Niched Co-Evolution Strategies to Address Non-uniqueness in Engineering Design

Emily M. Zechman
CB 7908, North Carolina State University
Raleigh, NC 27695
emzechma@ncsu.edu

S. Ranji Ranjithan
CB 7908, North Carolina State University
Raleigh, NC 27695
ranji@ncsu.edu

Li Liu
CB 7908, North Carolina State University
Raleigh, NC 27695
lliu2@ncsu.edu

ABSTRACT

When solving real engineering design problems, exploring the decision space for alternative and near-optimal solutions, and identifying non-unique solutions are useful. An efficient way to conduct this search is to identify alternative solutions that are as far apart as possible in the decision space. This will help not only provide meaningful design choices, if any are available, but also assess the degree of non-uniqueness present in the problem.

This paper describes the development and testing of a new evolutionary algorithm that extends evolution strategies to generate simultaneously a set of maximally different alternatives that are efficient solutions for engineering design problems. The new method, niched co-evolution strategies (NCES), is demonstrated for a groundwater pollutant source characterization problem that is known to have a high degree of non-uniqueness. Through an illustrative case study, NCES is shown to successfully assess and address the non-uniqueness for an instance of this relatively complex engineering design problem. Track: Evolution Strategies, Evolutionary Programming

General Terms

Algorithms, Design

Keywords

Evolution Strategies, Alternatives Generation, Engineering Design, Non-uniqueness

1. INTRODUCTION

Engineering design problems are solved by synthesizing domain knowledge of a particular system and identifying design choices that meet system constraints and requirements. Mathematical representations of system objectives, design choices, and constraints enable the use of mathematical search methods effectively in the design process. Many engineering systems require the execution of a complex simulation model to evaluate potential designs. Evolutionary algorithms (EA) can easily incorporate simulation models in solution evaluations and have been useful for design optimization in many engineering applications.

Engineering design problems often involve a large set of design choices, with complex interactions among decision variables. Further, the fitness landscapes for these problems are often complex and highly irregular. In addition to identifying the optimal solution, the design process may benefit if alternative and near-optimal solutions are generated. In such an exploration, solutions that are maximally different from each other, or far apart in the decision space from other good solutions, would provide more meaningful design alternatives. This would also help provide insight about the design problem and its solution space.

System identification is a class of design problems in which the characteristics of a system are determined using an inverse search technique. The goal is to identify system characteristics that will explain observations of the system. These problems often consist of a high degree of non-uniqueness, i.e., several solutions would result in an equally good match for the observations. Thus a single solution, the optimal solution, may provide misleading or insufficient description of the system. Exploring the decision space for maximally different system characteristics that provide similarly correct prediction of the observations would impart some understanding of the degree of non-uniqueness in the problem. For example, if many alternative system characteristics are found to describe the observations equally well, then one could conclude that the problem has more non-uniqueness than when no good alternatives are identified even though the decision space is effectively explored by a search for different solutions. Such an examination of non-uniqueness helps ascertain confidence in the prediction of the system characteristics. Subsequently the non-uniqueness could be resolved through additional confirmatory observations until no significantly different solution can be found.

A systematic approach to generating alternatives should seek solutions that are as different as possible in the decision space while performing well with respect to the stated fitness. The search for a set of maximally different alternative solutions has been formalized in [1], as the *modeling to generate alternatives* approach. A systematic exploration generates a small number of alternative solutions that perform similarly well and are located far apart in the decision space. An array of studies (e.g., [2], [3], [4], [5], [6], [8], and [9]) report the development of several alternatives generation procedures and their application to a number of realistic problems. In these studies, alternatives generation capabilities are used to address the presence of subjective preferences that cannot be included in mathematical objective statements, and to enhance the search for creative solutions for design problems.

Evolution strategies (ES) have been lauded as particularly useful for continuous parameter optimization problems [12]. ES were originally introduced for problem-solving for civil and mechanical engineering design problems and have efficiently optimized numerous engineering design problems [11]. Given the frequent use of ES in engineering design problems, which could benefit from exploring for design alternatives, the formal alternatives generation procedure is extended to ES in this paper. A new ES-based alternatives generation procedure is developed and tested. It is also successfully demonstrated for addressing the non-uniqueness present when solving a groundwater contamination source char acterization problem.

