



A Stochastic Nature-Inspired Metaheuristic for Modelling-to-Generate-Alternatives

Julian Scott Yeomans

OMIS Area, Schulich School of Business, York University, Toronto, ON, M3J 1P3 Canada

Abstract In solving many “real world” decision-making applications, it is generally preferable to formulate several quantifiably good alternatives that provide numerous, distinct approaches to the problem. This is because policy formulation typically involves complex problems that are riddled with incongruent performance objectives and possess incompatible design requirements that can be very difficult – if not impossible – to incorporate at the time supporting decision models are constructed. By generating a set of maximally different solutions, it is believed that some of the dissimilar alternatives will provide unique perspectives that serve to satisfy the unmodelled characteristics. This maximally different solution creation approach is referred to as modelling-to-generate-alternatives (MGA). This paper provides a stochastic biologically-inspired metaheuristic simulation-optimization MGA method that can efficiently create multiple solution alternatives to problems containing significant stochastic uncertainties that satisfy required system performance criteria and yet are maximally different in their decision spaces. The efficacy of this stochastic MGA approach is demonstrated on a municipal solid waste case study. It is shown that this new computationally efficient algorithmic approach can simultaneously produce the desired number of maximally different solution alternatives in a single computational run of the procedure.

Keywords Stochastic Modelling-to-generate-alternatives, Firefly Algorithm, Nature-inspired Metaheuristics

1. Introduction

“Real world” decision-making typically involves multifaceted stochastic problems that possess design components which are very difficult to incorporate into corresponding mathematical programming models and tend to be riddled with unquantifiable design specifications [2], [4], [5], [6]. While mathematically optimal solutions provide the best solutions to these modelled problems, they are generally not the best answers to the fundamental “real” problems as there are invariably unmodelled objectives and unquantified issues not apparent during model construction [3], [5]. Hence, it is generally considered desirable to generate a number of very dissimilar alternatives that supply completely distinct perspectives to the formulated problem [3], [5], [9]. These alternatives should possess near-optimal objective measures with respect to all known modelled objective(s), but be maximally different from each other in terms of their decision variable structures [4], [6]. Several approaches referred to as *modelling-to-generate-alternatives* (MGA) have been developed in response to this multi-solution creation requirement [1], [9].

In this paper, it is shown how to efficiently construct a set of maximally different solution alternatives by implementing a modified version of the computationally efficient Firefly Algorithm (FA) of Yang [7] in conjunction with a new stochastic MGA approach that employs simulation-optimization (SO) [9]. The MGA procedure provided in this study extends the earlier approach of Imanirad *et al.* [1] by employing a novel FA-driven SO approach to concurrently generate all of the desired number of solutions in a one-pass algorithm. Hence, this stochastic FA procedure is very computationally efficient from an MGA perspective. The procedure is demonstrated on a municipal solid waste (MSW) facilities expansion case study taken from Yeomans [8].



2. Firefly Algorithm for Function Optimization

While this section briefly outlines the FA procedure, more detailed specifications can be located in Yang [7] and Imanirad *et al.* [1]. The FA is a biologically-inspired, population-based metaheuristic with each firefly in the population representing a potential solution to the problem. An FA procedure employs three idealized rules: (i) All fireflies within a population are unisex, so that one firefly will be attracted to other fireflies irrespective of their sex; (ii) Attractiveness between fireflies is proportional to their brightness, implying that for any two flashing fireflies, the less bright one will move towards the brighter one; and (iii) The brightness of a firefly is determined by the value of its objective function. For a maximization problem, the brightness can simply be considered proportional to the value of the objective function. Yang [7] demonstrates that the FA approaches the global optima whenever the number of fireflies $n \rightarrow \infty$ and the number of iterations t , is set so that $t \gg 1$. In reality, the FA has been shown to converge extremely quickly into both local and global optima [1], [7]. The basic operational steps of the FA are summarized in the following pseudo code [7].

```

Objective Function  $F(\mathbf{X})$ ,  $\mathbf{X} = (x_1, x_2, \dots, x_d)$ 
Generate the initial population of  $n$  fireflies,  $\mathbf{X}_i$ ,  $i = 1, 2, \dots, n$ 
Light intensity  $I_i$  at  $\mathbf{X}_i$  is determined by  $F(\mathbf{X}_i)$ 
Define the light absorption coefficient  $\gamma$ 
while ( $t < \text{MaxGeneration}$ )
    for  $i = 1: n$ , all  $n$  fireflies
        for  $j = 1: n$ , all  $n$  fireflies (inner loop)
            if ( $I_i < I_j$ ), Move firefly  $i$  towards  $j$ ; end if
            Vary attractiveness with distance  $r$  via  $e^{-\gamma r}$ 
        endfor  $j$ 
    end for  $i$ 
    Rank the fireflies and find the current global best solution  $\mathbf{G}^*$ 
end while
Postprocess the results

