# 19A.2 THE APPLICATION OF AN EVOLUTIONARY ALGORITHM TO THE OPTIMIZATION OF A MESOSCALE MODEL

Dr. David Werth
Savannah River National Laboratory
Aiken, SC 29808

and

Dr. Lance O'Steen
Savannah River National Laboratory
Aiken, SC 29808

## ABSTRACT

Improvement in mesoscale atmospheric model simulations is often thought to be primarily a matter of finer spatial resolution. While this is generally true, there is a limit to the improvement one can obtain by simply decreasing the grid size of a numerical model. Further improvements in forecasts may be achieved with better model parameterizations, but this leaves the mesoscale modeler with the task of determining which parameterizations to use for a specific problem and what values to use for individual model parameters. The accuracy of a given numerical simulation is often a matter of a judicious choice of these values.

In this presentation, we show how a simple evolutionary programming (EP) algorithm can optimize a given set of parameters in a mesoscale atmospheric model with respect to agreement between simulation and observations. This is illustrated using the Regional Atmospheric Modeling System (RAMS). As an initial test case, data from a RAMS simulation with a default set of parameters, rather than actual data, is used to calculate an objective function. Ideally, the model parameters will evolve toward those of the default simulation as the objective function is minimized. This type of experiment also tests the ability of EP to find the global minimum of the objective function since the optimum is known.

Our primary goal was to demonstrate that an EP algorithm can provide a systematic and objective method for optimizing a mesoscale atmospheric model. We are now working towards exploring the path the model follows as it is pushed towards its target, and how such a scheme could be applied operationally.

**Keywords:** Meteorology, Genetic Algorithm, Optimization

## 1. INTRODUCTION

Recent atmospheric simulations of the Salt Lake Valley by Zhong and Fast (2003) demonstrate that there is a limit to the improvement one can obtain by simply decreasing the grid size of a numerical model. They conclude that "relatively large forecast errors can still exist even with sufficiently fine spatial resolution", and that further improvements in forecasts will require better model parameterizations. While this may not be too surprising, it does leave the mesoscale modeler with the task of determining which parameterizations to use for a specific problem and what values to use for individual model parameters.

Mesoscale atmospheric models typically have a large number of user-input and "hard-coded" parameters that exert considerable influence on the model behavior. Depending on one's viewpoint, this is either a blessing or a curse. In either case, decisions have to be made based on previous experience, intuition or expert advice. The majority of model parameters control surface processes and thus the greatest impact of these parameters is on the fluxes of mass, energy and momentum between the earth's surface and the surface-layer of the atmosphere. In addition, fluxes within the boundary layer are influenced by parameters that control turbulent length scales and radiative transfer in the atmosphere. The effect of unresolved topography on air flow within the boundary layer is also determined by model parameters. Even the size of the domain and location of domain boundaries, including nested grids, can be crucial. The accuracy of a given numerical simulation is often a matter of a judicious choice of model parameters. In fact, as most modelers know all too well, generating a simulation that simply runs without a catastrophic failure can depend on the choice of model parameters.

We show that a simple evolutionary programming (EP) process can optimize a given set of parameters in a

mesoscale atmospheric model with respect to agreement between simulation and observations. This is demonstrated using the Regional Atmospheric Modeling System (RAMS) developed at Colorado State University and synthetic 'observational' data derived from a 'target' RAMS simulation. (Synthetic data is used so that the global optimum will be known exactly.) The simulation focuses on the Salt Lake Valley, since it represents a complex terrain situation. An extensive meteorological database, the Vertical Transport and Mixing Experiment (VTMX) field campaign (Doran et al, 2002), is available for future work with real observations.

## 2. OPTIMIZATION EXPERIMENT

### 2.1 *The Regional Atmospheric Modeling System (RAMS)*

The Regional Atmospheric Modeling System (RAMS, Pielke et al. 1992), is a three-dimensional, non-hydrostatic finite-difference solution of the geophysical equations of motion. A wide range of atmospheric motions may be studied with RAMS due to its use of a two-way nested grid system. Initial and boundary conditions are typically generated from the gridded output of large-scale models, e.g. the GFS model of the National Centers for Environmental Predictions (NCEP). Topographic features are incorporated through the use of a terrain-following vertical coordinate system and topography is further modified with additional parameterizations to represent unresolved terrain. Soil and vegetation exchange energy and water with the atmosphere through the Land Ecosystem-Atmosphere Feedback (LEAF) model (Walko et al. 2000). In typical applications, the second-order turbulence parameterization of Mellor and Yamada (1982) is used to parameterize vertical diffusion. Clouds and radiative transfer are also heavily parameterized through a suite of sub-models. For the simulations performed here, a 4-grid system is used with a grid spacing of 45 km for the outer grid and 1.67 km for the inner most grid that encompasses the Salt Lake valley (Fig. 1).

### 2.2 *The Evolutionary Programming Algorithm*

A set of model parameters to be optimized are selected and initial values for these parameters are set randomly (but far from the target). Simulations with this initial parameter set are herein referred to as 'base case' simulations. These 'parent' parameters are then perturbed about their initial values based on a Gaussian distribution of variates with user-defined initial variances for each parameter. N sets of perturbed parameters were then generated, where N is the number of available processors. For the results presented here, 32 AMD

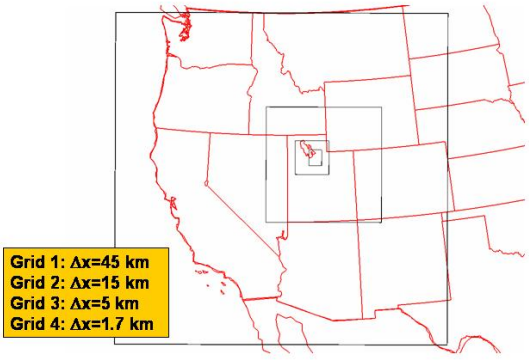


Figure 1 Simulated RAMS domain with spatial resolution of each.