2. METHODOLOGIES FOR GENERATING MAXIMALLY DIFFERENT ALTERNATIVES

2.1 Mathematical Background for Generating Alternatives

The mathematical definition of modeling to generate alternatives for a search problem has been provided in [1]. An engineering design problem can be represented as:

$$\text{Maximize } Z = f(X) \tag{1}$$

$$\text{Subject to } g_i(X) \geq b_i \forall i = 1, \ldots, M \tag{2}$$

where $f(X)$ is the function representing the objective statement, which is maximized, X is the vector of decision variables representing the design decisions, and Eqn. 5 represents the set of constraints. Let X^* be the best design identified, and Z^* is the corresponding maximum objective value. An alternative solution that is maximally different from X^* can be generated by solving the following model:

$$\text{Maximize } D = d(X, X^*) \tag{3}$$

$$\text{Subject to } f(X) \geq T(Z^*) \tag{4}$$

$$\text{Subject to } g_i(X) \geq b_i \forall i = 1, \ldots, M \tag{5}$$

where D is a difference function based on $d(X, X^*)$, which represents a "distance" measure between two solutions X and X^*, and T is a target that is specified in relation to the fitness value Z^*. T represents an allowable relaxation, if any, in the objective value.

The target, T, permits a small degredation in the objective value to provide exploration of very different design decisions. The target may be specified to allow no relaxation when exploring for alternate optima for highly non-unique problems. For other problems, a near-optimal solution may be acceptable, and the target may be set so that the search identifies solutions that are inferior to the optimal solution.

A set of alternatives can be identified in a sequential approach. Once the best solution X^* is found, the first alternative is identified by solving the model represented by Eqns. 3−4. To identify the second alternative, the difference function D can be modified to find the solution maximally different from both the best solution (X*) and the first alternative, while the target function remains the same. The difference function can be updated for each new alternative identified, and the search for new alternatives continues until no significantly different alternatives are found. Several algorithms have been designed for generating a sequence of maximally different alternative solutions to numeric optimization problems, based on mathematical programming search methods, including linear programming, nonlinear programming, integer/binary programming, and dynamic programming (e.g., [2] and [4]).

2.2 EA-based Approaches for Generating Alternatives

Several (EA)-based approaches have been developed for generating a set of alternatives. The most direct method is to use an EA sequentially. Initial execution of the algorithm identifies the best solution to the original modeled problem (Eqn. 1). For each sequential run, Eqn. 3 is used as the objective statement and the solution identified in the previous run is included in the distance function. This method is described and implemented for a GA in [5]. This approach may become computationally burdensome due to the repeated executions of an EA.

The niching operator [7] is a well-established EA-based approach for identifying a set of solutions for multimodal problems. Niching is designed to identify a set of peaks in a fitness landscape by derating the fitness function in regions of decision space where solutions are already located. Niching is useful for identifying a set of solutions that all perform well for the objective statement, but is not designed for specifically identifying the solutions that are most different from one another. Niching can be modified to generate a smaller number of more distant solutions by tweaking parameters, such as the niche count and the sharing distance, as explored in [6]. This approach involves extensive parameter tweaking or a priori knowledge of the decision space. In the context of ES, niching has been investigated to a limited extent [13].

In [6], a new GA-based procedure (GAMGA - Genetic Algorithms for Modeling to Generate Alternatives) extends niching to generate maximally different solutions. Niches are formed around the most different solutions in a population, and a neighborhood is defined to restrict migration between niches. The algorithm introduces several additional parameters and algorithmic steps that require careful tuning.

An EA-based approach, the Evolutionary Algorithm for Generating Alternatives (EAGA) [8], was developed to explicitly identify maximally different alternatives using a set of subpopulations. The first subpopulation converges to an optimal solution to the original modeled problem. The remaining subpopulations search for solutions that are distant from all other subpopulations while meeting the specified target on the objective. The target objective becomes more restrictive as the search progresses, tightening as the fitness of the individuals in the first subpopulation improves. Recombination, mutation, and selection operators are applied separately in each subpopulation, and migration is not allowed between subpopulations. EAGA is designed with a minimal number of additional parameters. The fitness definition of an alternative solution can be defined so that the search, though parallel, identifies alternatives in a sequential manner. The advantage of the simultaneous search capability provided by the subpopulations is that more exploration is allowed for the secondary subpopulations and they avoid local optima, or solutions that are only marginally different from other subpopulations but meet the target objective. As the structure of EAGA is independent of the search procedure employed in the subpopulations, it can be used with genetic algorithms or evolution strategies for numeric optimization problems.

3. NICHED CO-EVOLUTION STRATEGIES

Niched Co-Evolution Strategies (NCES) uses the basic concept of cooperative co-evolution to evolve a set of design solutions. A set of subpopulations is used to collectively search for different alternative solutions, where each subpopulation is guided toward a region in the solution space that is distant from other subpopulations. Information about the location of a subpopulation in the solution space (and therefore the set of common solution-characteristics of a subpopulation) is shared such that the subpopulations cooperate in co-evolving toward different regions of the solution space. Selection within each subpopulation depends upon how well the solutions perform with respect to the design optimization model (Eqns. 1-2), as well as upon how far they are from the other subpopulations. NCES is designed to search explicitly for a set of solutions that are as different as possible in the design choices and are within a target for the objective function (Eqn. 4). Building upon the procedure described for a genetic algorithm in [8], EAGA is extended to contruct a new algorithm NCES for generating alternatives using evolution strategies. The main steps of the algorithm are described below for a $(\mu + \lambda)$ ES, but could be easily adapted for any alternative ES configuration.