```

3. Modelling to Generate Alternatives with the Firefly Algorithm

Notwithstanding other fundamental limitations, most mathematical programming techniques have focused almost exclusively upon producing optimal solutions to single-objective problem formulations or generating noninferior solutions to multi-objective problem instances. While such algorithms may determine optimal solutions to the derived complex mathematical models, whether their results actually establish “best” to the underlying real problems is certainly questionable [5], [6]. In most “real world” decision problems, there are numerous system objectives and requirements that are never explicitly apparent or included at the decision formulation stage [5], [6]. Furthermore, it may never be possible to explicitly express all of the subjective considerations because there are frequently numerous incompatible, competing, design requirements and, perhaps, adversarial stakeholder groups [5]. Therefore, most subjective aspects of a problem remain unquantified and unmodelled in the construction of the resultant decision models. This is a common occurrence in situations where the final decisions must be constructed based not only upon clearly stated and modelled objectives, but also upon fundamentally subjective, political and socio-economic goals and stakeholder preferences [5], [9]. Numerous “real world” examples of this type of incongruent modelling duality are described in Loughlin *et al.* [3] and Yeomans [8].

When unmodelled objectives and unquantified issues exist, different approaches are required in order to not only search the decision space for the noninferior set of solutions, but also to explore the decision space for *inferior* alternative solutions to the modelled problem. In particular, any search for good alternatives to problems known (or suspected) to contain unmodelled objectives must focus not only on the non-inferior solution set, but also necessarily on an exploration of the problem’s inferior region. To illustrate the implications of an unmodelled objective on a decision search, assume that the optimal solution for a quantified, single-objective, maximization decision problem is \mathbf{X}^* with corresponding objective value $Z1^*$. Now suppose that there exists a second, unmodelled, maximization objective $Z2$ that subjectively reflects environmental/political acceptability.



Let the solution X^a , belonging to the noninferior, 2-objective set, represent a potential best compromise solution if both objectives could somehow have been simultaneously evaluated by the decision-maker. While X^a might be viewed as the best compromise solution to the real problem, it would clearly appear inferior to the solution X^* in the quantified model, since it must be the case that $ZI^a \leq ZI^*$. Consequently, when unmodelled objectives are factored into the decision-making process, mathematically inferior solutions for the modelled problem can be optimal for the real problem. Therefore, when unmodelled objectives and unquantified issues might exist, different approaches are required in order to not only search the decision space for the noninferior set of solutions, but also to simultaneously explore the decision space for inferior alternative solutions to the modelled problem. Population-based procedures such as the FA permit concurrent searches throughout a feasible region and thus prove to be particularly adept methods for searching through a problem's decision space.

The primary motivation behind MGA is to produce a manageably small set of maximally different alternatives that are quantifiably good with respect to modelled objectives yet are as different as possible from each other in the decision space. In doing this, the resulting alternative solution set is likely to provide truly different choices that all perform somewhat similarly with respect to the known modelled objective(s) yet very differently with respect to any unmodelled issues. By generating these maximally different solutions, the decision-makers can explore alternatives that may satisfy the unmodelled objectives to varying degrees of stakeholder acceptability. Obviously, solution-setters must conduct a subsequent comprehensive comparison of the alternatives to determine which options would most closely satisfy their very specific circumstances. Thus, an MGA approach should be viewed as decision support rather than explicit solution determination.