Opteron processors were used with a single simulation run on each processor. Each simulation was run for 30 hours, although only the final 24 hours were utilized for parameter optimization. (A single simulation required about 8 hours of CPU time.) Each set of 32 simulations represents a 'generation' in the evolutionary process.

After all the simulations are finished, the numerical results were compared to a set of synthetic Salt Lake Airport soundings extracted from the 'target' simulation during the final 24 hours. An rms error was calculated for each simulation with user-specified weights for each variable used in the optimization. The parameter set producing the best score was selected, and the entire process repeated with this parameter set as the new parent for the next generation. The variance for each parameter, however, is now reduced in a user-specified manner (described below). The RAMS parameters that were optimized are given in Table 1, along with their initial and target values.

The simulations used here employed the envelope orography procedure for creating the topography in RAMS, with EP1-4 as the topographic enchancement factors for this procedure. EP5 is a standard RAMS parameter for calculating a surface roughness height based on sub-grid topography, and EP6 dictates the vertical motion necessary to initiate convection. EP7 is the bare soil roughness. The albedos in Table 1 are for dry soil and saturated soil. (Albedos for intermediate soil moisture are linearly interpolated between these two values.)

| | | Initial Value | Target Value | | Initial Value | Initial Value | Target Value |
|---|---|---|---|---|---|---|---|
| EP1 | Topo Init for Grid 1 | 1.76 | .5 | EP13 | Soil Moisture (top) | -1.08 | -0.5 |
| EP2 | Topo Init for Grid 2 | .121 | .5 | EP14 | Soil Moisture (bottom) | -.529 | -0.5 |
| EP3 | Topo init for Grid 3 | .158 | .5 | EP15 | Convective Vert. Scale | .1512 | .1 |
| EP4 | Topo init for Grid 4 | 1.610 | .5 | EP16 | Nocturnal vert. Scale | .8194 | 1 |
| EP5 | SGS Topo Roughness | 0 | 0 | EP17 | Pielke LW down | 2.1522 | 1 |
| EP6 | Min. vertical velocity | .0174 | .005 | EP18 | Cloud scale factor | .9562 | 1 |
| EP7 | Soil Roughness | .0574 | 0.1 | EP19 | Lateral Nudging Timescale | 1001.5 | 900 |
| EP8 | Wet Albedo Scaling factor | .4121 | 0.6 | EP20 | Interior Nudging timescale | 7394.8 | 28800 |
| EP9 | Dry Albedo | .4926 | 0.31 | EP21 | Top Nudging Timescale | 5086.6 | 28800 |
| EP10 | Veg. Frac. #1 (Evergreen tree) | .1324 | 0.8 | EP22 | RAMS grid shift –X | -.4393 | 0 |
| EP11 | Veg. Frac. #2 (Short Grass) | .1555 | 0.8 | EP23 | RAMS Grid Shift - Y | .0445 | 0 |
| EP12 | Veg. Frac. #3 (Semi-Desert) | .1291 | 0.1 | | | | |

Table 1 Evolutionary Programming (EP) parameters

There are 31 classes of land-use/vegetation in RAMS and each has a set of properties which are used to describe their impact on surface fluxes of mass, energy, momentum and radiation. In this optimization procedure, only the vegetation fractions for the 3 classes with the largest population on the 2 innermost grids were optimized - evergreen needle-leaf tree, short grass and semi-desert. Soil moisture over the simulation domain was taken from the GFS simulation at $1^{o}$ resolution and 2 vertical soil levels. This data was then interpolated to 11 soil levels in the RAMS simulation. The two soil moisture parameters in Table 1 were used to scale all the soil moisture data before applying this data to the RAMS grids. This is very crude, but soil moisture is a high-impact parameter and is poorly known, thus it is an ideal candidate for optimization. EP parameters 15 and 16 are factors for the vertical length scales calculated for both convective and nocturnal turbulent diffusion. These are applied in the Mellor-Yamada turbulence parameterization that uses a TKE-scaled length for neutral /convective conditions and a length scale for stable conditions.

The Pielke radiation scheme was used for these simulations since clouds were not a factor. However, during the night, the simulations exhibited a strong tendency to the over cool the atmosphere near the surface and increasing the net radiation to the surface was the only way we could find to deal with this problem. Hence the use of a simple scaling parameter for the down-welling longwave radiation (EP17). While this might not be the source of the problem, it did prove to be and effective way to deal with it. EP19,20, and 21 are standard input parameters to RAMS that control the "nudging" of the lateral, interior and top grid points to the results from a global simulation that the RAMS simulation is nested within. The final two parameters are used to make small, horizontal adjustments to the interpolation of the global model data to the RAMS outer grid. This is used to account for small errors in the location of fronts or small scale features in the global data that are lead to significant inconsistencies with observations.

As noted previously, the default parameters in Table 1 were used to generate the "target" simulation from which the observational data was extracted, in this case synthetic soundings at the Salt Lake airport. A totally different set of parameters was used to initialize the EP. In both cases the initial and boundary conditions were

generated from global NCEP simulations for October 25, 2000. The objective function was simply the rms error between the observational data and the simulated sounding data at 6 hour intervals over a 24 hour period. The errors in wind speed, wind direction, temperature were weighted such that their contributions to the objective function were approximately equal while the dew-point was given a very low weight and consequently a small contribution to the objective function. Thus, parameters affecting humidity can be adjusted to improve the other contributions to the score without penalty due to dew-point errors. This approach seems reasonable if an accurate dew-point is not important for a particular application. However, as we shall see this logic can lead to unexpected, sub-optimal results.