3.1 Algorithmic Steps

Step 1. Create an initial population with P subpopulations (each with a population size of K), where P is the number of alternative solutions being sought. Let SP_p (p=1,..., P) represent the index for subpopulation p. The first subpopulation (SP_1) is dedicated to the search for the solution with the greatest objective function value.
Step 2. In each subpopulation SP_p (p=1,..., P), apply an adaptive mutation operator to generate λ offspring.
Step 3. In SP_1, evaluate the fitness (Eqn. 1), of each solution, and identify the solution in the subpopulation with the best fitness. This solution will serve as the benchmark for setting the relaxation constraint Eqn. 4.
Step 4. In SP_p (p=2,..., P), evaluate the fitness of each individual solution. Solutions that meet the target constraint Eqn. 4 are assigned a feasible flag, and solutions that fail to meet the target are labeled infeasible.
Step 5. For each solution k in subpopulation SP_p ($p \neq 1$), calculate the difference $D^{k,p}$ (defined in Section 3.2) in the solution space between that solution and other subpopulations.
Step 6. In each subpopulation SP_p, apply a selection operator. In SP_1, the selection is based on how well a solution maximizes the fitness. The solutions are ranked based on fitness and the top μ solutions survive to the next generation. In SP_p ($p \neq 1$), the selection is based on how well the solution meets the constraint Eqn. 3, as well as on the value of the difference function (as described in Section 3.2). Feasible solutions are ranked first from highest to lowest difference function. Infeasible soltuions are then ranked from best to worst fitness values.
Step 7. Check for termination criteria. Stop the algorithm if termination criteria (e.g., a maximum number of iterations) are met. Otherwise, go to Step 2.

3.2 Definition of Difference

The difference function for a solution is based on the distance of that solution to a set of subpopulations. For a numeric search problem, the distance d between two vectors of numbers may be easily represented, for example, as the Euclidean distance. The difference function, $D^{k,p}$, for solution k in subpopulation SP_p is the minimum of the distances between solution k and the other subpopulations SP_q, $q \neq p$. The distance from a solution in one subpopulation to another subpopulation is defined as the average of the distances between that solution and each solution in the other subpopulation, which is a representation of the centroid of the subpopulation. Thus $D^{k,p}$ is expressed as:

$$D^{k,p} = Min\{\frac{\sum_{j=1}^{K} d(X^{k,p}, X^{j,q})}{K}; q = 1, \ldots, P, q \neq p\} \quad (6)$$

where $d(X^{k,p}, X^{j,q})$ is the distance between the two solutions $X^{k,p}$ and $X^{j,q}$, K is the number of solutions in a subpopulation, and P is the number of subpopulations.

The *set difference* is a metric defined to evaluate the degree of difference among alternatives. The set difference can be represented as the minimum difference function value amoung a set of alternatives.

4. NCES FOR TEST PROBLEM

A test problem is used to demonstrate the use of NCES for generating a set of alternative solutions. The test function is a series of decreasing and increasing peaks (Fig. 1). For this maximization problem, the optimal solution is at $x = 0.015$, corresponding to $y = 1.0$. When the target relaxation is set at 75%, the two alternative solutions that are maximally different from the first solution and each other are located at $x = 0.97$ and $x = 0.77$. Both points correspond to the objective value $y = 0.75$, as shown in Fig. 1.

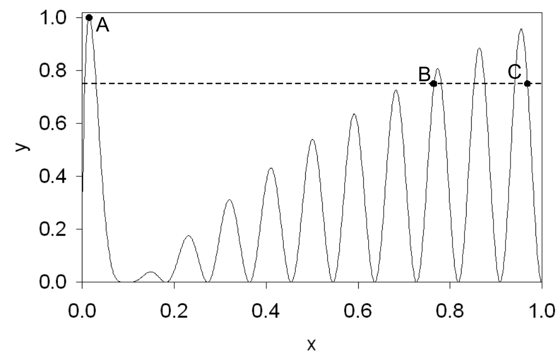

Figure 1: Test function. Dashed line represents the target relaxation of 75% of the best solution. A represents the best solution ($x = 0.015$), and B and C represent the set of two maximally different alternative solutions that meet the target relaxation.

NCES was implemented as described in Section 3, and executed to solve the test problem, using the algorithmic settings shown in Table 1. Twelve of 30 random trials successfully identified the three best solutions as shown in Fig. 1. Eleven trials misidentified the best solution in the first subpopulation as the peak at $x = 0.95$, but successfully converged to three maximally different peaks in the design space. The behavior of the algorithm is shown for one of the eleven successful trials in Fig. 2 over sixteen generations. The first subpopulation searched for the optimal solution, which is found by the forth generation. The second and third subpopulations converged to solutions that are distant from the first peak and from each other. By Generation 16, the second subpopulation converged to $x = 0.77$, and the third subpopulation converged to $x = 0.97$.