In order to properly motivate an MGA search procedure, it is necessary to provide a more formal definition of the maximal difference goals of MGA [3], [8]. Suppose the optimal solution to an original mathematical model is X^* with objective value $Z^* = F(X^*)$. The following model can then be solved to generate an alternative solution that is maximally different from X^* :

$$\begin{array}{ll} \text{Max } \Delta & = \sum_i |X_i - X_i^*| \\ \text{Subject to: } & X \in D \\ & |F(X) - Z^*| \leq T \end{array}$$

where Δ represents some difference function (shown as absolute in this instance) and T is a tolerance target specified in relation to the original optimal function value Z^* . T is a user-supplied value that determines how much of the inferior region is to be explored for alternative solutions. The FA-based MGA procedure is designed to generate a small number of good but maximally different alternatives by adjusting the value of T and using the FA to solve the corresponding, new maximal difference problem instance. In this approach, subpopulations within the algorithm's overall population are established as the Fireflies collectively evolve toward different local optima within the solution space. Each desired solution alternative undergoes the common search procedure of the FA. The survival of solutions depends upon how well the solutions perform with respect to the modelled objective(s) and by how far away they are from all of the other previously generated alternatives in the decision space.

4. FA-driven Simulation-Optimization Approach for Stochastic MGA

The stochastic FA-directed simulation-optimization (SO) approach for efficiently generating sets of maximally different solution alternatives consists of two alternating computational phases; (i) an "evolutionary phase" directed by the FA module and (ii) a simulation module. As mentioned earlier, an FA maintains a population of candidate solutions throughout its execution. The evolutionary phase considers the entire population of solutions during each generation of the search and evolves from a current population to a subsequent one. The quality of each solution in the population is found by having its performance criterion, F , evaluated by simulation. Because of the system's stochastic components, all performance measures are statistics calculated from the responses generated in the simulation module. After simulating each candidate solution, the respective fitness values are returned to the FA module to be utilized in the creation of the next generation of candidate solutions. One primary principle of an FA is that fitter solutions in the current population possess a greater likelihood for survival and progression into the subsequent generation. The FA module evolves the system toward improved



solutions in subsequent populations and ensures that the solution search does not become fixated at some local optima. After generating a new candidate solution set in the FA module, the new population is returned to the simulation module for comparative evaluation. This alternating, two-phase search process terminates when an appropriately stable system state has been attained.

An obvious approach to generate alternatives would be to iteratively solve the maximum difference model by incrementally updating the target T whenever a new alternative needs to be produced. Such an approach would need to run n times in order to produce n alternatives. The new stochastic MGA procedure is designed to concurrently generate its maximally different alternatives in a single pass of the FA procedure (i.e. the same number of runs as if FA were used solely for function optimization purposes) and its efficiency is based upon the concept of co-evolution [1]. Pre-specified stratified subpopulation ranges within the FA's overall population are established that collectively evolve the search toward the formation of the stipulated number of solution alternatives. Each desired solution alternative is represented by each respective subpopulation and each subpopulation undergoes the common operations of the FA. The survival of solutions in each subpopulation depends upon how well the solutions perform with respect to both the modelled objective(s) and by how far away they are from all of the other solutions in the decision space. This forces each subpopulation to co-evolve toward good but maximally distant regions of the decision space [1], [9]. By employing the co-evolutionary concept, it becomes possible to implement an FA-directed stochastic MGA procedure that concurrently produces alternatives which possess objective function bounds. Namely, the algorithm need be run only a single time to produce its entire set of alternatives irrespective of the value of n . Hence, it is a very computationally efficient process. The steps in the FA-driven stochastic algorithm are as follows:

1. Create an initial population stratified into P equally-sized subpopulations. The value for P must be established *a priori* by the decision-maker. P represents the desired number of maximally different alternative solutions within a prescribed target deviation from the optimal to be generated. S_p represents the p^{th} subpopulation set of solutions, $p=1, \dots, P$ and there are K solutions contained within each S_p .

