A Genetic Algorithm Method for Locating the Gulf Stream North Wall

ROBERT W. GRUMBINE * AND JAMESE D. SIMS

NOAA/NWS/NCEP/EMC/MMAB

* *Corresponding author address:* Robert W. Grumbine, #209 5200 Auth Rd., Camp Springs, MD 20706
E-mail: Robert.Grumbine@noaa.gov

ABSTRACT

Genetic algorithms are a tool for solving problems, particularly where there are many possible good solutions and it is relatively easy to tell when you have found one but hard to decide what will produce the best solution beforehand. Locating the North Wall of the Gulf Stream is such a problem. We demonstrate the application of Genetic Algorithms to this problem, as has been applied in NWS operations since 2003 in support of ocean forecasting.

1. Introduction

The location of the Gulf Stream, particularly the northern edge (or 'wall') is an important feature for, among other uses, ocean wave forecasting. Ocean wave models, for example NOAA WaveWatch III [Tolman 1991, Tolman et al. 2002], do not yet include ocean current fields. The currents are important because when waves are running against the currents, they steepen and are more dangerous to shipping [e.g. Phillips, 1977]. Steps in support of improving wave forecast guidance in NCEP has been to implement a regional ocean forecast system (ROFS) [Rivin et al., 2002] in 2001, and an improved model, the Real Time Ocean Forecast System-Atlantic (RTOFS-Atlantic) 13 Dec 2005 [MMAB, 2005 et seq.] which do provide guidance for ocean currents.

Nevertheless, it is desirable to have a more specific description of the location of the Gulf Stream north wall because not only is it a region of high currents, it is a region of warm water that can lead to unstable boundary layers in a cold air outbreak. Usually, the Gulf Stream and its north wall are located by a human analyst, examining satellite sea surface temperature observations and possibly other data sources. It is nontrivial for a human analyst to locate the Gulf Stream from data, consequently, e.g., it is done only 3 times per week by the Navy. Such a process would be difficult to apply to a daily run of a model, and for 5 days of forecast. One difficulty for analysis and users is that areas can be cloud covered for extended periods – as for cold air outbreaks, which is one of the periods of greatest need for such information. A second major drawback is that such analysis is typically only done for present observations. For making wave and marine weather forecasts, it is desirable to have automated guidance as to where the Gulf Stream north wall is going to be.

Since NCEP started running in 2001 an ocean forecast system for both nowcast and forecast purposes to 48 hours [Rivin et al, 2002] there have been objective fields on which algorithms could be run to locate the Gulf Stream north wall. An improved ocean model was implemented in December 2005 [MMAB, 2005 et seq.]. The challenge is to find a method which can do so in an automated way, which is stable, and which produces reasonable guidance for the intended user community of marine forecasters. While there are many approaches to locating the Gulf Stream documented in the literature, a common feature they have is reliance on very high resolution inputs – typically the 1.1 km of the AVHRR [e.g. Cornillon and Watts, 1987]. While the newer RTOFS model has a higher resolution than the ROFS, it is still only 5-10 km in the Gulf Stream region. A further question is how well such methods will perform in the event that the input (such as the ROFS model, we'll see below) is biased in some way, or if its biases change through time.

Rather than tackle the problem of perfecting extant methods, and then have to re-write or retune them when the model changes, we encountered and have applied a very different approach to the problem. Instead of thinking of a specific analytical method which *a priori* might be expected to locate the Gulf Stream north wall, we have applied genetic algorithms. Genetic algorithms are a means of evolving solutions to problems. If one has a general sense of a good starting point, the genetic algorithm (GA) can search a large parameter space efficiently to produce good solutions.

We will first describe some elementary aspects of using genetic algorithms. Then, we will apply them specifically to the problem of locating the Gulf Stream. Finally, we will discuss the results of applying the genetic algorithm in operations within the NWS in its first five years.

2. The Genetic Algorithm

A good introduction to evolutionary computing is Eiben and Smith [1998]. Genetic algorithms are only one such method. Terms are defined in a glossary. For our oceanographic and meteorological interests, it is useful to consider the genetic algorithm (GA) to be a means of evolving good parameter values. As with, for instance, neural networks, the initial inspiration for GAs is biology. The evolutionary process starts with a population – in our case, of parameter values. Then fitness is evaluated – the quality of the fit between manuallydrawn analyses and the automated analysis. The least ”fit” parameter values (those which produce the poorest scores) are eliminated from the population, while the most fit reproduce so as to fill out the allowed population. Reproduction, here, means a descendant parameter value inheriting part of the encoding from one parent, and part from another (crossover). Descendants may also acquire mutations.