Table 1: NCES algorithmic settings for the test problem and the groundwater source characterization problem

Parameter	Setting for test problem	Setting for groundwater problem
μ	15	100
λ	100	200
$\triangle\sigma$	0.2	0.2
Generations	20	100
No. alternatives	3	5
Target	0.75	1.5

5. GROUNDWATER POLLUTION SOURCE IDENTIFICATION PROBLEM

In managing groundwater systems, the identification and characterization of a pollutant source is an important step toward the remediation of a contaminated aquifer. When contamination is detected in an aquifer, measurements of contaminant concentration in the groundwater are taken at observation wells that receive signals from the source of pollution. If a mathematical model of the groundwater system is available, the time series of observations can be used to identify the location and concentration of the contaminant source in the aquifer. A simulation-optimization approach can be used to identify the source characteristics by coupling a groundwater flow and transport simulation model and optimization method (Fig. 3). The optimization method searches for source location, size, and concentration characteristics that minimize the prediction error in matching the observed concentration data to the simulated concentrations.

Evolutionary algorithms have been used for source identification by [10] for relatively complex and realistic groundwater systems. Identification of an accurate source characterization is often complicated by the presence of non-uniqueness. Several different source characterizations may result in similar observed concentration profiles. A simplified example is used to demonstrate non-uniqueness in such a source characterization problem. A homogeneous groundwater field is modeled as a two-dimensional grid of 100 meters ($x-$direction) by 60 meters ($y-$direction), shown in Fig. 4. Observation Wells 1 and 2 are used to collect time dependent concentration data of the contaminant plume. Fig. 5 shows the concentration profile at Observation Well 1 resulting from an instantaneous point source contamination at Source 1. This figure also shows the concentration profile resulting from an instantaneous point source contamination at Source 2. The two concentration profiles are very similar, although the two sources vary considerably in location, size, and concentration values, which are listed in Table 2. The non-uniqueness present in the problem thus complicates attempts to characterize the source based on the observations

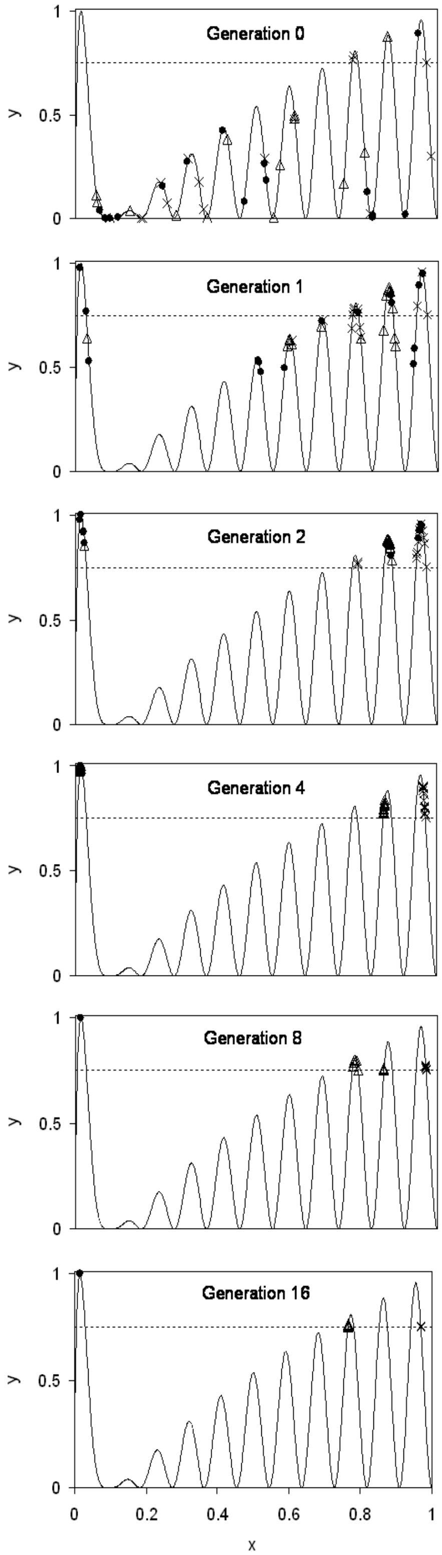

Figure 2: Convergence of NCES for the test function shown for a series of generations. The first subpopulation is identified by dark circles, the second subpopulation, by open triangles, and the third subpopulation, by $\times$'s.