### 3. RESULTS AND DISCUSSION

The first experiment was based on a reduction of $e^{-\beta*(1-F)}$, where F is the ratio of the current best score and the base score. The parameter β is used to determine how fast the variances of the parameter perturbations are reduced with improving score. This reduction in variance is important. If it is too rapid, the likelihood of finding the global minimum in the objective function is reduced. On the other hand, if it is too slow the search becomes inefficient.

The score results for the first optimization attempt are shown in Figure 2. Note that the scores have been normalized with respect to the initial parent (default parameters). During this initial optimization the variance factor β was manually adjusted to control/test the optimization procedure. Of course, interactive control is not a particularly desirable feature for optimization and one goal of this study was to determine how to best control the parameter variances based on the score history. In any case, over 231 generations (32 RAMS simulations/generation), the score decreases from 1 to about 0.04.

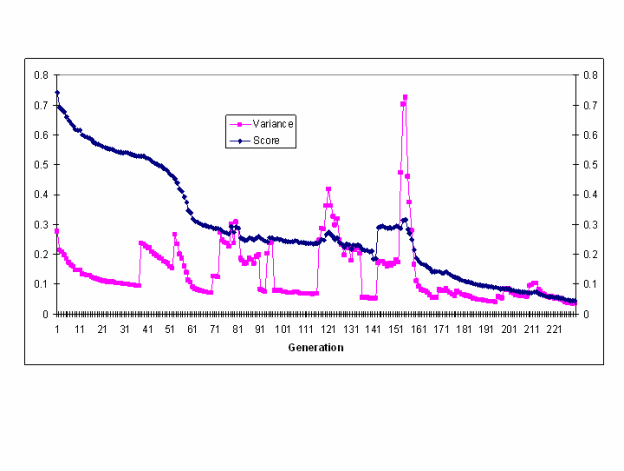


Figure 2 Variance and best score of the various model runs as a function of generation (Generation 1 not on plot).

Note the large drop in score during the first generation from 1 to 0.74. This was either a matter of good fortune or a poor choice of initial parameters. As the score fell the initial decline in the variance was too large and it appeared that the optimization had become trapped in a local minimum of the objective function with a score of 0.52. The variance was increased at generation 41 and again at generation 55 in an attempt to force the optimization out of the assumed local minimum and accelerate the convergence to the global minimum of 0. This resulted in a rapid decrease in score to about 0.32 at generation 61. At this point the decline in the score abated substantially decreasing from 0.32 to 0.18 over the next 80 generations. Attempts to accelerate this decline in score by changing the variance had little impact. At this point examination of the EP parameters indicated that the near surface soil moisture was far too high (Fig. 3) and consequently the diurnal variation of the surface temperature was too small (Fig. 4) and, more important, the surface dew-point temperatures were far too high (Fig. 5).

It appears that although the error in the surface temperature (first 2 RAMS model levels) was relatively large, this was having a small impact on the overall score and consequently the EP having trouble determining that soil moisture was the key to further improvements in the score. Even relatively large increases in the variance (generation 121) did not seem to help. However, it does seem clear that the objective function without a dewpoint error has a relatively deep local minimum (valley). Based on the above findings, a dewpoint error was added to the objective function, although humidity is not typically a major concern for most applications.

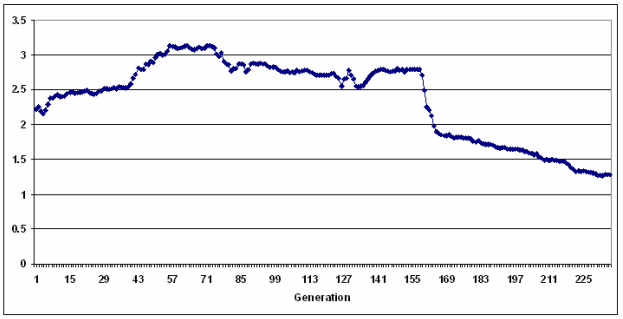


Figure 3 Upper level soil moisture (normalized by the target value) as a function of generation.

The goal was to alter the objective function surface such that the global minimum would be easier to find. This was done at generation 141. The result was an immediate increase in the score from 0.18 to 0.29 due to the additional dewpoint error term. This produced an immediate but small increase in the score that was

followed by a dramatic decrease that continued to the end of the optimization. The improvement in soil moisture can be seen in Figure 3 as the error drops following the change in dewpoint weighting.

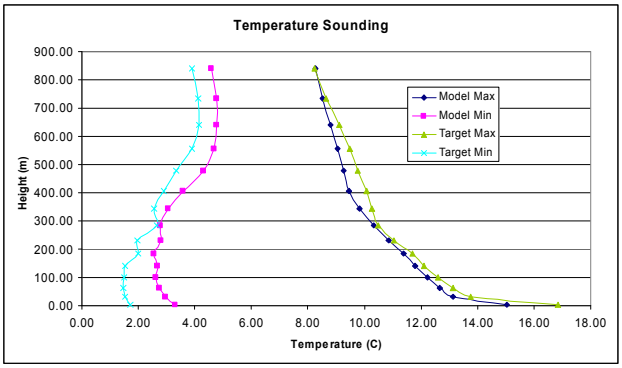


Figure 4 Maximum and minimum temperature soundings at SLC after 93 generations.

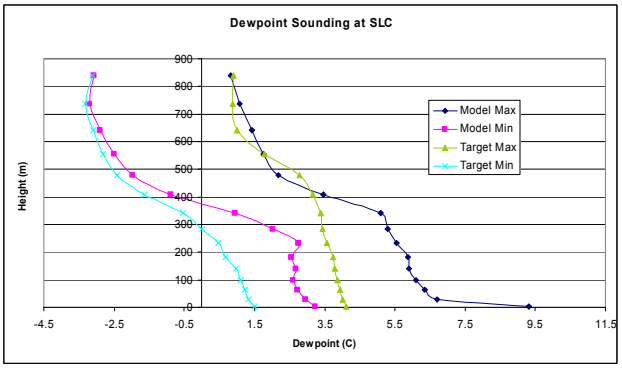


Figure 5 As in Fig. 4 but for dewpoint temperature.