Then the process is repeated for another generation (iteration): evaluate fitness, select the best, reproduce, repeat. What was a good parameter value (or, at least, good enough to be carried forward to the next iteration) in one generation may not be in a later generation. Not because it is any worse, but because the competition is better. In carrying out the generation by generation evaluations, what we've described is known as an 'elitist' scheme – the current best values are retained unchanged in the population. This is unlike real biological evolution, but can be useful in a computational situation.

If the parameter space is sufficiently small, we could simply evaluate all possible values and select the best one. If the fitness function were sufficiently smooth, we could apply a familiar Newton's or other such method to locate the optimum. The evolutionary methods have their greatest value in situations where neither of those applies – a large parameter space in which the quality function may not only be unsmooth, but have multiple local optima. Given the number of different methods tried for locating the Gulf Stream north wall automatically, it seems likely that multiple local optima with unsmooth quality function is our case.

Given, again, that there have been multiple attempts at an automated method, it also seems likely that this is a difficult problem, and one in which we are likely to be best served by not presuming much about what the Gulf Stream will look like to the automated finder. Common elements in attempting to automate finding the location of the Gulf Stream include looking for critical values (such as 14 °C at 400 m [Halkin and Rossby, 1985] for instance), and applying edge detection algorithms [e.g. Cornillon and Watts, 1987]. We will use a critical temperature (not in the initial implementation, but added in 2005), the sea surface height, its gradient squared and laplacean (which should be related to 'edges'), and let the GA find optimal weighting parameters for the combined set. We will look, then, for the weights to a function such that the maxima of the function lie along the north wall of the Gulf Stream. The first three terms in the trial function are those implemented in the first operational finder (2001). They are based largely on experiments with sea surface height – we look for, potentially, a location where the height is perhaps large (depending on the size of n1), where the gradient in height is large (a typical way of thinking of the Gulf Stream north wall), or where the laplacean is large (in a geostrophic system, this would correspond to large vorticity, a different view of the Gulf Stream). The final term is based on searching for a critical temperature, as has been used previously (e.g. Halkin and Rossby, 1995) but we do not limit to either the depth or a simple integer value. Our trial function is, then:

$$f(H) = (H^{N1} + a\nabla^2(H)^{N2} + b\nabla^2(H))/(1 + a + b) + d/(0.1 - (T - T_C)^2) \qquad (1)$$

where H is any scalar. The GA will be searching for values of N1, a, N2, b, and d and T_C. The division by the sum of the weights is so that the magnitude of the function is limited. In some cases it may be desirable to permit the GA to be less bounded. In our case, however, we are looking for a relative maximum in f, so that there is nothing gained by having a higher maximum. H, $(\nabla H)^2$ and $\nabla^2 H$ are scaled to not exceed 1 in the model domain, each day.

For initial operational use, a requirement was to use surface fields only. Sea surface height was found to be the best among surface salinity, temperature, and height. A later effort found that using surface temperature as well as the sea surface height fields improved the analysis by about 10% and was implemented operationally in May 2005. Temperatures at depths of 200 and 400 m were also attempted, and gave inferior results.

3. An Example

Let us consider a simple example of a genetic algorithm before moving on to the full complexity of the Gulf Stream Finder – locating the maximum sea surface temperature on a half degree analysis for a given day [15 June 2006; Gemmill, Katz, Li, 2007]. The genome will be a string of bits which represent the i and j coordinates of the maximum. These coordinates range from 1 to 720 and 1 to 360 and require 10 and 9 bits, respectively, to represent as binary integers. One may work with real numbers instead. Values here over the upper limit are folded, such that a bit string which represents, say, 785 for the i coordinate is folded to 65.