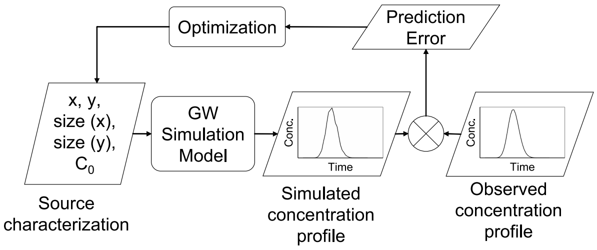

Figure 3: Source characterization for the groundwater pollutant problem using a simulation-optimization approach.

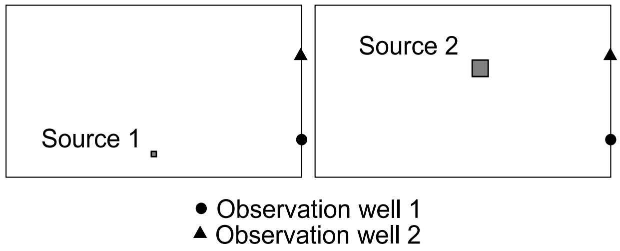

Figure 4: A groundwater contaminant characterization example.

taken at Observation Well 1.

A time series of concentration data may also be taken at Observation Well 2. Though at Observation Well 1 the concentration profiles from both sources are similar, the concentration profiles at Observation Well 2 from the two sources are highly dissimilar (as shown in Fig. 6). If observations from more wells are available, the non-uniqueness in the problem can be reduced and the source of the pollutant can be accurately identified through the use of additional information.

In the presence of non-uniqueness, a single-solution search may misidentify the source, resulting in a solution that has a low prediction error, but does not necessarily match the true source characteristics. A set of possible solutions will provide insight for the problem and a representation of the range of sources that could cause the observed contamination. At the outset of solving such source characterization problems and other inverse problems, it may not be possible to determine whether the problem is highly non-unique or a unique solution exists. An alternatives generation approach can be used to resolve this non-uniqueness issue.

The computational resources required to use a simulation-optimization approach for population-based searches may become impractical. To avoid excessive run-times for this investigation, a neural network approach was taken to identify a surrogate model for the finite element simulation model. A set of data was generated using the groundwater model and was used to train the neural network surrogate model. The neural network surrogate model provides quick approximations of the response surface of the model and is able to imitate the true groundwater model relatively accurately. A small error is introduced to the prediction error due to the imperfect approximation of the neural networks and contributes to the non-uniqueness present in the problem.

5.1 Identifying a Set of Alternative Source Characteristics using NCES

NCES is applied to a synthetic groundwater system management problem to generate a set of alternative solutions. A contaminant source is modeled using a groundwater model, and observations at a well are generated, which are treated as the set of synthetic observation data. NCES is coupled with the neural network surrogate model to identify the source characterization that matches this observation data. The groundwater field is modeled as a two-dimensional grid of 100 meters by 60 meters. The true source is the source characterization used as input to the model to generate the observation data. It is four meters by four meters, and the

Table 2: True Characteristics of Sources 1 and 2

Source. #	Loc. (x)	Loc. (y)	Size (x)	Size (y)	Conc. (mg/L)
1	53.0	9.0	2.0	2.0	34.0
2	57.0	39.0	6.0	6.0	100.0

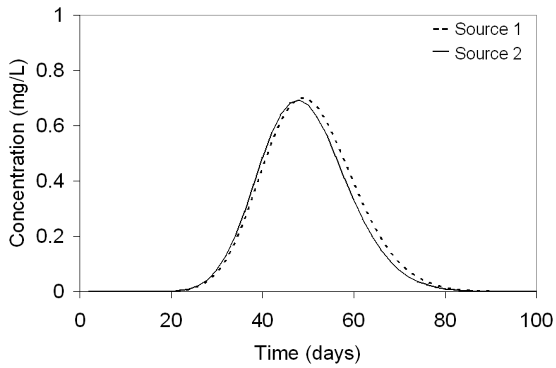

Figure 5: Concentration profiles at Observation Well 1.

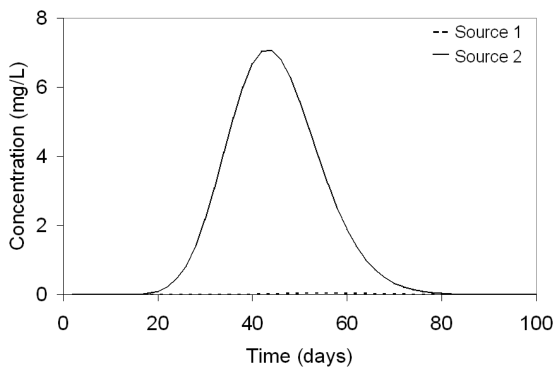

Figure 6: Concentration profiles at Observation Well 2.

centroid of the source is located at 54 meters on the x-axis and 34 meters on the y-axis. The concentration of the true source is 70 mg/L. Observations are simulated at twenty time-steps.