2. Evaluate all solutions in S_p , $p=1, \dots, P$, with respect to the modelled objective using the simulation module. Solutions meeting the target constraint and all other problem constraints are designated as *feasible*, while all other solutions are designated as *infeasible*.

3. Apply an appropriate elitism operator to each S_p to preserve the best individual in each subpopulation. In S_p , $p=1, \dots, P$, the best solution is the feasible solution most distant in decision space from all of the other subpopulations (the distance measure is defined in Step 6). Note: Because the best solution to date is always placed into each subpopulation, at least one solution in S_p will always be feasible.

4. Stop the algorithm if the termination criteria (such as maximum number of iterations or some measure of solution convergence) are met. Otherwise, proceed to Step 5.

5. Identify the decision space centroid, C_{ip} , for each of the $K' \leq K$ feasible solutions within $k=1, \dots, K$ of S_p , for each of the N decision variables X_{ikp} , $i=1, \dots, N$. Each centroid represents the N -dimensional centre of mass for the solutions in each of the respective subpopulations, p . As an illustrative example for determining a centroid, calculate $C_{ip} = (1/K') * \sum_k X_{ikp}$. In this calculation, each dimension of each centroid is computed as the straightforward average value of that decision variable over all of the values for that variable within the feasible solutions of the respective subpopulation. Alternatively, a centroid could be calculated as some fitness-weighted average or by some other appropriately defined measure.

6. For each solution $k=1, \dots, K$, in each S_q , calculate D_{kq} , a distance measure between that solution and all other subpopulations. As an illustrative example for determining a distance measure, calculate $D_{kq} = \text{Min} \{ \sum_i |X_{ikp} - C_{ip}|; p=1, \dots, P, p \neq q \}$. This distance represents the minimum distance between solution k in subpopulation q and the centroids of all other subpopulations. Alternatively, the distance measure could be calculated by some other appropriately defined measure.

7. Rank the solutions within each S_p according to the distance measure D_{kq} objective – appropriately adjusted to incorporate any constraint violation penalties. The goal of maximal difference is to force solutions from one subpopulation to be as far apart as possible in the decision space from the solutions of each of the other subpopulations. This step orders the specific solutions in each subpopulation by those solutions which are most



distant from the solutions in all of the other subpopulations.

8. In each S_p , apply FA “change operations” to the solutions and return to Step 2.

5. Case Study of Stochastic MGA for the Expansion of Waste Management Facilities

The ability of the stochastic FA-directed MGA procedure to concurrently produce maximally different alternatives will be illustrated using the MSW facilities expansion case study taken from [8]. The region in the facility expansion planning problem consists of three separate municipalities whose MSW disposal needs are collectively met by a landfill and two waste-to-energy (WTE) incinerators. The planning horizon consists of three separate time periods with each of the periods covering an interval of five years. The landfill capacity can be expanded only once over the entire 15-year planning horizon. Each of the WTE facilities can be expanded by any one of four possible options in each of the three time periods. The expansion costs escalate over time to reflect anticipated future conditions and are discounted to present value cost terms for use in the objective function. The MSW waste generation rates and the costs for waste transportation and treatment vary both temporally and spatially. The case requires the determination of the preferred facility expansion alternatives during the different time periods and the effective allocation of the relevant waste flows in order to minimize the total system costs over the planning horizon. Yeomans [8] produced a single best solution to the expansion problem costing \$600.2 million. In order to create three maximally different planning alternatives, (for example, the optimal solution and three alternatives generated by target values of, 2%, 5%, and 8%, respectively), the stochastic FA-directed MGA procedure described in the previous section was run once to produce the objectives for the 4 alternatives shown in Table 1.

