Metaheuristic Design Pattern: Candidate Solution Repair

Krzysztof Krawiec* Institute of Computing Science Poznan University of Technology 60965 Poznan, Poland krawiec@cs.put.poznan.pl

ABSTRACT

In metaheuristic algorithms applied to certain problems, it may be difficult to design search operators that guarantee producing feasible search points. In such cases, it may be more efficient to allow a search operator to yield an infeasible solution, and then turn it into a feasible one using a repair process. This paper is an attempt to provide a broad perspective on the candidate solution repair and frame it as a metaheuristic design pattern.

Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search—*Heuristic methods*

Keywords

metaheuristic algorithms; solution repair; search operators; feasibility

1. PROBLEM STATEMENT

Many problems solved with metaheuristic algorithms are *constrained*, i.e., problem formulation includes conditions (usually logical predicates) that have to be fulfilled by a candidate solution to be *feasible*. As a result, the entire search space S of candidate solutions is partitioned into *feasible* solutions $F \subset S$ and infeasible solutions $I \subset S$, $F \cap I = \emptyset$.

As the goal is to find the optimal and feasible solution (or, more often in practice, a well-performing suboptimal and feasible solution), it is desirable to traverse S using search operators that are guaranteed to produce feasible solutions (i.e., $s : S \to F$). However, designing operators with this guarantee can be challenging, for instance when feasible solutions are few and far in between (i.e., the ratio |F|/|S| is small), or when the structure of S is such that the feasible solutions are frequently neighbored by the infeasible ones. Moreover, even if such an operator can be designed, then it may be biased in an undesirable way (in the sense of the probability distribution of produced candidate solutions).

This makes it necessary to equip an algorithm with some means of handling the infeasible solutions. Discarding them is rarely a good idea for at least two reasons. Firstly, with the discarded candidate solutions some knowledge gathered in hitherto search is also lost. Secondly, this may prevent the algorithm from visiting the yet unexplored and potentially promising parts of the search space. As an alternative, one may extend the definition of the objective function so that it works 'reasonably' beyond the feasible region; for instance, an infeasible candidate solution to a knapsack problem may be evaluated according to how much it violates the knapsack capacity ([4], p. 149). For many problems however coming up with such an extended and uncontroversial objective function may be difficult.

Another way of handling infeasible solutions is penalization by introducing appropriate terms into objective function f, or by defining additional objectives that implement constraint violation measures. This proceeding assumes however that f can be applied to unfeasible solutions, i.e., that the constraints are in some sense *soft*. Though this is often technically true (e.g., in many continuous optimization benchmarks f is defined in almost entire \mathbb{R}^n space), f applied to an unfeasible solution may return a nonsensical value and potentially deceive the search process. This can be solved by not applying f to infeasible solutions altogether, i.e., by relying on penalties alone for such candidate solutions, but this typically implies that an infeasible solution will be always considered as worse than any feasible solution, which can again prevent exploration of certain parts of the search space.

2. THE SOLUTION

In the face of above controversies, it has been often recommended to repair the unfeasible solutions. A repair operator is any operator O guaranteed to fulfill

$$O(x) \in F$$

for any given infeasible solution $x \in I$.

There are at least two major ways in which such an operator can be employed:

1. In the most obvious scenario, the resulting repaired solution x' = O(x) replaces the infeasible one, and the latter one is discarded. This proceeding is useful in problems where the objective function is capable of evaluating infeasible solutions. For instance, in

^{*}Research conducted during K. Krawiec's stay at the ALFA Group, Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions @acm.org.

GECCO'14, July 12-16, 2014, Vancouver, BC, Canada.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2881-4/14/07 ...\$15.00.

http://dx.doi.org/10.1145/2598394.2609847.

the Traveling Salesperson Problem (TSP), one can still calculate route length even if it visits some cities more than once [10].