In another experiment, the variance of the child parameters were perturbed, not according to the score of the best child simulation, but by the number of 'successful' children per generation or 'success rate'. Successful children are defined here as those simulations which generate a score lower than the previous generation's best score (their parent). For this optimization, the EP algorithm maintains a record of the number of successful children from the previous three generations (for a total of 96 children). If that number falls below a preset value (in this case, 3), the variance is reduced by 10% from the current value. The reasoning behind this approach is that very low success rates (over several generations) imply that the current parent resides in a minimal region of the response surface with respect to the parameter perturbation variance. Since the search was initiated with a large variance, we assume that the parent is in a global (or at least, a deep) minimum. Therefore the variance should be reduced so that solutions far beyond that point will not be explored. This systematically localizes the search, forcing the score deeper into the minimum.

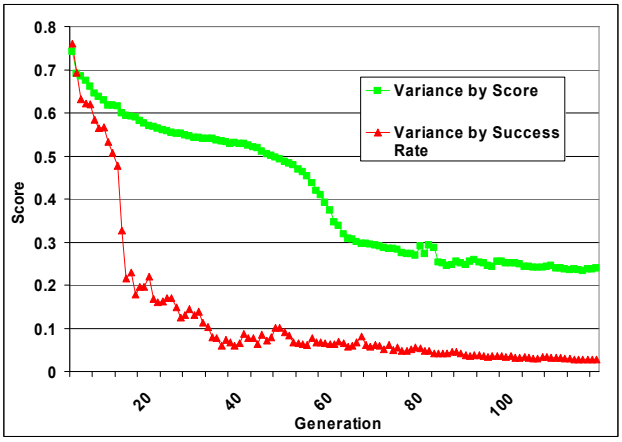


Figure 6 Best score for the model runs for the experiments with variance reduction based on score and success rate.

This approach proved to be very effective as can be seen in Figure 6 where, for comparison, the score for the first experiment is shown again. The normalized score falls very rapidly, reaching 0.2 in only 15 generations. The detailed behavior of the normalized scores is interesting. Both exhibit a rapid drop initially to about 0.5 and then a precipitous fall to the 0.2 – 0.25 level. This is followed by a slow, steady decline to the final score of approximately 0.04. The difference is in the number of generations required to achieve these score levels, in particular the 0.5 and 0.2 - 0.25 levels. Early in the optimizations, before any manual variance adjustments in the first experiment, there are significant differences in optimization efficiency. The optimizations that maintain larger parameter perturbation variances exhibit faster declines in score with generation number. Variance reduction based on child success rate is clearly the best method of those examined here. The success rate method depends on the choice of a cut-off success rate and an averaging period, but reasonable values for both of these are fairly easy to select. The variance scale factor holds steady at 1.0 until generation 18. At this point, the number of successful children from the previous 3 generations has fallen below 3, indicating that the optimization is closing in on a minimum and the search needs to become more localized. The variances of the child parameter perturbations are therefore reduced. This process is repeated numerous times, yielding a sequence of 10% reductions in the variance at various intervals. The process reaches an asymptote of 0.04 at about generation 85, much sooner than the previous experiment.. because it had the best result, all subsequent analysis is for the second experiment.

Fig. 7 shows that the initial temperature profile is too cool throughout the column relative to the target. After the 117th generation is run, the model profile matches the target profile very closely. The same is true for the wind direction – the model and target profiles differ by as much as 8º below 1500m and over 10º above this level. After 117 generations are run, the two differ by little more than 1º, matching through several wind shifts as the height increases from the surface to 3km aloft.

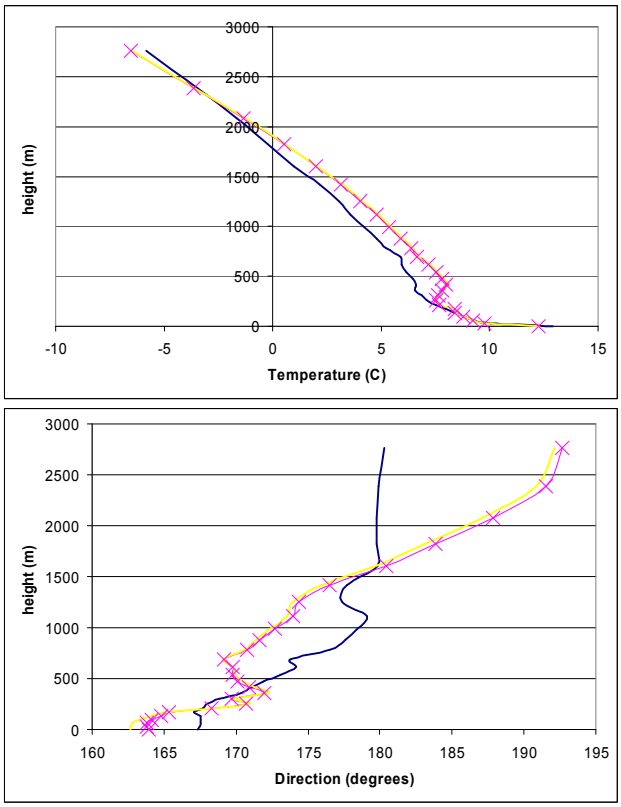


Figure 7 Temperature (top) and wind direction (bottom) soundings on 1800 UTC 26 October at the SLC station for EXP3. The curves are for the 1st generation (solid blue line), the 117th generation (yellow dashed line), and the target value (pink crossed line).

