Searching for a Permutation in a Haystack

Vincent A. Cicirello
Computer Science and Information Systems
School of Business, Stockton University
101 Vera King Farris Drive, Galloway, NJ 08205
`http://www.cicirello.org/`

Abstract

We provide an overview of the Permutation in a Haystack problem, which serves as a tool for fitness landscape analysis, as well as in search algorithm design. The Permutation in a Haystack enables defining a search landscape for permutation optimization problems that abstracts the permutation features, which impact solution fitness, such as absolute element positions, edges, and pairwise element precedences. In this way, it allows the local search algorithm designer to explore a wide variety of landscape characteristics in a problem independent manner. We introduced the Permutation in a Haystack problem in a paper recently published in *IEEE Transactions on Evolutionary Computation* (Cicirello 2016). In that paper, we also introduced the Search Landscape Calculus, an analytical framework for analyzing the fitness landscape induced by a local search operator via local rates of change of fitness. We then applied the tools to an analysis of several common permutation mutation operators on a wide variety of permutation search landscapes, and validated the findings empirically using simulated annealing.

We define the Permutation in a Haystack problem, Haystack(δ, N): Find the permutation p such that $p = \arg\min_{p' \in S_N} \delta(p', p_N)$, where S_N is the set of all permutations of the set $\{0, 1, \ldots, (N-1)\}$, $p_N = [0, 1, \ldots, (N-1)]$ is a permutation with the elements in increasing order, and δ is a permutation distance metric that serves as the optimization objective function. The optimal solution is $p = p_N$, the "needle" in our figurative "haystack." However, the choice of distance metric affects the terrain of the landscape, and thus search performance. Distance metrics are available to capture the essence of a variety of optimization problems, such as (Ronald 1998; 1997; Kendall 1938); and thus, δ can be chosen based on the permutation features under analysis. For example, to study fitness landscapes where element adjacency impacts fitness, such as the traveling salesperson problem, use a distance metric that interprets a permutation as a set of edges. Thus, Haystack(δ, N) is parameterized to enable experimenting with local search behavior for permutation optimization problems with a variety of features.

References

Cicirello, V. A. 2016. The permutation in a haystack problem and the calculus of search landscapes. *IEEE Transactions on Evolutionary Computation* 20(3):434–446.

Kendall, M. G. 1938. A new measure of rank correlation. *Biometrika* 30(1/2):81–93.

Ronald, S. 1997. Distance functions for order-based encodings. In *Proc. IEEE CEC*, 49–54. IEEE Press.

Ronald, S. 1998. More distance functions for order-based encodings. In *Proc. IEEE CEC*, 558–563. IEEE Press.