The steps after initializing a population are:

Evaluate fitness

Select population from which to reproduce

Reproduce with mutation and crossover

Repeat

The population is initialized by random assignment of a 0 or 1 to each of the 19 bits in the genome, for each member of the population (200 of these genomes – the figure may be varied). For simplicity, we will let fitness be the temperature itself, in Celsius. So to evaluate fitness, the first 10 bits are transcribed into an integer representing the i coordinate of the point to be examined, the next 9 transcribed to the j, and then the SST for that location is read out. The SST itself is the fitness score in the example.

The candidate genomes are those whose fitness is greater than the average of the whole population. This is an adjustable parameter in evolutionary computing. One can be much more selective [Eiben and Smith, 1998]. If it should happen that more than half the genomes have a higher fitness than the average, the median score is used instead. Again, this can be adjusted.

The selection scheme used here, for the simple example, and the Gulf Stream north wall finder, is elitist. All of these most fit (highest scoring) genomes are preserved intact into the next iteration (generation). Reproduction, then, only generates new genomes for the difference between the total population of 200 and the reproducing population. In this implementation of elitist strategy, this could be as few as 100 new genomes in the next iteration (generation).

Given a population from which to reproduce (generate new genomes), we must still decide how reproduction shall be done. There are many possible methods, see Eiben and Smith [1998]. Here we use roulette weighting – the probability that a genome from the parent population is used is proportional to its score divided by the total score of all parent candidates. More fit genomes, therefore, reproduce more often.

In analogy to reproduction like that for humans, two parents are used (diploidy). The resultant genome is identical to the first parent (selected by the roulette-weighting) for the first N bits, and the second (another roulette selection) for the remainder. N itself is a uniformly random variable from 1 to 19 (in this case – it ranges across the entire length of the genome in general). This procedure is known as crossover. It need not be applied to generate every possible descendant. Instead, some fraction of the time, one may generate a completely random new genome. In our case, the crossover proportion is 0.5. Half of the new genomes are random. The random genomes help provide a continued source of new genes (diversity), so as to help avoid premature convergence to a local optimum.

The other diversity-preserving measure is that after new genomes are created, they are subject to mutation. The probability of mutation is 1 in 19 in this case, more generally 1 in N where N is the number of bits in the genome. This is a common value in evolutionary computing [Eiben and Smith, 1998]. Some genomes, therefore, will be unmodified, while others will receive multiple mutations. We expect 1 mutation per genome.

Then we repeat the process, from evaluation through generating descendants until a sufficiently good result is found, or we reach a limit in number of iterations. For the example, we used a limit of 20 generations. At all times, we have a list of currently best genomes.

Consequently, we have a running estimate of the best locations. This is something of a (computational) biological counterpart to ensemble methods.

In the first generation, the results are as for strictly random search. The warmest temperature seen is 30.00 °C, at 14.25 °N, 117.75 °E in the South China Sea. The 15th best is 28.49 °C. Given that warm waters cover much of the ocean area, it is unsurprising that there is little falloff in score. After 20 generations (and 1918 evaluations, it turned out), the best (warmest) water is found to be 32.38 °C (vs the actual maximum of 32.59 found by looking at all 259,200 grid points), at 56.25 °E, 24.75 °N in the Persian Gulf. Table 1 gives the top 15 score and their locations.

Though the genetic algorithm did not find the absolute maximum, it did locate several areas of warm water on the global ocean. This points to a couple of features. One is, as with other optimizers, one may arrive at a local optimum rather than the global optimum. On the other hand, one could also combine the genetic algorithm with a local gradient climber. This would then bear some resemblance to simulated annealing [Metropolis et al., 1953] except that the random stage would be evolved. Another is that we may often want to know the many locations in the parameter space in which there are high scores (warm SST in this case, but in a moment it will mean good ways of locating the north wall of the Gulf Stream).

Figure 1 and 2 show the population evolution through 400 generations of selection, displaying generations 1, 50, 75, 200, and 400. In the first generation, candidates are all over the globe. The search did not exclude land points, though that could have been done, because the sea surface temperature analysis fills in land points. In later generations, we see increasingly many points are focused to the Pacific warm pool, the Red Sea, and the Persian Gulf. Nevertheless, even in later generations, some points appear far from these high quality locations. These are the points which result from larger mutations, or large effect crossover. In more general cases than this simple example, they prevent premature convergence to local optima, and ensure sampling of remote points which might also be highly fit. Figure 2 is as for figure 1, but focuses on the Arabian Peninsula, which is where the warmest water is.