The synthetic Salt Lake City sounding was the only data included in the cost function. How will the rest of the model domain respond? Surface plots of temperature, dewpoint and wind vectors for grid 1 in Figs. 8-10 reveal the effect that optimizing the model with one vertical profile at SLC has on the entire domain. We see that the errors in temperature, dewpoint and wind are all significantly reduced after optimization

In the initial run, the temperature field (Fig. 8) shows cool errors in California and a warm bias over most of Utah. The warm temperature bias appears to be a grid 2 effect and is probably caused by incorrect initial values of some of the grid-dependent parameters. Optimization removes this grid 2 bias completely. The cold bias still exists in California over the central Sierra Nevada but its magnitude has been reduced. A few new but small errors have been introduced in northern California.

A similar response to optimization occurs with the dewpoint (Fig. 9). An initial dry bias exists in the west, particularly in Arizona and Nevada. As with temperature, this is greatly reduced (though not eliminated) after optimization. The wind errors (Fig. 10) are large for the initial run, with the most serious problems along the west coast extending into southern California. After optimization, the errors are much smaller throughout the domain, and are mostly limited to Nevada and southern California.

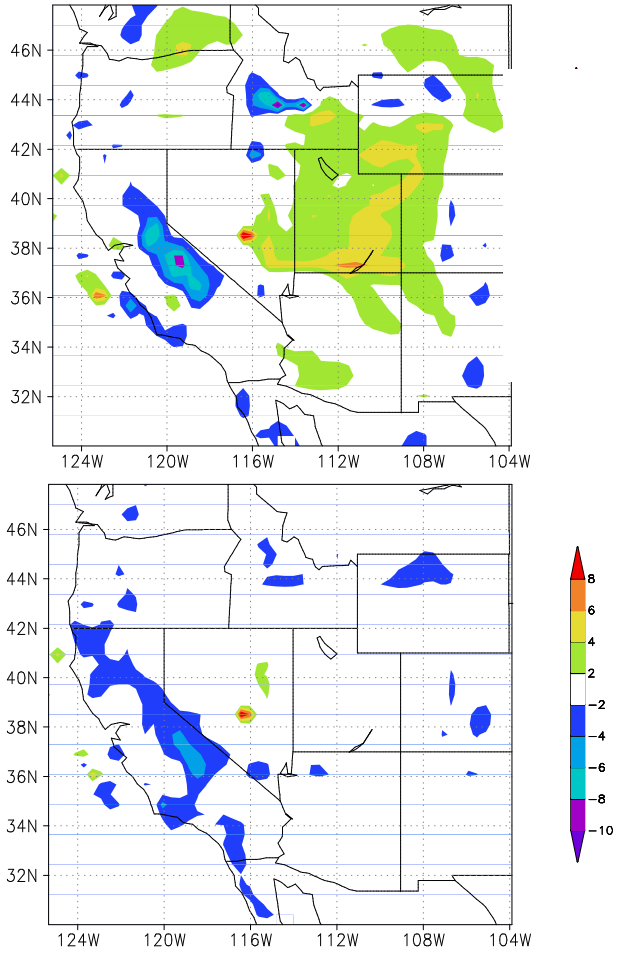

Figure 8 Surface temperature (ºC) at 26 UTC October 2000 for the original run (top) and the final run (bottom).

The fact that the cost function or score (based on a single vertical profile) approaches 0 does not guarantee that the optimized parameters converge to their target values. That this is so can be seen in the parameter evolution plots presented in Fig. 11. In these figures, the optimized parameters are normalized by their target values (if those values are non-zero) and plotted versus

generation number. Thus the target parameter values are either 0 or 1. Note that while the score does not monotonically decrease with generation number, a faster approach to 0 or 1 implies a more efficient optimization. These plots also reveal sensitivity of the model to different parameters since changes in parameters to which the model is very sensitive will result in large changes in

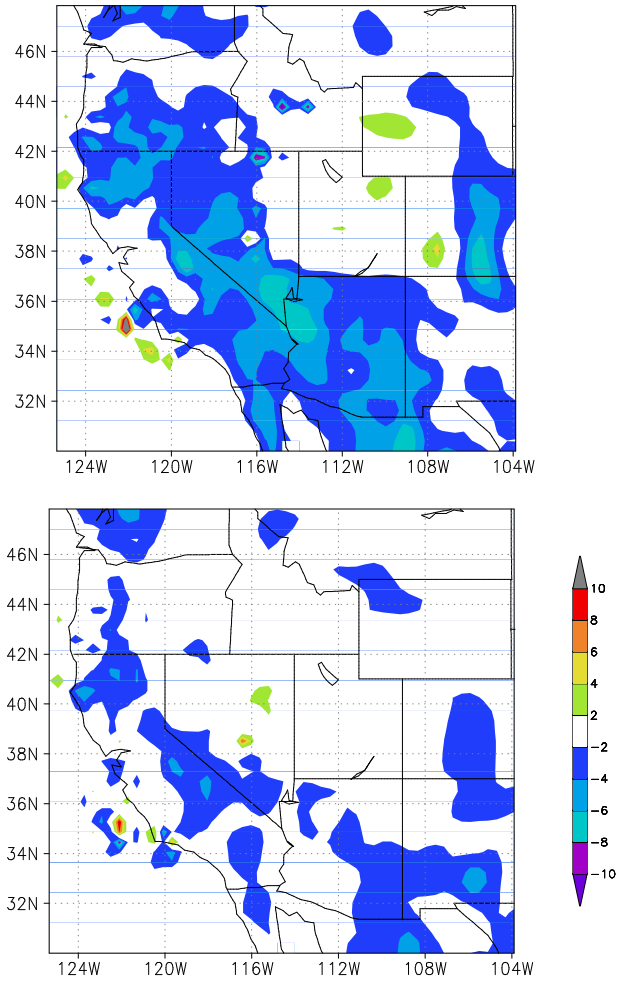


Figure 9 As in Figure 8 but for dewpoint temperature (ºC).

the model solution, and hence, the score. This will encourage the rapid adjustment of these parameters as the algorithm strives to improve the score as quickly as possible. Conversely, parameters that have little effect on the model score will feel little selection pressure and be allowed to drift more or less randomly. The evolution of the parameters being optimized can also be very insightful with respect to model deficiencies. Unusual or unrealistic parameter values after optimization can be an indication that the model does not have a parameterization that can adequately reproduce the observed meteorology (or that the parameterization has not been included in the parameter set for optimization).

