Abstracts of Papers Presented at SoCS 2017 in the Previously Published Paper Track

Alex Fukunaga and **Akihiro Kishimoto**
Editors

Abstract

This document gathers the abstracts for the papers that were presented as part of the *Previously Published Paper Track*.

Goal Recognition Design with Stochastic Agent Action Outcomes

Christabel Wayllace, Ping Hou, William Yeoh, and Tran Cao Son New Mexico State University

Goal Recognition Design (GRD) problems involve identifying the best ways to modify the underlying environment that the agents operate in, typically by making a subset of feasible actions infeasible, in such a way that agents are forced to reveal their goals as early as possible. Thus far, existing work assumes that the outcomes of the actions of the agents are deterministic, which might be unrealistic in real-world problems. For example, wheel slippage in robots cause the outcomes of their movements to be stochastic. In this paper, we generalize the GRD problem to Stochastic GRD (S-GRD) problems, which handle stochastic action outcomes. We also generalize the worst-case distinctiveness (wcd) measure, which measures the goodness of a solution, to take stochasticity into account. Finally, we introduce Markov decision process (MDP) based algorithms to compute the wcd and minimize it by making up to k actions infeasible. A full version of this paper appeared in the Proceedings of IJCAI-16 (Wayllace et al. 2016).

Grid Pathfinding on the 2^k Neighborhoods

Nicolás Rivera, King's College London
Carlos Hernández, Universidad Andrés Bello
Jorge A. Baier, Pontificia Universidad Católica de Chile

Grid pathfinding, an old AI problem, is central for the development of navigation systems for autonomous agents. A surprising fact about the vast literature on this problem is that very limited neighborhoods have been studied. Indeed, only the 4- and 8-neighborhoods are usually considered, and rarely the 16-neighborhood. to utilizing extended neighborhoods is that admissible heuristics akin to the Manhattan or octile distances are not known for neighborhoods beyond 16. This paper describes three contributions that enable the construction of effective grid path planners for extended 2^k-neighborhoods. First, we provide a simple recursive definition of the 2^k-neighborhood in terms of the 2^{k-1}-neighborhood. Second, we derive distance functions, for any $k > 1$, which allow us to propose admissible heurisitics which are perfect for obstacle-free grids. Third, we describe a canonical ordering which allows us to implement a version of A whose performance scales well when increasing k. Our empirical evaluation shows that the heuristics we propose are superior to the Euclidean distance (ED) when regular A is used. For grids beyond 64 the overhead of computing the heuristic yields decreased time performance compared to the ED. We found also that a configuration of our A-based implementation, without canonical orders, is competitive with the "any-angle" path planner Theta* both in terms of solution quality and runtime. A full version of this paper is included the Proceedings of the 31st AAAI Conference on Artificial Intelligence (Rivera, Hernández, and Baier 2017). An application of the 2^k neighborhood to subgoal graphs appears in this proceedings (Hormazabal et al. 2017).

Online Bridged Pruning for Real-Time Search with Arbitrary Lookaheads

Carlos Hernández, Universidad Andrés Bello
Adi Botea, IBM Reserach, Ireland
Jorge A. Baier, Pontificia Universidad Católica de Chile
Vadim Bulitko, University of Alberta

Real-time search is a core part of Artificial Intelligence research, with applications in robotics and games. Despite a steady progress over the last several decades, potential poor-quality solutions remain a key obstacle to a wider adoption of real-time search in applications. This is mostly due to their seemingly irrational re-visitation of states. Among the techniques proposed to mitigate such state "scrubbing", on-line state pruning works by removing states irrelevant to a good solution from the search graph. While promising in principle, past pruning techniques have been limited to the myopic local search space with only one expanded state, which limits their impact in practice. We introduce a new pruning approach that works with arbitrarily large local search spaces. Our method is easy to implement and can be used with modern high-performance real-time search algorithms that build a local search space at each decision step. Given a local search space of arbitrary size, our technique aggressively prunes away all states in its interior, possibly adding new edges to maintain the connectivity of the search space frontier. An experimental evaluation shows that our pruning often improves the performance of a base real-time search algorithm by over an order of magnitude. This allows our implemented system to outperform state-of-the-art real-time search algorithms used in the evaluation. The full version of this paper is available in the Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI 2017 (Hernández et al. 2017).

Searching for a Permutation in a Haystack

Vincent A. Cicirello, University of Stockton

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.

Hernández, C.; Botea, A.; Baier, J. A.; and Bulitko, V. 2017. Online bridged pruning for real-time search with arbitrary lookaheads. In *Proceedings of the International Joint Conference on Artificial Intelligence, IJCAI-17*.

Hormazabal, N.; Díaz, A.; Hernández, C.; and Baier, J. A. 2017. Fast and almost optimal any-angle pathfinding using the 2^k neighborhoods. In *Proceedings of the 10th Symposium on Combinatorial Search (SoCS)*. Pittsbugh, Pennsylvania: AAAI Press.

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

Rivera, N.; Hernández, C.; and Baier, J. A. 2017. Grid pathfinding on the 2^k neighborhoods. In Singh, S. P., and Markovitch, S., eds., *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 891–897. San Francisco, California: AAAI Press.

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.

Wayllace, C.; Hou, P.; Yeoh, W.; and Son, T. C. 2016. Goal recognition design with stochastic agent action outcomes. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, 3279–3285.