If we disable the mutation and crossover, giving us a purely random search, but otherwise leave the evolutionary program alone (so that it does retain most fit genomes), in 20 generations and 2659 evaluations, the best found is 32.20 °C and 15th is 30.42 °C. This problem has a large fraction of the search space giving very high scores, and evaluates a large fraction of the search space, so random searching can do relatively well.

The model here demonstrates some of the language and character of genetic algorithm (GA) methods, rather than to demonstrate its superiority to other methods. Nevertheless, it is a reminder that in some problems random selection may be valuable, or that exhaustive searching may be appropriate. In the case at hand, arriving at a good answer after searching 1% of the parameter space (approx 2500 samplings vs. 259,200 grid points) is no great savings as the fitness evaluation is extremely inexpensive. Further, this is a rather high portion of the parameter space to be searching.

In the real problem of interest, the north wall finder, the parameter space is one of 42 bits, or $4{\cdot}10^{12}$ values. Sampling 1% of that space is extremely expensive, the more so as evaluating the fitness is a much more expensive proposition than a mere memory lookup. In practice, the finder evaluates fitness about 25,000 times, about $6{\cdot}10^{-9}$ of the space, in 3 minutes on a current desktop computer. Exhaustive evaluation would require approximately 900 years at that pace.

4. Representing the Genetics for the North Wall Finder

We will let the 4 parameters, N1, N2, a, and b be represented by floating point numbers in the range [-4:4) with 8 bits precision (steps of 1/32), d is 6 bits in the range 0 to 1 and T_C is given 4 bits, in the range 16 to 20. This is the genome. Even steps this small are already distinguishable in the fitness score. This gives us a parameter space of approximately 4 trillion members, in excess of what can be explored profitably by brute force. The fitness function is also non-smooth and contains many local optima.

Parents are selected with roulette weighting. The probability that the ith potential parent is selected is proportion to $\frac{P_i}{\Sigma P_j}$, where the sum is over all parent scores P_j.

The crossover rate is 0.5. If the random number (uniform on [0:1)) is less than this, we cross genomes between the selected parent and another.

The mutation rate is 1/42 – given the 42 bits in the genome, we expect 1 bit to flip within each genome. After new genomes are generated, we step through each genome and test whether to mutate (flip a 1 to a 0 or vice versa) each bit, with this probability of doing so.

All parents survive in to the next generation (complete elitism).

The population size is 100. The limiting number of generations is 200. One could set an earlier stopping criterion, such as the best fit being better than some tolerance. But in this case, running to the limit number of generations was more satisfactory.

Population size, generations, stopping criteria, crossover rate, mutation rate, and parent selection rules are all things which can themselves be experimented with. See Eiben and Smith [1998] and references therein for some considerations. The fitness score is the inverse of rms distance between the Navy manual north wall Gulf Stream and the maximum in the GA. For initial simplicity, the maximum was sought separately along each line of longitude in the model output. A result of this is that eddies and other local features disconnected from the main body of the Gulf Stream sometimes are identified instead of the Gulf Stream. Discussion with the Ocean Prediction Center [Sienkiewicz, pers. com. 2002] lead to the conclusion that this was a constructive feature. The OPC interest is in correcting wave model guidance for current interactions, and these other points are definitely active and in need of wave forecaster correction. Consequently this method was continued for this utility, though it did result in limiting the degree to which the automated method could reproduce the manual analysis.

We must also select a range of longitudes along which to assess the quality of the analysis. The eastern boundary is set by the analysis, which stops a 65 °W. (Analysis by the finder, however, is extended to the eastern boundary of the domain.) The western boundary is set at 77 °W, where the Gulf Stream begins significant eastward motion. West of this, the Gulf Stream flows nearly northwards (hence introducing artefacts to the algorithm searching along lines of longitude for a maximum in the genetic function; Auer [1985] refers to this as the non-orthogonality problem). Parameter sets which work well to the east also work well to the west, but the converse is not necessarily, or even often, the case.

