the human eye might have trouble discerning. We are also rescanning the data at higher resolution, which should improve the fits significantly. For the present type of data it is believed that the crowding seen in some regions of the set is more than can be reliably reduced at the current resolution.

5. Conclusions and Further Work

It can be concluded that simulated annealing offers a viable way of optimizing function fits to data, in particular as it applies to astronomical photometry.

The main drawback remains speed. By its own nature simulated annealing is slow, randomly exploring regions of parameter space that other methods may never visit. But it is also here were its strength lies, since in doing so it is able to come up with results other methods may miss. So it is not a matter of rejecting it based on the longer execution time but rather it should be judged on the need for better results it can provide.

Speed, however, may still be increased significantly. In particular new methods have been reported [16] that make use of the shape of the χ^2_υ surface to guide the random walk. This is achieved by combining simulated annealing with simplex, which has the ability of choosing to move in a way consistent with the surface. We are also experimenting with a method for distributing the annealing problem across a network of workstations, which will decrease the real time needed to solve one problem by a factor of 20.

More complete data and tabular results are available for electronic transfer.

This work has been supported by the National Sciences and Engineering research Council of Canada.

References

[1] Metropolis, N. et al, **Equation of state calculation by fast computing machines,** *The Journal of Chemical Physics,* Vol 21, No. 6, 1953. Pp 1087-1092, 1953.

[2] van Laarhoven, P.J.M., and Aarts, E.H.L., **Simulated annealing: Theory and applications,** *Kluwer academic publishers,* 1989.

[3] Bohachevsky, I.O., Johnson, M.E. and Stein, M.L., **Generalized simulated annealing for function optimization,** *Technometrics,* Vol. 28, Pp. 209-217, 1986.

[4] Basu, A. and Frazer, L.N., **Rapid determination of the critical temperature in simulated annealing inversion,** *Science,* 1990, Vol. 249, Pp. 1409-1412,.

[5] Kirkpatrick, S., Gelatt, C.D. Jr., and Vecchi, M.P., **Optimization by simulated annealing,** *Science,* Vol. 220, No. 4598, 1983, Pp. 671-681,

[6] Kirkpatrick, S., **Optimization by simulated annealing: quantitative studies,** *Journal of Statistical Physics*, Vol. 34, No. 5/6, 1984. Pp. 975-986,

[7] Corana, A., Marchesi, M., Martini, C., and Ridella, S., **Minimizing multimodal functions of continuous variables with the ‘‘simulated annealing’’ Algorithm,** *ACM Trans. on Mat. Soft.* , Vol. 13, No. 3, 1987. Pp. 262-280,

[8] Bohm-Vitense, E., **Introduction to stellar astrophysics, (Vol. 1),** *Cambridge University press,* 1989.

[9] Bohm-Vitense, E., **Introduction to stellar astrophysics, (Vol. 2),** *Cambridge University press,* 1989.

[10] Stetson, P.B., **DAOPHOT: A computer program for crowded-field stellar photometry,** *Pub. A. S. P.* , Vol. 191, 1987. Pp. 191-222,

[11] Moffat, A.F.J., **A theoretical investigation of focal stellar images in the photographic emulsion and application to photographic photometry,** *Astron. & Astophys,* Vol. 3, 1969. Pp. 455-461,

[12] Alcaino, G. et al, **The CM diagram of the nearby globular cluster NGC 6397,** *Astr. J.* , Vol. 94, No. 4, 1987. Pp. 917-1107,

[13] Alcaino, G., **The globular cluster NGC 6397,** *Astron. Astrophys. Suppl.* , Vol. 29, 1977. Pp. 397-405,

[14] Alcaino, G. and Liller, W., **The main sequence of the metal-poor globular cluster NGC 6397,** *Astr. J.* , Vol. 85, No. 6, 1980. Pp. 680-694,

[15] Alcaino, G., and Liller, W., **Photoelectric UBVRI sequences in the region of the galactic globular clusters NGC 2808, NGC 4372, and NGC 6397,** *Astr. J.* , Vol. 91, No. 1, 1986. Pp. 87-88,

[16] Press, W.H. and Teukolsky, S.A., **Simulated annealing optimization over continuous spaces,** *Computers in Physics, Vol. 5, No. 4, 1991. Pp. 426-429,*