The settings for the algorithm are given in Table 1. A solution is represented as a vector of five real numbers between zero and one, and the allowable ranges of values for the decision variables are given in Table 3. The distance between two solutions is the sum of the normalized absolute differences between the $x-$location, the $y-$location, the size, and the concentration. NCES was executed to identify the best (optimal) solution and another four alternative solutions. The objective value of a solution is the prediction error, which is the maximum absolute difference between the observed concentration values and the predicted concentrations for the data set. The distance for a set of solutions, i.e., the set distance, is defined as the minimum of the difference measure between all pairs of the four alternative solutions in the set.

Table 3: Allowable ranges of decision variable values

Variable	Range
Location $x-$axis	0-80 m
Location $y-$axis	0-60 m
Size $x-$axis	2-8 m
Size $y-$axis	2-8 m
Concentration	0-100 mg/L

5.2 Results

5.2.1 Source Characterization Using Observation Well 1

NCES was executed for 30 random trials, and each trial generated five alternative source characterizations. The first subpoplulation searches for the solution that minimizes prediction error. The four remaining subpopulations identify solutions most distant from the first subpopulation and the other secondary subpopulations and with prediction errors within a specified range of the best solution in the first subpopulation. Eighteen of the 30 trials identified four alternative solutions with errors within the specified range of the best solution. The average error of the solutions identified by the first subpopulation is 0.007 with a standard deviation of 0.002.

A successful trial, one in which all solutions had prediction errors within the target relaxation and the set difference was relatively high, is randomly selected to view the behavior of NCES. As shown in Fig. 7, the prediction error of the best solution in subpopulation 1 decreases monotonically. As the first subpopulation converges, the secondary subpopulations (subpopulations 2, 3, 4, and 5) identify error values within 1.5 times the error of the best solution in the first subpopulation. The secondary subpopulations evolve toward distant parts of the decision space, as seen by the increasing trend of the distance measure for the best solution in the secondary subpopulations (Fig. 8). The final set distance for the five alternatives is 0.796. This number represents the difference among the solutions identified, which is visualized in the 2-D layout of the solutions Fig. 9. Table 2 lists the characteristics of each alternative and the true source characterization, as well as the prediction error associated with each alternative solution. As shown in Fig. 10, the concentration profile of each alternative solution matched the observations very closely, even with a generous target relaxation applied to the alternatives. Three of the five solutions are proximate to the true solution, with Alternative #4 appearing closest to the true solution. Had a single-solution search been conducted, the solution identified by the first subpopulation, Alternative #1, would be misleading as it is spatially farthest from the true solution. The generation of alternatives provides candidate source locations that, in this case, are closer to the true source.

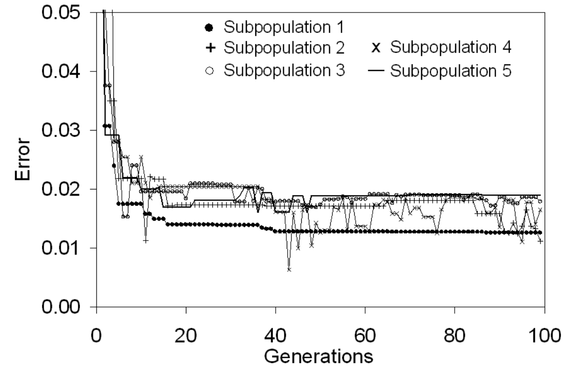

Figure 7: Convergence of the fitness of the best solutions in the five subpopulations.

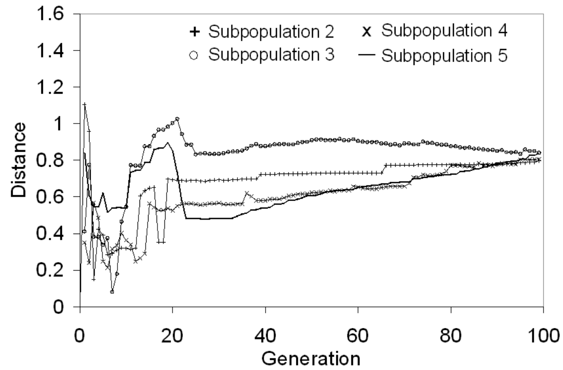

Figure 8: Evolution of difference metric for best solutions in subpopulations 2, 3, 4, and 5.