5. Performance of the Finder

The original north wall finder was implemented operationally in NCEP on 24 January 2003. The improved version, including the surface temperature in the fitness function, was implemented operationally 20 May 2005. The discussion here has included the full function, and certain coding improvements and population changes. For this paper, we re-ran the finder with consistent (current) parameters for the genetics. As the finder is computationally inexpensive, this provides a consistent scoring, and demonstration of the current method's skill.

Figure 3 shows the quality of the best found north wall, as inverse rms grid point differences, on those days which had Navy analyses for comparison. The horizontal axis is days since 31 January 2003. Since the finder is permitted to locate eddies and active areas away from the Gulf Stream itself, the score will not be as good as it might be. Conversely, some portions of the North Wall are not observable in every analysis, so that the error may be in the analysis. Nevertheless a lower rms error gives more confidence in the procedure. We see in this figure that the median rms error is 35 km, satisfying the original design criterion of 50 km. Somewhat troubling is that the scores worsened through time with the ROFS, approaching 50 km late in 2006. The ROFS model was retired from operations in late 2007.

Figure 4 displays an ensemble of the 15 best north walls and the Navy analysis they were evolved to fit. The analysis extended only to 65 W. Matches were quite close from 77 to 72

W, but as the Gulf Stream curves and recurves more actively, the fit declines. We see that the different members mostly make the same errors between 65 and 69 W, which suggests a model bias that the algorithm could not overcome. Outside the region, 77 to 65 W, in which the evolution was applied, the different members diverge more from each other. Some members apparently picking up on Gulf Stream recirculation.

In another offline run, searching for critical temperatures at depth (in correspondance with methods like Halkin and Rossby [1985]), we find that there is a trend in model temperatures with time, and that the ROFS model is increasingly cold. Rather than 14 °C at 400m [Halkin and Rossby 1985] for the Gulf Stream axis, the best (median) for the north wall is 10.2 °C at 400m, and this figure becomes colder as time moves towards the present. We expect, then, that the reason for worsening scores in the North Wall finder is that ROFS is experiencing climate drift. While the finder remains useful in spite of the climate drift, there are obviously limits to which the method can adapt to drifting input.

6. RTOFS-Atlantic Version

The advent of a new, higher resolution, real time ocean forecast system for the Atlantic ocean (RTOFS-Atlantic [MMAB 2005 et seq.]) at NCEP made it possible and necessary to re-examine the genetic algorithm method used for locating the Gulf Stream north wall. If evolutionary computing, of which genetic algorithms are an example, is to live up to its promises, it must be able to adapt to a new situation easily. As we also wish to expand capabilities with the advent of more capable models, we want the system to be easily modified without either great computational cost or programming effort. Finally, as our concerns evolve, it is necessary that we be able to readily change the fitness definition(s) used by the algorithm to represent more accurately the desired features of the analysis.

Evolutionary methods rely critically on what the user decides is the best result, the 'fitness' function which we will discuss again. With the new model, RTOFS-Atlantic, for input to the North Wall finder, we again defined fitness to be inverse (inverse so that higher score is better) rms distance between the north wall line as found by the automated system and the Navy analysis. But now the automated analysis, rather than sweeping along lines of longitude in search of maxima, traces down the spine of maxima in the evolved 'north wall function'. This enables it to follow recurving stretches of the north wall. The initial points on the north wall are found by scanning for maxima along lines from 82 °W 28.2 and 28.0 °N, to 77 °W 28.2 and 28.0 °N. These then define points on the north wall, and a line of travel along which to look for maxima. The direction of travel can be reversed to find an estimate for the location of the Loop Current in the more extensive model domain of the RTOFS(Atlantic).

The distance is then computed between each point on the automated analysis' line and the nearest point on the manual analysis line. The manual analysis line is first examined to find the point (call it point N) which is nearest to the automated analysis point in question. Then the distance between the automated analysis and the lines formed by points (N,N-1) and (N,N+1) are computed and the minimum taken. If one examines all line segments, distances between automated points and the manual analysis line segments can be extremely small as the recurved segments of the Gulf Stream project lines through much of the ocean even far from the Gulf Stream. Evolutionary systems are very good at locating loopholes in definitions, so that the implementation of this simple concept was elaborated after some very bad (to human eyes) analyses were given very good scores by that loophole.

