

Genetic Algorithm Optimization of Superresolution Parameters

Barry Ahrens
Rockwell Scientific Company
University of Idaho
P.O. Box 87227
Fayetteville, NC 28304-7227 USA
(01) 910-488-0252
bahrens@rwc.com

ABSTRACT

Superresolution is the process of producing a high resolution image from a collection of low resolution images. This process has potential application in a wide spectrum of fields in which navigation, surveillance, and observation are important, yet in which target images have limited resolution. There have been numerous methods proposed and developed to implement superresolution, each with its own advantages and limitations. However, there is no standard method or software for superresolution. In this paper a genetic algorithm solution for determining the registration and point spread function (PSF) parameters for superresolution is proposed and implemented, and a superresolved image is generated using genetic algorithm optimization of an existing superresolution method.

Categories and Subject Descriptors

I.4.0 [Image Processing and Computer Vision]: General – *Image processing software.*

General Terms

Algorithms, Measurement, Performance.

Keywords

Genetic Algorithms, Superresolution, Registration, Point Spread Function.

1. INTRODUCTION

The solution space for superresolution problems is usually very large, even for relatively small images. The problems are highly computationally complex and have highly domain-dependent features. The dimensionality of the superresolution problem includes that of the unknown high-resolution image as well as that of a measurement vector, both of which may be in the hundreds of thousands of values. The goal of any superresolution method must be to implement a model that will simplify the problem.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '05, June 25-29, 2005, Washington, DC, USA.

Copyright 2005 ACM 1-59593-010-8/05/0006...\$5.00.

Genetic algorithms combine the utility and complexity of problem-solving computer programs with the power and simplicity of natural selection. Genetic algorithms are stochastic search techniques based on the principles of evolution. Extensive research has been performed exploiting the robust properties of genetic algorithms and demonstrating their capabilities across a broad range of problems. These evolutionary methods have gained recognition as general problem solving techniques in many applications, including function optimization, image processing, classification and machine learning, training of neural networks, and system control.

Instead of attacking the superresolution problem directly with highly computationally complex algorithms, genetic algorithms can be applied to reduce the problem to the three genetic operators—mutation, recombination, and selection. Not only can genetic algorithms reduce the complexity of superresolution significantly, but genetic algorithms may also be easily mapped onto distributed computing systems, which are in essence powerful inexpensive supercomputers.

An overview of existing superresolution techniques is in section 2. Section 3 outlines the methodology of this genetic algorithm experiment to solve the superresolution problem, and section 4 presents the results. Conclusions are in section 5; future work is discussed in section 6.

2. SUPERRESOLUTION

Superresolution is the derivation of a high-resolution image from a collection of relatively noisy, blurred low-resolution images of the same scene. Superresolution seeks to overcome the limitations of the imaging device. Because each of the low-resolution images is slightly different from the others, by horizontal shift, vertical shift, or rotation, these images are ideally combined to extract the high frequency details which cannot be captured in any of the low-resolution images.

2.1 Difficulty of Superresolution

Superresolution is an ill-posed inverse problem [6]. In superresolution the task is to use a set of low-resolution images, each of which is a representation of an original target object, and determine the pixellation of the original object on a more refined pixel grid. The original target image that was the object of the low-resolution images through the camera lens is unknown. In the standard computer graphics problem, if the camera settings, camera orientation to the image, lighting, background, original target and so on are known, then the camera image of the object

(the low-resolution images) could be determined. In such a case, we can figure out what image the camera creates from the target object. However, in the inverse problem of superresolution, we know only the low-resolution images, we may know the camera settings, camera orientation, but not the details of the original target object. This inverse problem is highly computationally intensive. The problem is ill-posed as the solution is not unique, and there is considerable information loss in the generation of low-resolution images. To help compensate for the loss of information, a priori information is used to place constraints on the solution set, such as using smoothness operators.

The blurring, or degradation, of an image can be caused by many factors, such as by movement during the image capture process, by the camera or, when long exposure times are used, by the subject. Other causes include out-of-focus optics, use of a wide-angle lens, atmospheric turbulence, or a short exposure time, which reduces the number of photons captured. A blurred image can be approximately described by the following equation.

$$y = Hx + n \quad (1)$$

The blurred image is represented by y , H is the distortion operator, x is the original true image, and n is the additive noise, introduced during image acquisition, that corrupts the image. The distortion operator H is also called the point spread function (PSF). This function, when convolved with the image, creates the distortion. The unknown image x represents the "perfect" image of unbounded resolution.