Fig. 11 (top) shows the normalized parameter evolution for the grid-dependent topographic enhancement factors (TEF) employed in RAMS using the "envelope orography" option for grid 4. We see how the value varies but is drawn to unity with each generation. The TEF for grid 5 was added as a check (there was no grid 5) and as expected it wanders aimlessly (not shown).

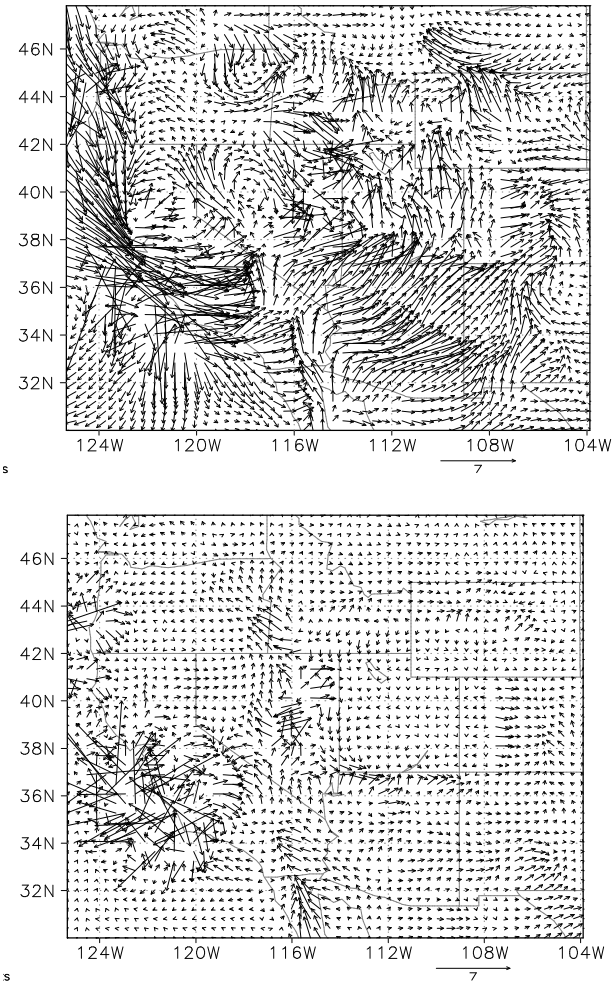


Figure 10 As in Figure 98 but for wind (m/s).

Certainly the most significant parameter in the optimization was the scale factor applied to the downwelling long-wave radiation used in the Pielke radiation parameterization. These results are shown in Fig. 11 (bottom). This parameter found its target value of 1 quickly and exhibited little variability about that value. This is not too surprising since low level temperatures and flows under stable conditions are dominated by the net radiation flux at the surface.

## 4. PHASE SPACE PATH

How does the EP algorithm act on the different variables, given that they are stochastically driven towards a predetermined target? We can explore this by

taking the time series of the 23 model parameters and using them to create a 23-dimensional generalized phase space. At the start of the experiment, the initial state of the model will occupy a specific position in this space, and with each subsequent generation follow a path through this space until the last generation, which will ideally be near the target.

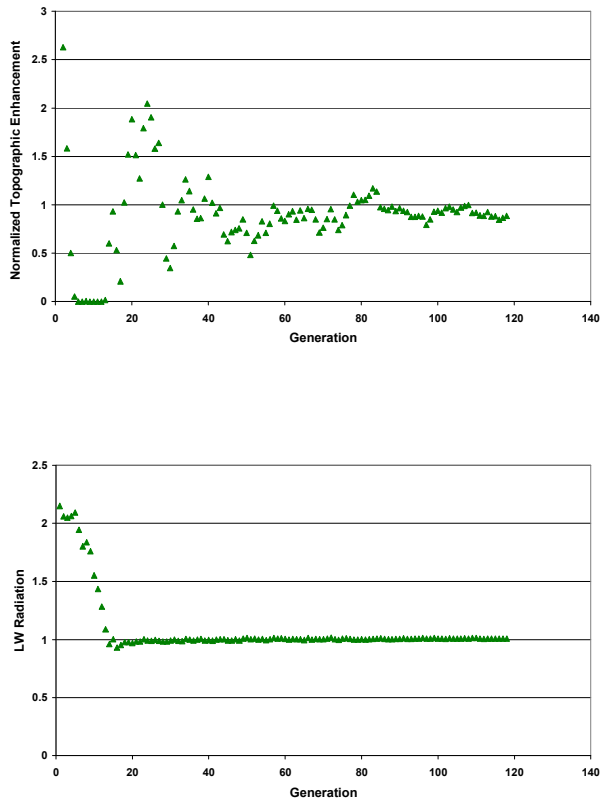


Fig. 11 Evolution of model parameters as a function of generation for the grid 4 topographic factor (top) and the Pielke longwave factor (bottom).

At each point in phase space, a unique score can be assigned. The way that the score varies with position will determine the path - the tendency will be to go 'downhill', avoid going 'uphill' but drifting in directions in which the score varies little.