2. One can alternatively use the repaired solution x' = O(x) only as a means for obtaining the value of f(x), which otherwise could not be calculated [5]. More specifically, x remains a working solution, but receives the objective value of its repaired version, i.e.:

$$f(x) = f(O(x)) \tag{1}$$

3. CONSEQUENCES

The availability of a repair operator for a given problem has several implications. Firstly, it gives the experimenter more freedom in designing search operators, which do not have to always produce feasible solutions anymore. In some cases, it may pay off to design search operators that *notoriously* yield infeasible solutions, in which case repair becomes a norm rather than an exception, and the repair operator gets to be a more important 'search driver' than the primary search operator(s).

Using repair only to obtain an 'evaluable proxy' (case 2 in Section 2) has particularly interesting consequences. In principle it allows conducting search as if there were no constraints, i.e., in the space S of all candidate solution rather than in the subspace F of feasible solutions only. As a consequence, the search trajectory can traverse the subspace of infeasible solutions I and remain there for several iterations (provided the search operators can work there). This potentially allows the algorithm to find 'shortcuts' to parts of Fthat would be hard (or even impossible) to reach via a path of feasible solutions only. On the other hand, such a policy involves the risk of diverging from F so far that the evaluation proxies obtained by repair (O(x) in Eq. (1)) cease to provide useful information about the search gradient. In general, whether this policy pays off depends on the structure of the search space, characteristics of repair operator, and informativness of fitness function.

Though usually a necessity, solution repair can be seen in a broader context as a way to make a search process more *responsive* to the feedback resulting from the evaluation of candidate solutions. On the other hand, the information about constraint violation can be considered as a manner of making the search process more *informed*.

Note that repair in the above sense can be seen as yet another search operator, capable of operating in I (as opposed to conventional search operators which normally are required to work only in F). Therefore, the characteristics of a repair operator becomes a (potentially important) part of algorithm's *search bias*.

In the context of evolutionary computation, the two modes of operation listed in the previous section can be respectively deemed *Lamarckian repair* and *Baldwinian repair*, as the former one employs the feedback received from the environment (i.e., the information about violating a constraint) to explicitly modify (replace) the solution, while in the latter scenario the only effect is in the objective (fitness) value.

Sometimes it may be difficult to delineate repair from other components of a metaheuristic algorithm. In some cases, repair is an indispensable element of genotype-phenotype mapping. For instance in grammatical evolution [8], solutions (genotypes) are vectors of numbers, each of them determining the choice of a grammar production to be used when deriving the phenotype (program). When such a number is greater than the number of productions available at a given stage of program derivation, it is to be treated by a modulo n operator. That operator can be seen as a repair operator in this context.

A general concept related to repair is *embedding*. The space of feasible solutions can be seen as embedded in a greater space of all candidate solutions reachable by means of the search operator(s). For instance an n-argument Boolean function can be uniquely represented by a look-up table of 2^n values 0 and 1. One can embed this space in the $[0,1]^2$ hypercube, in which only the corners of the hypercube correspond to the feasible solutions. In this setting, repair is the collapse from the latter space to nearest lower space representative (where the notion 'nearest' would be explicitly or implicitly implemented by the repair algorithm). In general, the embedding space can feature a different number of dimensions (typically higher) than the embedded space. However, such embeddings can be considered also for dimensionless (i.e., non-vector) spaces, for instance spaces of programs.

4. EXAMPLES

- Numerical optimization abounds in various methods for solution repair (see, e.g., [5] for a review).
- A simple form of solution repair is clamping (trimming) the value of a modified variable to a predefined range (domain). Though in most cases the ranges are part of problem statement, sometimes the reasons are more subtle. For instance in [9], we clamp the variables representing the parameters of Othello strategies. However, we do so not because the fitness function is *inherently* unable to handle the values outside the interval, but to avoid an uncontrolled growth of variables that may not bring qualitative changes in players' behavior (because candidate solutions are functionally equivalent up to linear scaling). Thus, infeasibility is not the only motivation for repair.
- Crossover can be a significant source of infeasible offspring, in particular in metaheuristic search algorithms with discrete representation encodings. For example, in Grouping Genetic Algorithms [1], crossover of two parents with permutation encodings is highly likely to produce an infeasible offspring. In such cases, solution repair becomes an integral part of the crossover operation.
- Intron removal in genetic programming (GP, e.g., [2]) can be considered as a form of repair (notwithstanding that it has been shown that some amount of introns may be beneficial for GP search).
- Simplification algorithms used in certain variants of genetic programming to reduce solution size (e.g., [6, 7]) can be seen as as a form of repair: they do not affect the semantics of a program, but change its syntax and thus make it possible for it to meet, e.g., an upper limit imposed on program size.
- Methods have been proposed that avoid solution repair by using the information on constraint violation to guide the search process [3].