A different change made for the RTOFS(Atlantic) version of the north wall finder was to work on the model's native horizontal grid. This model, as was the case for ROFS [Rivin et al., 2002] uses a variably-spaced grid. In conducting the initial work on the north wall finder, the precision of the manual analysis (0.1 degree digitization) and operational forecast concern (to be better than 50 km analysis – approximately the width of lines on mariner guidance maps [Sinekiewicz, 2001]) were such that the ease of working with a regular latitude-longitude grid was significant. In the present case, now that the prior generation of finder had established a typical rms error of only about 35 km in spite of the artefacts noted above, it was decided to attempt all possible precision. The RTOFS(Atlantic) grid varies from approximately 5 to 10 km spacing over the Gulf Stream region [MMAB 2005 et seq.]. Locations are reported to 0.01 degree by interpolation from these grids. As future generations of model improve resolution, this precision will become more meaningful. In the mean time, it ensures that the closeness of fit between manual and automated analysis is not limited by the grid spacing (which alone guaranteed a 0.05 degree minimum difference between manual and automated analysis in the ROFS version as the manual is rounded to 0.1 degree, and automated grid was staggered such that points were on x.05, x.15, ... degrees latitudes and longitudes).

In the ROFS version of the finder, sea surface height, its gradient squared, and its laplacean were each rescaled every day to fit the range [-1,1]. The hope was that this would lead to a series of weights in the north wall function (equation 1) which would become slowly varying in time, as the relative magnitude of the Gulf Stream signal (in these terms, compared to farther afield values) would be relatively slowly varying, versus the detailed values themselves. Such was not the experience. So in this edition, values for gradient squared of sea surface height, laplacean of sea surface height, and gradient squared of sea surface temperature are scaled (multiplied) by constant values (10^5, 10^9, 10^3), sufficient that each is of approximately the same order of magnitude. The sea surface heights used are the 25 hour averages centered on the valid time (nowcast, 1 day forecast, etc.), except for the final forecast date when the last 25 hours of the forecast period are used. This provides improved tide suppression over the ROFS version, where only 24 hour averages were possible. The variations in parameters then are attempting to find universal parameters descriptive of the north wall in the model, irrespective of far field behavior of the ocean. Our initial scaling of the fields to approximately equal magnitudes lets the GA decide the relative importance easily. We could leave the fields unscaled, but if it turned out (as it does) that the relative contributions are within a factor of 10 of each other, we would need to permit the GA to search over ranges of plus or minus 10^6 as opposed the factor of 10 we do permit. This gives us a much more rapid convergence in the evolutionary process.

As in the ROFS version of the finder, we defined a function which will constructed to have a maximum along the north wall of the Gulf Stream, then seek parameters to that function which will make this be the case. For the RTOFS version, with a vastly larger model grid (approximately 2 million points, vs. 63 thousand), we use a function simlar to ROFS version, but computationally faster to evaluate over the grid:

$$a(\nabla H)^2 + b\nabla^2(H) + c(\nabla T)^2 \tag{2}$$

where H is the sea surface height, and T is the sea surface temperature. a, b, c are all represented with 8 bits. c is in the range [0:10), a and b are in the range [-10:10).

The genetic algorithm now searches for the best weights a, b, c, with best defined as producing the minimum rms error in location. The population size is 100, and 50 generations are evaluated. Crossover rate and mutation rate are unchanged, at 0.5 and 1/N (1/24).

Recent results available show rms error of about 8 grid points, approximately 45 km, and originally on implementation it was about 35 km. Given the number of changes in how errors are computed, this figure is not strictly comparable to that from the ROFS finder. Nevertheless, it is reassuring that the differences are comparable to the ROFS system, and within the tolerances required for operational use even with this much more stringent comparison to the analysis. The system for this vastly larger model output is slower, about 10 minutes for the RTOFS version versus 3 for the ROFS version. But this compares quite favorably to the 100-fold increase in grid points between the two systems.

7. Conclusions

This paper introduced two topics: genetic algorithms and the NCEP operational automated Gulf Stream north wall finder.

Genetic algorithms have shown their utility and flexibility in our specific application. While remaining computationally tractable, they search a large parameter space to locate objective functions which can be used to trace the Gulf Stream's north wall. The computational requirement is 3 (ROFS) or 10 (RTOFS) minutes on recent desktop computers for those days which have an analysis for comparison. They further managed to produce meritorious results in spite of a significant model drift in the ROFS. The technique also satisfied our requirement that it be readily applicable to new models as they become available.