To create a phase space that best describes the changes in the 23 parameters, a principle component analysis is applied. By treating each of the 23 parameters as a time series (actually a generation series), we apply a PC analysis to the data to create a series of eigenvectors, each of which represents a set of values for each of the 23 parameters. Since the values of the different parameters vary over different orders of magnitude, each time series is normalized by its standard deviation before applying the PC analysis.

$$x(i,t) = \sum_{e=1}^{23} PC_e(t) V_e(i)$$

Where x(i,t) if the value of parameter i at time t, $V_e(i)$ is the value of i for the $e^{th}$ eigenvector, and $PC_e$ is the value of the $e^{th}$ principal component at time t. In this way, the eigenvectors will better describe the motion through the phase space.

The 1st eigenvector (not shown) is dominated by the vegetation fractions and the nudging times, while the second is dominated by soil moisture. The third is controlled by the topographic factors. Note that the first three eigenvectors contain contributions from all model parameters, so it cannot be determined from this analysis that any are unnecessary.

The low-order principle components demonstrate variability on various timescales (Fig. 12, top). We see that PC1, which explains most of the motion, is characterized by a fairly steady progression until generation (g) =25, following which it remains steady. PC2, however, decreases downward until g=25, when it starts climbing until it too reaches an equilibrium, this time at g=37. PC3 oscillates until about g=61, following which it too starts moving towards its final value. These motions can be seen better in a plot showing 2-D slices of the path in the 23-d space (Fig. 13). We see than PC1 moves steadily from -11 to 1. PC2 falls during this time, but once PC1 reaches 1, PC2 starts to rise, giving us the inverted sawtooth pattern seen in the plot. The path in the PC2-3 space (not shown) forms a spiral pattern until PC2 reaches its minimum of -6, when it follows a straighter path to the final point.

PC4,5,6 follow similar patterns (Fig. 12 bottom) - they oscillate widely until PC1-3 reach equilibrium (at g=61), then they start making their way towards their final values. This can be interpreted as motion parallel to one PC until it gets close to the target value, during which time the motion in the higher order directions is more random. After the dominant PC hits its target, however, motion commences in the direction of another eigenvector.

What can we infer about the shape of the score in the parameter space, based on the single trajectory we have? For that matter, why does the trajectory take the path that it does? After all, a straight line from the initial point to the target would represent the shortest route. Why doesn't the EP algorithm follow such a line to reduce the score as rapidly as possible? Are there complicated nonlinearities in the space, forcing it to follow a convoluted path?

One likely reason is that the sensitivity of the score to the position in the parameter space is not isotropic, but varies with direction. For example, assume an idealized scenario in which we have only 2 parameters, and the score varies symmetrically with distance from the 'target' (the center) (Fig. 14, top). Start with a position in the lower left corner as indicated, and apply an algorithm that

finds the lowest point within the vicinity of the current solution, and makes that point the new solution before starting the process again. We see that the point will indeed move straight to the target. Now, assume that the score varies asymmetrically, so that the gradient in 1 direction is different from that in another (Fig. 14, bottom). We see now that as the solution follows the direction of greatest slope, it will follow a curved path, leaning towards 1 direction at first, then another (similar to the behavior of the PCs in Fig. 12).

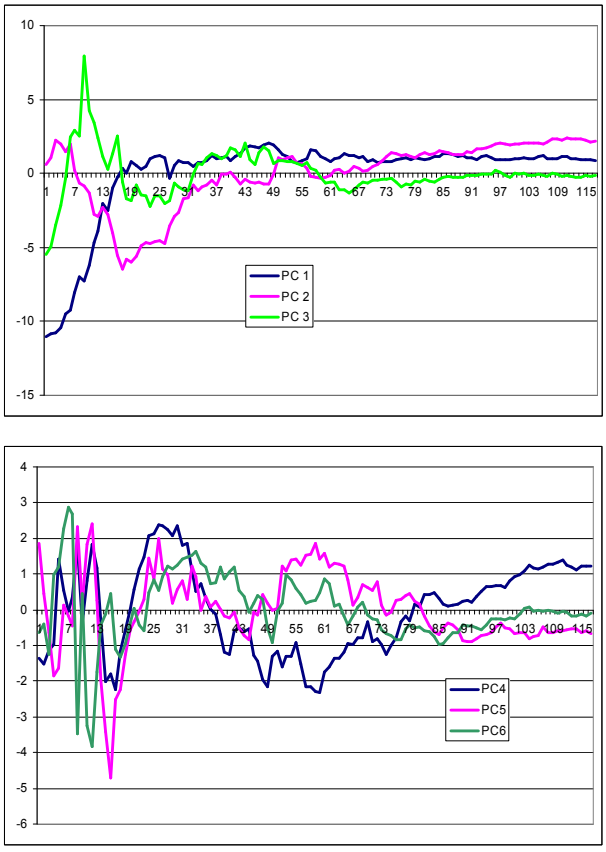


Figure 12 Principal component magnitudes

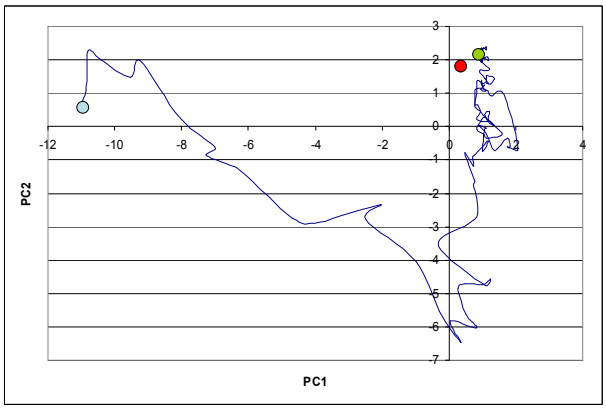


Figure 13 Scatter plot of PC1 vs. PC2. The blue dot is the starting point, the green dot is the final point, and the red dot is the target.

