Learning Crossover Operators in Genetic Algorithms: Application to the Capacitated Vehicle Routing Problem

Amaury Guichard, Quentin Cappart

Polytechnique Montréal, Montreal, Canada
amaury.guichard@polymtl.ca,quentin.cappart@polymtl.ca

Abstract. Industries frequently face optimization challenges such as vehicle routing and scheduling, where even minor improvements can lead to substantial economical and ecological benefits. The Capacitated Vehicle Routing Problem (CVRP) involves finding optimal routes for vehicles with limited capacity to meet client demands. Vidal (2022) introduced a Hybrid Genetic Search that combines local search intensification with genetic algorithm diversification. However, in this method, among all possible children that can be generated by the crossover, a random one is picked. We propose an enhancement to HGS by refining the crossover process so that only children considered good by a defined criterion are generated. Our approach consists in training a Graph Neural Network (GNN) to predict effective crossover subsequences and shows improved performance on smaller instances in terms of number of iterations, but at the expense of an increased execution time. This work was realized under the supervision and advice of Pr. Cappart Quentin at Polytechnique Montréal.

Industries encounter some problems frequently such as drawing routes for their vehicles or timetables for their employees. Such task often involve optimizing an objective function and a small gain in performance can lead to significant economical or ecological gains. The CVRP is a standard constrained optimization problem [4]. Given a set of n clients each with a demand, the goal is to trace the routes of k vehicles with a limited capacity Q so that each route starts and ends at a depot, each client is visited once only and the sum of the demands on a route cannot exceed the capacity of a vehicle. Vidal (2022) [3] proposed a state-of-the-art method to solve this problem known as Hybrid genetic search. This approach combines the intensification capabilities of a local search with the diversification strategies of a genetic algorithms. A population of solution, feasible or infeasible, is maintained. An aggressive local search is applied on each individual of the population to lower the solution cost. During the reproduction phase, two individuals, called parents, are selected and crossed to generate a new individual. A local search is then applied on the new individual before being added to the population. Doing so will increase the population at each iteration. To prevent that, when the population is reaching a certain size, it is reduced by removing the worst individuals and the duplicates. The crossover

OX [1] is used to cross the parents. A subsequence from the first parents forms the basis of the child which is completed with no repetition by the sequence of the second parent.

However the subsequence is extracted randomly in the original method leading to the generation of an individual that can be of no use for the population. We propose a targeted approach to select this subsequence to generate the optimal child according to a specific criterion, thereby improving the method's convergence. The criterion is a characteristic we want to find in our child after the crossover and local search step. This criterion reflects the goal for the population, low solution cost and high diversity. To achieve this, we normalize the cost and multiply it by the diversity of the child with regards to the population. Diversity is computed using the edges of the individual. If the individual has edges that are frequent in the population, the individual is not considered diverse and conversely. The result is called the crossover score and is equal to $NormalizedCost \times Diversity$. The higher the score the better the individual is with regards to our criterion. To verify this criterion, we implemented an exhaustive method. At each iteration, every possible child resulting from the crossover of the two parents is generated and a local search is applied to it. The crossover score is computed for each of them and only the one with the highest score is added to the population. However, there are $\frac{n^2}{2}$ possible children for a given couple where n is the number of clients. This leads to runtime growing quadratically. Because of this, the method can only be tested on small instances with a number of clients ranging from 100 to 250. The same instances used by Vidal (2022) where used [2]. The results on those instances showed that our method was able to solve them to optimality in a few hundreds iteration compared to multiple thousands for the original method. However the exhaustive method takes multiple hours while the original one only takes a few seconds.

The main idea of this work is to use a graph neural network to predict the good crossover points instead of generating all children. Our motivation is to achieve similar performances as the exhaustive test but with a significantly decreased execution time. The GNN is trained to predict for each node, representing the clients, the best crossover score that can be obtained with this node. This way, after a forward pass, the two nodes with the highest crossover score represent the beginning and ending of the subsequence to extract. Our experiments, carried out on on instances of size 100 to 600, showed that the original method was outperforming ours. When limiting the execution to a thousand iterations, we realized that our method found better solution faster but was then stuck in local optima. This was mainly due to the fact that our criterion intensified too much, leading to great results on small instances but poor results on bigger instances.

To conclude, we were able to train a model to generate children based on a defined criterion, influencing the execution of the algorithm. To further improve our method, a better criterion should be determined. Moreover, the duration of an iteration must be reduced for the method to be viable.

References

1. Oliver, I.M., Smith, D.J., Holland, J.R.C.: A study of permutation crossover operators on the traveling salesman problem. In: Proceedings of the Second International Conference on Genetic Algorithms on Genetic Algorithms and Their Application. p. 224–230. L. Erlbaum Associates Inc., USA (1987)
2. Uchoa, E., Pecin, D., Pessoa, A., Poggi, M., Vidal, T., Subramanian, A.: New benchmark instances for the capacitated vehicle routing problem. European Journal of Operational Research **257**(3), 845–858 (Mar 2017). https://doi.org/10.1016/j.ejor.2016.08.012, http://dx.doi.org/10.1016/j.ejor.2016.08.012
3. Vidal, T.: Hybrid Genetic Search for the CVRP: Open-Source Implementation and SWAP* Neighborhood. Computers & Operations Research **140**, 105643 (Apr 2022). https://doi.org/10.1016/j.cor.2021.105643, http://arxiv.org/abs/2012.10384, arXiv:2012.10384 [cs]
4. Vidal, T., Laporte, G., Matl, P.: A concise guide to existing and emerging vehicle routing problem variants. European Journal of Operational Research **286**(2), 401–416 (Oct 2020). https://doi.org/10.1016/j.ejor.2019.10.010, http://dx.doi.org/10.1016/j.ejor.2019.10.010