A feature of this automated method, versus manual analyses of observations, is that it can be applied to forecast fields as well, and without any signficant further cost (the main cost is in developing the latest genomes; evaluating the North Wall location once the genomes have been found is a negligible cost). Further, in working from the model, there are never observation gaps.

8. Glossary

Crossover : Inheritance process in which two parents are selected and the first N bits of the first parent are used in the descendant, and the remaining genes are from the second parent. n.b. there are many other crossover methods

Descendant : A genome which is developed by crossover and mutation from some other (parent) genome

Elitism : Retaining into a subsequent generation the parents of the current generation

Encoding : The translation process between the bits in the genome and the parameter values.

Evolutionary Computing : EC; The collection of methods, of which GAs are one, which use evolutionary methods and approaches to solve problems.

Fitness : The quality of the result Fitness function : a function which assesses the quality of a parameter set (genome)

Fitness landscape : The magnitude of the fitness function displayed, or considered, over all parameter values (dimensions)

GA : Genetic Algorithm; a method of using evolutionary processes to develop improved parameter sets

Generation : A complete iteration of evaluating the fitness of a population, selecting parents, and replacing by descent from the selected parents those genomes which were rejected for parenthood.

Gene : A collection of one or more bits which represent a value

Genome: A collection of one or more Genes, a parameter set

Inheritance : The process by which genes are passed from parents to descendants

Mutation : Flipping a bit from 1 to 0 or vice versa.

Parent : A genome used as a basis for developing more genomes (descendants)

Population: The set of all genomes under consideration

Reproduction; The process by which new genomes are generated from old

Source codes used for the RTOFS and ROFS North Wall finders is available at http://polar.ncep.noaa.gov/mmab/models/tnXXX/

Auer, Stephen J. Five-Year Climatological Survey of the Gulf Stream System and its Associated Rings *J. Geophys. Res.*, **92**, 11,709-11,726, 1987.

Cornillon, Peter and D. Randolph Watts Satellite Thermal Infrared and Inverted Echo Sounder Determinations of the Gulf Stream Northern Edge *J. of Atmos. and Oceanic Tech.*, **4**, 712-723, 1987.

Eiben, A. E. and J. E. Smith, *Introduction to Evolutionary Computing*, Springer-Verlag, Berlin, 299 pp., 1998.

Gemmill, William, Bert Katz and Xu Li Daily Real-Time Global Sea Surface Temperature - High Resolution Analysis at NOAA/NCEP NOAA / NWS / NCEP / MMAB Technical Note Nr. 260, 39 pp, 2007.

Halkin, Daniel, and T. Rossby The Structure and Transport of the Gulf Stream at 73 °W *J. Phys. Oceanography* **15**, 1439-1452, 1985.

Metropolis, N., Arianna Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller Equation of State Calculations by Fast Computing Machines *The Journal of Chemical Physics*, **21**, 1087-1092, 1953.

MMAB Real Time Ocean Forecast System (Atlantic) http://polar.ncep.noaa.gov/ofs/ 2005, last accessed 13 May 2008.

Phillips, O. M., The Dynamics of the Upper Ocean, 2nd ed., Cambridge Univ. Press, 336 pp., 1977.

Rivin, Ilya, Lawrence Burroughs, James Partain, and John Kelley, NOAA Coastal Ocean Forecast System, TPB 489, MMAB TN 212, 2002.

Sienkiewicz, Joseph personal communications 2001, 2003.

Tolman, H. L. A third-generation model for wind waves on slowly varying, unsteady and inhomogeneous depths and currents *J. Phys. Oceanography*, **21**, 782-797, 1991.

Tolman, H. L., B. Balasubramaniyan, L. D. Burroughs, D. V. Chalikov, Y. Y. Chao, H. S. Chen, and V. M. Gerald Development and implementation of a wind-generated ocean surface wave models at NCEP *Weather and Forecasting*, **17**, 311-333, 2002.