## 5. CONCLUSIONS

A simple evolutionary algorithm employing only a mutation operator is capable of optimizing a large set of model parameters for a meteorological simulation based on agreement with a single vertical profile of synthetic observational data (wind speed, wind direction, temperature and dewpoint). These synthetic data were derived from a separate meteorological simulation and included no measurement error. Thus the model is capable of reproducing these 'observations' exactly. For the optimization studies performed here, the cost function (based on a weighted RSM error) was reduced by 96-99%. The results presented here suggest that all possible data types should be included in the optimization if they are deemed reliable. Even though a particular data type (e.g. dewpoint) may be considered less important with respect to the application at hand, it can still be critical to finding the global optimum.

An analysis of the system evolution reveals a complicated response surface that describes how the model score responds to changes in the parameters. This could possibly be exploited by applying the EP process to a subset of the 23 eigenvectors instead of the full set of 23 parameters, reducing the amount of computation necessary to find the optimal solution.

Assuming that the optimization efficiency for real observations is comparable to that for synthetic data, operational implementation should be possible. Operational implementation will likely be based on a semi-continuous optimization in which a new generation is produced each day, using a target that is updated each day and forcing the parameters to remain up to date in the face of changing conditions (such as the transition from winter to spring). Whichever child best predicted today would serve as the model to produce tomorrow's forecast and also become the parent for the next generation. By periodically updating the 'target', the simulations will always be near optimal as conditions change with time.

This also mitigates the problem of the parameters being fixed at incorrect values, since, even if these give a good solution initially, as time goes on the errors will become more apparent and the algorithm will be forced to adjust them away from the incorrect values to improve the simulation.

This optimization procedure should be very efficient since it is incremental in nature, i.e. the initial guess will be at least near optimal and changes in the optimal parameter set should be small if the update cycle is frequent enough. It should be very interesting to see how the optimized paramters evolve with time and explore the reason for the changes.

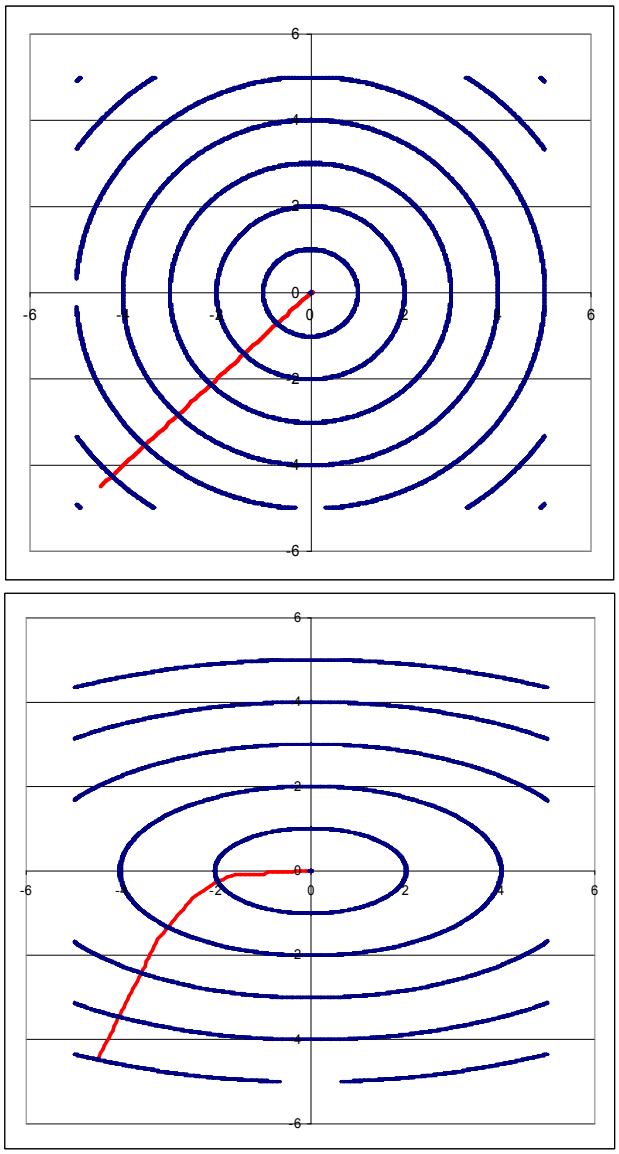


Figure 14 Phase space paths (red lines) for idealized response surfaces (blue lines represent the 'score' isolines). See text for details.

## REFERENCES


Doran, J.C., J.D. Fast, and J.Horel, 2002: The VTMX 2000 campaign. *Bull.Amer.Meteor*. Soc., **79**,265-283.

Mellor G. L. and T. Yamada, 1982: Development of a turbulent closure model for geophysical fluid problems. *Rev. Geophys. Space Phys*. **20** (10), 851-875.

Pielke R. A., W. R. Cotton, R. L. Walko, C. J. Tremback, W. A. Lyons, L. D. Grasso, M. E. Nicholls, M. D. Moran, D. A. Wesley, T. J. Lee, and J. H. Copeland, 1992: A comprehensive meteorological modeling system--RAMS. *Meteor. Atmos. Phys.*, **49**, 69-91.

Walko, R.L., L.E. Band, J. Baron, T.G.F. Kittel, R. Lammers, T.J. Lee, D. Ojima, R.A. Pielke, C. Taylor, C. Tague, C.J. Tremback, and P.L. Vidale, 2000: Coupled Atmosphere–Biophysics–Hydrology Models for Environmental Modeling. *J. Appl. Meteor.*, **39**, 931–944.

Zhong,S. and J.D. Fast, 2003: An evaluation of the MM5,RAMS and Meso-Eta models at subkilometer resolution using VTMX field campaign data in the Salt Lake valley. Mon. Wea. Rev., **131**, 1301-1322.


## ACKNOWLEDGEMENTS


This research was supported by the National Nuclear Security Administration of the Department of Energy.

This research used resources of the National Center for Computational Sciences at Oak Ridge National Laboratory, which is supported by the Office of Science of the Department of Energy under Contract DE-AC05-00OR22725.