• Repair is applicable also in presence of weak constraints. For instance in GP, the program size limit is intended to avoid expensive evaluation and excessive memory occupancy rather than to principally exclude certain programs as infeasible. A program that slightly violates this constraint but commits low error can be still considered attractive. In this context, repair interestingly becomes an operator that transforms the 'less desirable' candidate solutions into the more desired ones.

Acknowledgment

The author would like to thank the reviewers for constructive remarks and acknowledges financial support from the Polish-U.S. Fulbright Commission and from National Science Centre grant no. DEC-2011/01-/B/ST6/07318.

5. REFERENCES

- Emanuel Falkenauer. Genetic Algorithms and Grouping Problems. John Wiley & Sons, Inc., New York, NY, USA, 1998.
- Thomas Haynes. Duplication of coding segments in genetic programming. In Proceedings of the Thirteenth National Conference on Artificial Intelligence, volume 1, pages 344–349, Portland, USA, 4-6 August 1996. AAAI Press / MIT Press.
- [3] Nanlin Jin, Edward Tsang, and Jin Li. A constraint-guided method with evolutionary algorithms for economic problems. *Applied Soft Computing*, 9(3):924–935, 2009.
- [4] Sean Luke. Essentials of Metaheuristics. lulu.com, first edition, 2009. Available at

 $http://cs.gmu.edu/{\sim}sean/books/metaheuristics/.$

[5] Z. Michalewicz. A survey of constraint handling techniques in evolutionary computation methods. In John R. McDonnell, Robert G. Reynolds, and David B. Fogel, editors, *Proc. of the 4th Annual Conf. on Evolutionary Programming*, pages 135–155, Cambridge, MA, 1995. MIT Press.

- [6] Alberto Moraglio, Krzysztof Krawiec, and Colin Johnson. Geometric semantic genetic programming. In Christian Igel, Per Kristian Lehre, and Carsten Witt, editors, *The 5th workshop on Theory of Randomized Search Heuristics, ThRaSH'2011*, Copenhagen, Denmark, July 8-9 2011.
- [7] Alberto Moraglio, Krzysztof Krawiec, and Colin G. Johnson. Geometric semantic genetic programming. In Carlos A. Coello Coello, Vincenzo Cutello, Kalyanmoy Deb, Stephanie Forrest, Giuseppe Nicosia, and Mario Pavone, editors, *Parallel Problem Solving from Nature* - *PPSN XII*, volume 7491 of *Lecture Notes in Computer Science*, pages 21–31. Springer, 2012.
- [8] Conor Ryan, J. J. Collins, and Michael O'Neill. Grammatical evolution: Evolving programs for an arbitrary language. In W. Banzhaf, R. Poli, M. Schoenauer, and T. C. Fogarty, editors, *First European Workshop on Genetic Programming 1998*, pages 83–95, Berlin, 1998. Springer.
- [9] M. Szubert, W. Jaśkowski, and K. Krawiec. On scalability, generalization, and hybridization of coevolutionary learning: A case study for othello. *Computational Intelligence and AI in Games, IEEE Transactions on*, 5(3):214–226, 2013.
- [10] Tim Walters. Repair and brood selection in the traveling salesman problem. In Agoston E. Eiben, Thomas Bäck, Marc Schoenauer, and Hans-Paul Schwefel, editors, *Parallel Problem Solving from Nature – PPSN V*, pages 813–822, Berlin, 1998. Springer. Lecture Notes in Computer Science 1498.