List of Tables

Table 1. Temperature (Score), i,j coordinate of point on grid, and Latitude-Longitude of the location

Temperature °C	I	J	Longitude °E	Latitude ° N
32.38	112	130	56.25	24.75
32.14	105	131	52.75	24.25
30.97	200	162	100.25	8.75
30.87	86	151	43.25	14.25
30.76	93	148	46.75	15.75
30.64	285	169	142.75	5.25
30.57	196	184	98.25	-2.25
30.56	292	172	146.25	3.75
30.54	107	138	53.75	20.75
30.48	210	160	105.25	9.75
30.44	216	146	108.25	16.75
30.43	274	165	137.25	7.25
30.42	312	182	156.25	1.25
30.42	210	169	105.25	5.25

List of Figures

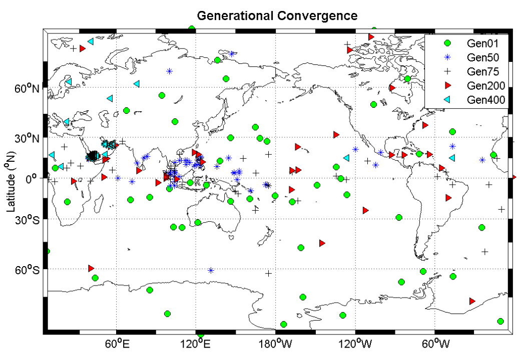

FIG. 1. Figure 1: Population of evolutionary candidates, by generation

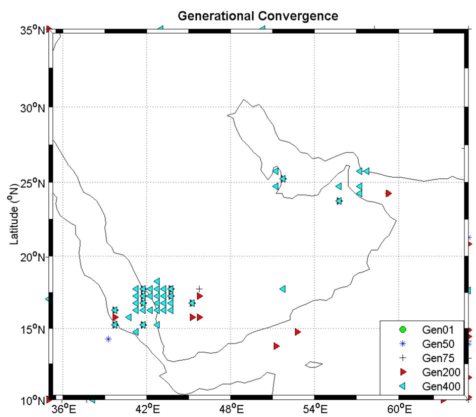

FIG. 2. Figure 2: Population of evolutionary candidates, by generation. Focused in area of highest temperatures

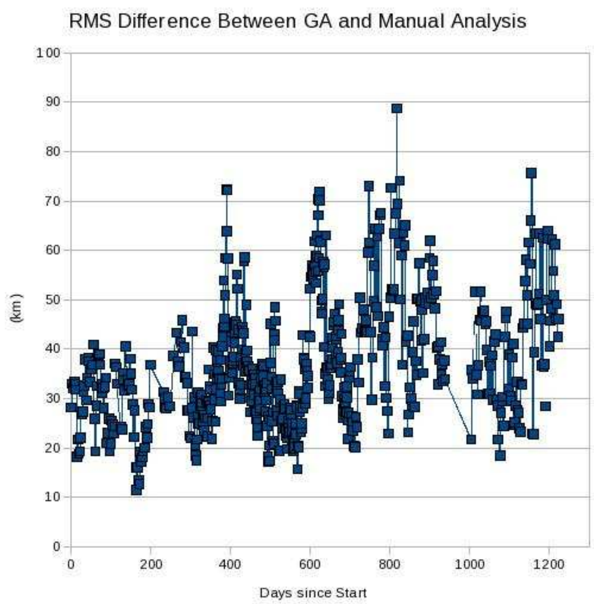

FIG. 3. Figure 3: RMS difference in north wall location in the ROFS model as diagnosed by the genetic algorithm as compared to manual analysis of satellite data. Days since 31 January 2003

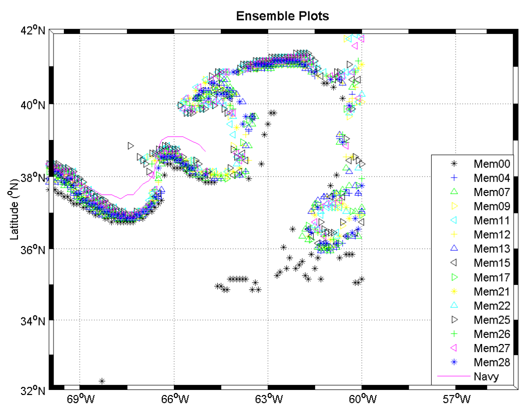

FIG. 4. Figure 4: Ensemble plot of the 15 best fit genetic algorithm north wall edges