5.2.2 Source Characterization Using Observation Wells 1 and 2

NCES was applied again to solve the same source characterization problem based on observations from Observation Wells 1 and 2. Thirty random trials were executed to generate five alternative source characterizations. Nineteen of the 30 trials identified four alternative solutions with errors within the specified range of the best solution. The average error of the solutions identified by the first subpopulation is 0.013 with a standard deviation of 0.003. The average set distance over the 30 trials using Observation Wells 1 and 2 is 0.002. The average set distance over the 30 trials using Ob-

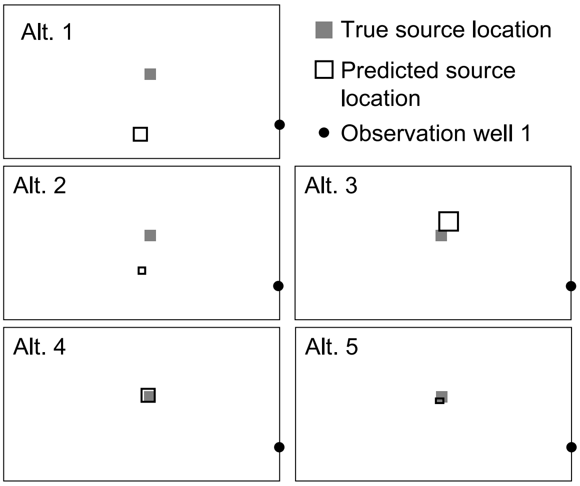

Figure 9: Five alternative source locations predicted using Observation Well 1.

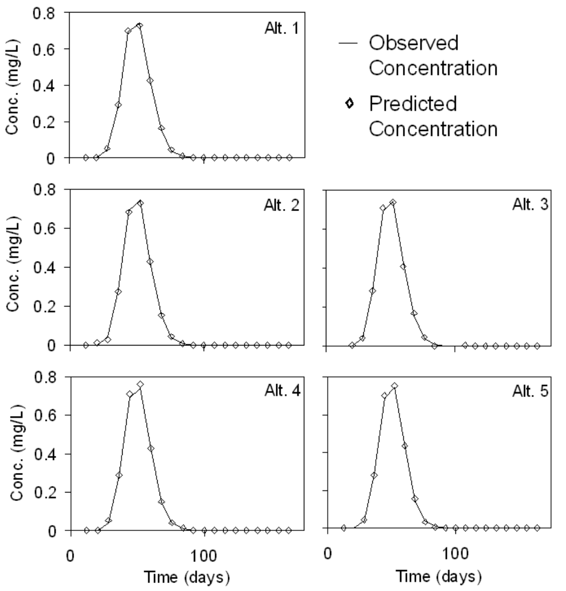

Figure 10: Concentration profiles of the five alternative source characteristics.

servation Well 1 (as discussed in Section 5.2.1) is 0.423. The amount of decision space that will yield similar prediction errors has decreased significantly by increasing the number of observations used to identify the source. By identifying relatively similar solutions as the set of maximally different solutions for the trials using two observation wells, NCES quantifies the decrease in non-uniqueness.

One successful trial of the thirty was chosen for analysis. The five alternative solutions generated had a set distance of 0.005, which is much lower than the set differences for the case where the source characterization was based only on Observation Well 1. Table 3 lists the characteristics of each alternative identified and the true source characterization. The five alternative sources are all located in approximately the same place in the $x - y$ space, as shown in Fig. 11.

While the inclusion of observations from the second observation well helps significantly resolve the location non-uniqueness, the sizes and concentrations of the alternatives still vary widely. Though this non-uniqueness remains unresolved, the resolution of the non-uniqueness in location is valuable and important for remediation purposes. This enables a more accurately targeted contaminant mitigation action, increasing the effectiveness and confidence of success of a remediation management strategy.

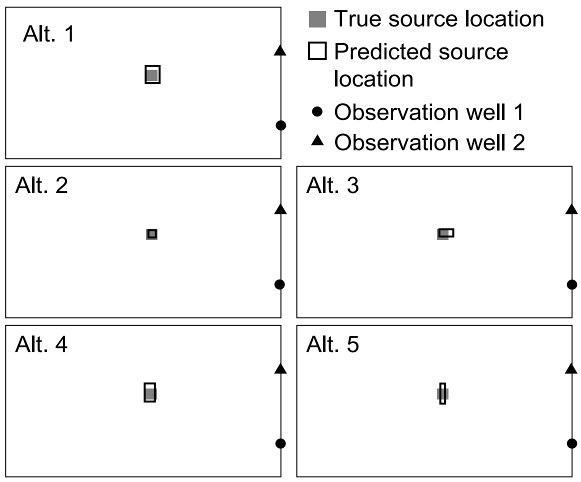

Figure 11: Five alternative source locations predicted using Observation Wells 1 and 2.