Table 1: System Expansion Costs (\$ Millions) for the 4 Alternatives

	Overall “Optimal” Solution	Best 2% Solution	Best 5% Solution	Best 8% Solution
System Expansion Costs	600.19	602.57	611.96	616.44

6. Conclusions

In this paper, a stochastic FA-directed MGA algorithm was introduced that demonstrated how the computationally efficient, population-based FA could be exploited to concurrently generate multiple, maximally different, near-best alternatives via its co-evolutionary solution process. In this MGA capacity, the FA-directed approach produces numerous solutions possessing the requisite problem characteristics, with each generated alternative guaranteeing a very different perspective. The computational example underscored several important findings with respect to the stochastic FA-based MGA method: (i) The FA can be employed as the underlying optimization search routine for SO purposes; (ii) The co-evolutionary capabilities within the FA can be exploited to concurrently generate more good alternatives than planners would be able to create using other MGA approaches because of the evolving nature of its population-based solution searches; (iii) By the design of the MGA algorithm, the alternatives generated are good for planning purposes since all of their structures will be as mutually and maximally different from each other as possible; (iv) The approach is very computationally efficient since it need only be run once to generate its entire set of multiple, maximally different, good solution alternatives (i.e. to generate n solution alternatives, the MGA algorithm needs to run exactly the same number of times that the FA would need to be run for function optimization purposes alone – namely once – irrespective of the value of n); and, (v) The best overall solutions produced by the stochastic MGA procedure will be very similar, if not identical, to the best overall solutions that would be produced by the FA for function optimization alone. Since FA-directed techniques can be adapted to solve a wide variety of problem types, the practicality of the stochastic FA-directed SO MGA approach can clearly be extended into numerous “real world” applications. These extensions will become the focus of future research.

References

- [1]. Imanirad, R., Yang, X.S., & Yeomans, J.S. (2012). A Co-evolutionary, Nature-Inspired Algorithm for the Concurrent Generation of Alternatives. *Journal on Computing*, 2(3): 101-106.



- [2]. Janssen, J.A.E.B., Krol, M.S., Schielen, R.M.J., & Hoekstra, A.Y. (2010). The Effect of Modelling Quantified Expert Knowledge and Uncertainty Information on Model Based Decision Making. *Environmental Science and Policy*, 13(3): 229-238.
- [3]. Loughlin, D.H., Ranjithan, S.R., Brill, E.D., & Baugh, J.W. (2001). Genetic Algorithm Approaches for Addressing Unmodeled Objectives in Optimization Problems. *Engineering Optimization*, 33(5): 549-569.
- [4]. Mowrer, H.T. (2000). Uncertainty in Natural Resource Decision Support Systems: Sources, Interpretation and Importance. *Computers and Electronics in Agriculture*, 27(1-3): 139-154.
- [5]. van Delden, H., Seppelt, R., White, R., & Jakeman, A.J. (2012). A Methodology for the Design and Development of Integrated Models for Policy Support. *Environmental Modelling & Software*, 26(3): 266-279.
- [6]. Walker, W.E., Harremoes, P., Rotmans, J., Van der Sluis, J.P., Van Asselt, M.B.A., Janssen, P., & Kraymer von Krauss, M.P. (2003). Defining Uncertainty – A Conceptual Basis for Uncertainty Management in Model-Based Decision Support. *Integrated Assessment*, 4(1): 5-17.
- [7]. Yang, X.S. (2010). *Nature-Inspired Metaheuristic Algorithms*, Luniver Press, Frome, UK, 2nd Ed.
- [8]. Yeomans, J.S. (2012). Simulation-Driven Optimization in Waste Management Facility Expansion Planning. *Journal of Computational Methods in Sciences and Engineering*, 12(1/2): 111-127.
- [9]. Yeomans, J.S., & Gunalay, Y. (2011). Simulation-Optimization Techniques for Modelling to Generate Alternatives in Waste Management Planning. *Journal of Applied Operational Research*, 3(1): 23-35.