2.2 Superresolution Methods

Superresolution was first proposed by Tsai and Huang [6], with the utilization of the frequency domain approach to reconstruct one improved resolution image from several downsampled images, based on spatial aliasing. Many other methods have been proposed for superresolution. In the projection onto convex sets (POCS) method [13], an initial estimate of the high-resolution target image is updated iteratively based on the error measured between the observed and synthetic low-resolution images obtained by simulating the imaging process with the initial estimate as the input. In the Schultz and Stevenson method [12], the maximum a priori (MAP) estimator is used with the Huber-Markov Random Field (HMRP) prior.

In the iterative method of Irani and Peleg [8], the first step is to define a reference low-resolution frame, such as frame 0 in a collection of N images. This frame is "upsampled" (expanded) to the desired superresolved image resolution, calculating the pixel values in the larger image by any one of several methods, such as bilinear interpolation. Then the information from the rest of the low-resolution images is used to produce a more accurate high-resolution image. In the iterations of this algorithm, the error of each low-resolution image is back-projected to the best guess high-resolution image.

There have been several attempts at superresolution using Bayesian methods. One Bayesian method was developed by Cheeseman at NASA for reconstruction of planetary images [2]. Bayesian methods use Bayes' Theorem, which states that the probability of B given A can be found if we know the probability of A given B , the probability of A and the probability of B separately. In the inverse problem of superresolution, in which we know the low-resolution images of the camera, and want to find

the original target, the Bayesian method assumes that there are different images, which could have produced the low-resolution images. It is the object of the Bayesian method to find the most probable target surface which created the low-resolution images. This Bayesian method requires a high degree of computational complexity, and is not likely to be modeled and implemented efficiently on an ordinary personal computer.

Another Bayesian method was proposed by Tipping and Bishop [14]. The key development of this method is the use of Bayesian, rather than simply maximum a priori (MAP), techniques by marginalizing over the unknown high resolution image in order to determine the low resolution image registration parameters. However, the computational cost for this method is so great that in the original experiment, patches from the center of the low resolution image were used to find the horizontal shift, vertical shift, rotation, and point spread function values. These image patches were only of size 9 by 9 pixels.

2.3 Genetic Algorithms and Superresolution

Genetic algorithms have been used to solve many problems, and it is not beyond their scope to solve an ill-posed inverse problem such as superresolution. Evolutionary computation techniques are powerful optimization techniques and are suitable for optimization problems that arise when solving practical inverse problems [10]. Unlike the traditional, gradient based search techniques, one characteristic of genetic algorithms is the effectiveness and robustness in coping with uncertainty, insufficient information and noise. Likewise, genetic algorithms are better at handling integer variables than continuous variables, because of the structure of the genomes which are implemented. A variable in a genetic algorithm is usually implemented within a range of values using a binary string of characters. Modeling a similar value as a continuous variable greatly increases the number of possible values. As all pixel values in the superresolution problem are integers, the superresolution model works well with genetic algorithms. The values are easily constrained with upper and lower bounds, as pixel values may range from 0 to 255 for an 8-bit grayscale image.

3. METHODOLOGY

3.1 Application of Genetic Algorithms

Genetic algorithms employ a search technique with simple operators, thus simplifying the computational complexity of the search, as well as abstracting the search for a solution from the specific details of the problem. To solve the superresolution problem with genetic algorithms, several different techniques are employed. In this experiment, it is assumed that neither the point spread function of the camera nor the registration of the low-resolution images is known. The iterative method of Irani and Peleg is modified for genetic algorithms in order to produce a superresolved image. The concept of optimizing the point spread function and the superresolved image simultaneously is used, as found in the Tipping and Bishop method.

3.2 Superresolution Genome

In solving the superresolution problem, genetic algorithms are applied in two steps. First, a genetic algorithm is employed to find the optimal set of registration parameters. Next, another genetic

algorithm uses the registration parameters to determine the point spread function parameter and the superresolved image.

In the image registration step, the registration values are arranged in an array of integers to create the genome for this experiment. For example, for a set of 16 low-resolution images, the genome consists of 49 values that are used to code a possible solution to the optimization problem. There is one rotation value, one horizontal shift value and one vertical shift value for each image, and a single PSF value for all images (see Figure 3).