6. FINAL REMARKS

The new method NCES extends evolution strategies to generate simultaneously a set of maximally different alternative solutions to engineering design problems. Exploration of such alternatives with distinctly different solution characteristics would help identify meaningful design choices. For problems with non-uniqueness issues, NCES is useful to address non-uniqueness efficiently. NCES was tested and illustrated for a multi-modal test function. The method was also applied to a groundwater contaminant source characterization problem that involves an evolution strategies-based search. As non-uniqueness is a serious issue in such inverse problems, NCES was employed to successfully address it. Results demonstrate that when the problem contains a high degree of non-uniqueness due to limited observations, NCES

Table 4: Characteristics of the true contaminant source and the five alternative source characterizations identified by NCES

Sol.	Loc. (x)	Loc. (y)	Size (x)	Size (y)	Conc. (mg/L)	Error (mg/L)
True Source	54.0	34.0	4.0	4.0	70.0	–
Alternatives identified based on Observation Well 1						
Alt. 1	51.0	10.3	5.3	4.8	11.9	0.013
Alt. 2	50.6	20.0	2.0	2.0	32.5	0.011
Alt. 3	56.1	39.2	6.9	6.6	92.9	0.018
Alt. 4	54.1	34.0	5.4	4.6	53.1	0.016
Alt. 5	53.2	31.9	2.5	2.0	100.0	0.019
Alternatives identified based on Observation Wells 1 and 2						
Alt. 1	53.8	34.0	4.4	6.8	43.9	0.015
Alt. 2	54.0	33.9	6.1	5.2	42.9	0.023
Alt. 3	54.4	33.9	3.1	3.0	98.2	0.023
Alt. 4	54.2	34.1	7.1	2.0	69.7	0.023
Alt. 5	53.6	34.2	3.5	8.0	44.0	0.023

is able to identify a set of different solutions and indicate the presence of non-uniqueness. As more observation data are included, the results show that the non-uniqueness in location is resolved, although the size and concentration of the source are still non-unique. Knowing the location alone is useful in pinpointing the site for remediation action. While NCES is demonstrated for a specific application, it is a generically applicable method to any problem that is solved using evolution strategies. Use of NCES in an ES-based solution of engineering design problems is expected to be greatly valuable in exploring the decision space for alternative designs. By structure, NCES identifies designs that are as different as possible from each other with respect to the design choices included in the different solutions. This provides design innovation in engineering design problems and insight to the solution space.

7. ACKNOWLEDGEMENTS

The research leading to the reported results was supported partly by grants from the National Science Foundation (BES-0312841). The viewpoints presented here are those of the authors and do not necessarily reflect those of the funding agencies.

8. REFERENCES

[1] E. D. Brill, Jr. Use of optimization models in public-sector planning. *Management Science*, 25(5):413–422, 1979.

[2] E. D. Brill, Jr., S.-Y. Chang, and L. D. Hopkins. Modeling to generate alternatives: The HSJ approach and an illustration using a problem in land use planning. *Management Science*, 25(2):221–235, 1982.

[3] E. D. Brill, Jr., J. M. Flach, L. D. Hopkins, and S. Ranjithan. MGA: A decision support system for complex, incompletely defined problems. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(4):745–757, 1990.

[4] S.-Y. Chang, E. D. Brill, Jr., and L. D. Hopkins. Use of mathematical models to generate alternative solutions to water resources planning problems. *Water Resources Research*, 18(1):58–64, 1982.

[5] L. Harrell. *Methods for Generating Alternatives to Manage Water Quality in Watersheds*. PhD thesis, North Carolina State University, Raleigh, NC, USA, 1998.

[6] D. H. Loughlin, S. Ranjithan, J. W. Baugh, and E. D. Brill, Jr. Genetic algorithm approaches for addressing unmodeled objectives. *Engineering Optimization*, 33(5):549–569, 2001.

[7] S. W. Mahfoud. *Niching Methods for Genetic Algorithms*. PhD thesis, University of Illinois at Urbana-Champaign, Urbana-Champaign, Illinois, USA, 1995.

[8] E. M. Zechman and S. Ranjithan. An evolutionary algorithm to generate alternatives (EAGA) for engineering optimization problems. *Engineering Optimization*, 36(5):539–553, 2004.

[9] E. M. Zechman and S. Ranjithan. Multipopulation Cooperative Coevolutionary Programming (MCCP) for enhancing creative design. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 1641–1648. ACM, 2005.

[10] G. K. Mahinthakumar and M. Sayeed. Hybrid genetic algorithm - Local search methods for solving groundwater source identification inverse problems. *Journal of Water Resources Planning and Management*, 131(1):45–57, 2005.

[11] I. C. Parmee. Engineering Design and Evolutionary Algorithms. Springer, 2001.

[12] T. Back. Evolution Strategies: An alternative evolutionary algorithm. In *Selected Papers from the European conference on Artificial Evolution*, pages 3–20. Springer-Verlag, 1995.

[13] O. M. Shir and T. Back Dynamic niching in evolution strategies with covariance matrix adaptation. In *Proceedings of the Congress on Evolutionary Computation*, pages 2584–2591. IEEE, 2005.