Figure 1. Superresolution Genome

Constraints are built into the genomes. It is important to limit the search using all available a priori knowledge. Not only does this speed up the search, but it also helps prevent populations from stagnating at local minima with improper values. In this experiment, it is known that the point spread function is greater than zero. An upper bound can also be placed on the point spread function. In like manner, reasonable upper and lower bounds are placed on the rotation and horizontal and vertical shift.

3.3 Fitness Function

In the first genetic algorithm to determine image registration, each genome is composed of an array of values describing the motion of each of the low-resolution images. The fitness function uses the difference in sub-pixel values of two low-resolution images to which the motion parameters have been applied. In a given population, one image is compared with each of the other images, and the sum of the mean squared error is the fitness value.

In the second genetic algorithm to determine the point spread function and superresolved image, the mean squared error is determined for the difference between the superresolved image and the upsampled low-resolution images. To upsample the low-resolution images, each image is expanded to the desired resolution, and the point spread function and registration parameters are used to create a modified image. Major differences between this image and the superresolved image of the genome deserve special attention, so some of these pixels are selected for mutation.

3.4 Input Data

The data for the experiment is taken from a digital camera with the same camera settings for each frame. These are like the first

stop sign image in Figure 6. The camera is moved slightly from frame to frame. In this way, the data contains multiple low-resolution samples of the real-world image, each of which provides additional information about the object. The point spread function for all images is assumed to be the same.

3.5 Crossover

In the superresolution genome, 2-dimensional crossover was implemented for the 2-dimensional array. In order to preserve the schema in the arrays, rectangular patches in the parents are selected for crossing over. Because this could lead to irregular results, it was found that the size of the rectangular patches had to be constrained to 3 x 3 pixels.


3.6 Mutation

In the superresolution genome, the best guess for the superresolved image is represented in a 2-dimensional array of 8-bit pixel values. In the mutation step, if pixels are mutated arbitrarily, without constraints, they will create a salt and pepper effect. This could be removed at a later stage using a low-pass filter. Another solution is to have constraints on the mutation. In this experiment constraints were placed on mutation. Pixels were selected for mutation in two ways: (1) in a random selection, and (2) when large differences were found in the fitness function between the low-resolution images and the superresolved image. The possible values for randomly selected pixels were constrained to correspond to the range of values of neighboring pixels. The possible values for pixels selected from the fitness function were constrained to correspond to the low-resolution image values.

4. RESULTS

There were over 250 trials in this experiment, with significant differences in genetic algorithm parameters, regularization parameters, and general methodology. The genetic algorithm parameters include the size of the population, selection method used, probability of crossover, and probability of mutation. The regularization parameters include all types of constraints placed on the genomes and calculated values. For instance, the “salt and pepper” image was created when no constraints were placed on the mutation values—with a high rate of mutation, this resulted in many pixels out of normal range. The general methodology includes the method used to calculate the superresolved image. A thresholding operation was implemented to produce the best result, in order to produce more refined edges. The size of the superresolved image varied greatly in different trials. The best image had a size of 44x44 pixels, improved from a low-resolution size of 27x27 pixels.

The image registration and PSF were determined simultaneously. Then the superresolved image was produced. The optimization was toward a minimum error of the actual superresolved pixel values compared to pixel values of the registered low-resolution images. The convergence for best sample trial is also displayed.




Low-Resolution Image
(27x27 pixels)

Superresolved Image
(81x81 pixels)

Figure 2. Sample Trial: Salt and Pepper Image

Early trials demonstrated the need for extensive experimentation with the smoothing prior. A lax smoothing prior allowed mutation to produce a “salt-and-pepper” effect (see Figure 2). This was the result of a relatively high mutation rate, coupled with the failure to constrain new mutation values to normal pixel values. Pixel values for the grayscale image range from 0 to 255. However, a mutated superresolved pixel value must conform to the low-resolution values in the near neighborhood.

On the other hand, a smoothing prior that was too strong produced an image lacking the high-frequency details. For instance, if a pixel value is constrained to be exactly the same value as that of its neighbors, or a value that is very close to its neighbors, then the high definition values will be lost. The pixel values must be allowed to vary significantly in the range defined by the values of the low-resolution pixels. A sample result of a trial in which the smoothing prior was too constrained is displayed in Figure 5.3. The convergence of this trial was fast compared with other trials (see Figure 5.4).



Low-Resolution Image
(27x27 pixels)

Superresolved Image
(400x400 pixels)

Figure 4. Sample Trial: Unconstrained PSF


The best trial was 44x44 pixels, improved from 27x27 pixels. The high frequency details were preserved, and the edges in the image were well-defined. The superresolution process resulted in a much brighter, sharper image (see Figure 5.11).



Low-Resolution Image
(27x27 pixels)

Superresolved Image
(44x44 pixels)

Figure 5. Sample Trial: Best Image of 250 Trials



Low-Resolution Image
(27x27 pixels)

Superresolved Image
(81x81 pixels)

Figure 3. Sample Trial: Excessively Smooth Image

The constraints placed on the PSF were very important. When the PSF was allowed to be very large, the results would quickly converge. However, the superresolved image is just a large blur. The PSF indicates the blurring of an image from the discretization process of the camera. For a very large PSF, all surrounding pixels blend into each other. This can produce a circular effect or a wavy effect. A sample trial without constraints on the PSF is displayed in Figure 4.

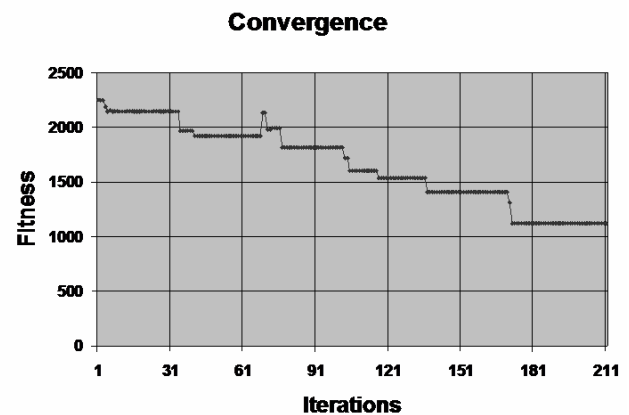


Figure 6. Convergence for Best Image

Parameters for the best sample trial displayed in Figure 5 are listed below.

Population Size:	75
Selection:	Roulette Wheel Selection
Mutation Rate:	0.02
Crossover Rate:	0.8
Low-resolution image:	27x27 pixels
Superresolved image:	44x44 pixels

5. CONCLUSIONS

Superresolution is an inverse ill-posed problem, for which there has been limited success in solving in a timely and efficient fashion. To achieve superresolution restoration of image sequences containing anything more complex than global translational motion, a spatial-domain technique is required. Furthermore, spatial-domain techniques are able to utilize the a-priori information necessary to constrain the search space in inverse ill-posed problems such as superresolution. In this experiment, a genetic algorithm method is implemented as a spatial-domain superresolution technique.

Many trials were performed on a single set of low-resolution images in order to find a set of parameters and constraints that would efficiently produce a superresolved image. Genetic algorithms have been successfully implemented for other inverse ill-posed problems, and they are especially useful in solving the superresolution problem. The ability to incorporate various constraints and the abstraction of the genetic operators from the complex details of superresolution provide a good framework for superresolution.

Several results are displayed to demonstrate the ability of the genetic algorithm programmer to add various constraints to handle different aspects of the superresolution problem. In order to best solve the superresolution problem, all available a priori knowledge should be applied to the algorithm. For example, one challenge was to balance a data fidelity constraint and a smoothing constraint. If a smoothing constraint is too strong, the high-frequency content will be limited (as seen in Figure 5.1). Likewise, as noted by Baker and Kanade [1], if the resolution enhancement factor is too large, the smoothness prior will produce a restored image with very little high-frequency content.

Some of the limitations in superresolution in general concern the existence and uniqueness of the solution. There are instances in which no solution exists, as when observation noise leads to conflicting constraints. Likewise there are instances in which the solution is not unique. It is necessary to determine conditions under which the iteration converges and whether or not that solution is unique. A powerful method for addressing the problems of nonexistence, nonuniqueness and continuity failings of ill-posed problems is through the use of a-priori knowledge which constrains the solution space. To elegantly and effectively tackle ill-posed problems, consider methods which utilize a-priori constraints to obtain regularized solutions to ill-posed inverse

problems. This inclusion of a-priori constraints makes high-quality solutions to even ill-posed inverse problems possible.

The genetic algorithm experiment is compared with the spatial domain method POCS. The advantages of the POCS restoration framework include simplistic framework. Some of the disadvantages of POCS include the possibility of the nonuniqueness of the solution, when the intersection of the convex constraint sets is not a point. Other disadvantages include the dependence of the solution on the initial estimate, and the expensive computational cost, with slow convergence.

Like POCS, the genetic algorithm method has a simplistic framework. The operations of the genetic algorithm are domain-independent. Unlike the POCS method, the genetic algorithm method is not as susceptible to nonuniqueness of the solution. Like POCS, the performance of the GA method is dependent on a good initial estimate. Also like POCS, the computational cost can be rather high, and convergence can be rather slow for such a complex problem.

The genetic algorithm solution for superresolution shares the advantages of other spatial-domain methods. The method allows flexible observation models, including degradation models resulting from problems such as missing pixels. The method also allows for inclusion of a-priori information for regularization, which is necessary for solving inverse ill-posed problems.

The genetic algorithm solution also shares several disadvantages of other spatial-domain methods. The method is more complex and computationally expensive than frequency-domain methods.

6. FUTURE WORK

Genetic algorithms and artificial intelligence have each been employed to solve many complex problems. In future work it would be ideal to combine the strengths of both genetic algorithms and artificial intelligence. In an evolutionary system, successful genes may propagate faster if the individuals are capable of learning. This principle, known as the Baldwin Effect, has been demonstrated in an artificial evolutionary system by [2]. The evolutionary system would benefit from having genomes with the ability to learn.

The superresolution system described in this experiment will be made more sophisticated by adding parameters to take into account the noise in the images. Also, the system can be expanded to test a more diverse collection of images, such as those in which the appearance of the target object begins to change more significantly as the camera orientation changes, as from a moving aircraft. More varied techniques will also be explored for producing the final superresolved image using registration and point spread function parameters.

7. REFERENCES

- [1] Baker, S. and Kanade, T.. Limits on super-resolution and how to break them. In IEEE Transactions on Pattern Analysis and Machine Intelligence, volume 24, pages 1167–1183 (2002).
- [2] Cheeseman, P., Kanefsky, B., Kruft, R., Stutz, J. and Hanson, R. Super-Resolved Surface Reconstruction From Multiple Images, NASA Technical Report FIA-94-12, December, 1994.

- [3] French, R. and Messinger, A. *Genes, Phenotypes and the Baldwin Effect: Learning and Evolution in a Simulated Population*, Artificial Life IV (1994), 277-282.
- [4] Hadamard, J., *Lectures on the Cauchy Problem in Linear Partial Differential Equations*. Yale University Press, New Haven, CT (1923).
- [5] Holland, John H. *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor (1975).
- [6] Huang, T. S. and Tsai, R. Y. *Multiple Frame Image Restoration and Registration*. In *Advances in Computer Vision and Image Processing*, vol. 1 (1984), 317-339.
- [7] M. Irani and S. Peleg, Motion analysis for image enhancement: Resolution, occlusion and transparency. *Journal of Visual Communications and Image Representation*, volume 4, pages 324-335 (December 1993).
- [8] Irani, M. and Peleg, S. Improving Resolution by Image Registration, *Graphical Models and Image Processing*, Vol. 53 (May 1991), 231-239.
- [9] Keller, J., Inverse problems. *American Mathematical Monthly*, 83(2): 107-118 (February 1976).
- [10] Mera, N. S., Elliott, L. and Ingham, D. B. *On the Use of Genetic Algorithms for Solving Ill-posed Problems*, *Inverse Problems in Engineering*, 11, 2 (2003), 105-121.
- [11] Schultz, R., Meng, L. and Stevenson, R., Subpixel motion estimation for superresolution image sequence enhancement. *Journal of Visual Communication and Image Representation*, special issue on High-Fidelity Media Processing, 9(1): 38-50 (March 1998).
- [12] Schultz, Richard R., and Stevenson, Robert L. *A Bayesian Approach to Image Expansion for Improved Definition*. *IEEE Transactions on Image Processing*, Vol. 3, No. 3 (1994), 233-242.
- [13] Stark, H. and Oskoui, P. *High-resolution Image Recovery from Imageplane Arrays, Using Convex Projections*. *Journal of the Optical Society of America*, vol. 6 (1989), 1715-1726.
- [14] Tipping, M. E. and Bishop, C. M. *Bayesian Image Super-resolution*. *Neural Information Processing Systems 15* VS05, 2002.