INFORMATION
Volume 29, Number 1, pp.101-107
ISSN 1343-4500 eISSN 1344-8994

A Mixture-of-Leaders with Teleportation QPSO and its Performance Evaluation for TSP/CVRP

Tatsuya Kiryu and Lei Li

Hosei University, Graduate School of Science and Engineering, Tokyo, Japan.
E-mail:Tatsuya.kiryu.9j@stu.hosei.ac.jp

Abstract

We propose Mixture-of-Leaders with Teleportation QPSO (MoLT-QPSO), a novel variant of Quantum-behaved Particle Swarm Optimization (QPSO) aimed at solving large-scale instances of the Traveling Salesman Problem (TSP) and Capacitated Vehicle Routing Problem (CVRP). Classical QPSO guarantees global convergence and has better search ability than standard PSO, but it still tends to lose diversity and get trapped in local optima . MoLT-QPSO addresses this issue by introducing a multiple leaders strategy and a teleportation mechanism to maintain search diversity. Additionally, we employ a problem-specific solution encoding and integrate local search to handle the discrete nature of TSP/CVRP. In experiments on challenging TSP (up to 11,849 cities) and CVRP (up to 524 customers) benchmarks, MoLT-QPSO consistently outperformed the standard QPSO, finding shorter routes (closer to known optima) and improving average solution quality. Notably, MoLT-QPSO achieved near-optimal results on medium instances (within 0.3–0.5% of optimum) and significantly narrowed the optimality gap on very large instances compared to basic QPSO. These results demonstrate that MoLT-QPSO effectively alleviates premature convergence of QPSO and offers a robust approach for large-scale combinatorial optimization.

Keywords: QPSO, TSP, CVRP, Mixture-of-Leaders

1. Introduction

The Traveling Salesman problem (TSP) and the Capacitated Vehicle Routing Problem (CVRP) are classic NP-hard combinatorial optimization problems with wide real-world applications in logistics and routing. TSP asks for the shortest possible tour that visits each city exactly once and returns to the start, while CVRP extends TSP by assigning customers to multiple vehicle routes with capacity constraints on each vehicle. Both problems become intractable to solve exactly as their size grows (the number of possible solutions grows explosively), so metaheuristic algorithms are often used to find high-quality solutions within reasonable time. Well-known approaches include local search heuristics (e.g., 2-opt, 3-opt), genetic algorithms, ant colony optimization, and Particle Swarm Optimization (PSO) variants.

In particular, PSO and its extensions have drawn attention for large-scale problems due to their ability to handle high-dimensional search spaces.

Quantum-behaved PSO (QPSO) is a PSO variant inspired by quantum mechanics, which replaces the classical velocity update with a stochastic "quantum" update towards attractors. QPSO has fewer parameters and is proven to converge globally under certain conditions. However, a known issue with the standard QPSO (sometimes called *classic QPSO*) is that all particles converge toward a single global best attractor, causing a rapid loss of population diversity and a tendency to stagnate in suboptimal solutions. In other words, premature convergence in QPSO is also inevitable without additional diversification strategies. To overcome this, many PSO variants introduce mechanisms like nonlinear inertia weight schedules, neighborhood topologies (local leaders instead of one global leader), random re-initializations, or hybridization with local search. Building on these ideas, we propose an improved QPSO-based approach for combinatorial problems that specifically targets the diversity issue.

Mixture-of-Leaders with Teleportation QPSO (MoLT-QPSO) is our proposed algorithm designed to enhance QPSO's performance on discrete routing problems. The core idea is to combine a multiple leader guidance strategy with a teleportation restart mechanism, along with problem-specific solution representations and local improvements. By using multiple leaders, the swarm is not guided by just one global best, reducing search bias and helping particles explore different promising regions. The teleportation mechanism periodically "resets" stagnating particles to random positions, which helps escape local optima and inject new genetic material into the population. We also encode TSP/CVRP solutions in a PSO-friendly way (as continuous "giant tours" that are later split into routes) and incorporate domain-specific local search (e.g. , 2-opt edge swaps within tours, or swapping customers between routes for CVRP) to refine solutions. These enhancements allow MoLT-QPSO to maintain search diversity and intensify the search around good solutions when appropriate. MoLT-QPSO was experimentally compared against the standard QPSO on benchmark instances of TSP and CVRP to evaluate its effectiveness. The following sections describe the proposed method in more detail and present the results of these comparisons.

2. Proposed Method: MoLT-QPSO

Multiple Leaders (Mixture-of-Leaders): We maintain multiple leader particles instead of one global best. At certain intervals, the swarm's p_best solutions are clustered in the continuous solution space (using a projection to a lower-dimensional feature space to group similar solutions). Each cluster yields a leader (the best in that cluster). When updating positions, each particle is randomly assigned one of these leaders as its guiding attractor (rather than everyone using the single best solution). We apply a weighted selection so that better clusters have higher chance to guide a particle, but every leader influences some particles. This mixture-of-leaders approach diversifies guidance: particles explore around different promising solutions in parallel, which reduces the risk of the entire swarm following a poor gbest. The cluster assignments are periodically recomputed (e.g., every fixed number of generations) to allow dynamic adaptation of leaders as the search progresses.

Teleportation (Random Restart): MoLT-QPSO includes a teleportation mechanism to reintroduce diversity when stagnation is detected. If the global best has not improved for a specified number of iterations (and the recent improvement rate falls below a small threshold), a fraction of the particles are "teleported" to new random positions in the search space. This is analogous to a random restart or injecting new random individuals. We trigger teleportation in two ways: (1) Event-driven teleportation when stagnation criteria are met (no improvement for some generations), and (2) Periodic teleportation every fixed interval (e.g., every 300 iterations) as a preventive measure. The teleportation rate (percentage of particles reset) can be tuned per problem; in our setting we used a higher rate for TSP than CVRP, reflecting that TSP's search space is larger relative to population size, thus requiring more aggressive diversification. Teleported particles may either be placed uniformly at random or guided by distributing them around the positions of leaders from other clusters (to exploit information from well-performing regions while still providing randomness). This mechanism allows the swarm to escape from any stagnant configuration and continue exploring fresh regions of the solution space.

3. Experimental Results

We evaluated MoLT-QPSO against the standard (basic) QPSO on two TSP instances and two CVRP instances of varying sizes. The TSP instances are taken from TSPLIB: d2103 (2103 cities, best known tour length ≈80,450) and rl11849 (11,849 cities, optimal tour length 923,288). These represent a medium-size and a very large TSP instance, respectively – rl11849 is so large that even state-of-the-art exact methods cannot solve it optimally in reasonable time. The CVRP instances are from the CVRPLIB (Uchoa et al. 2017 benchmarks): X-n101-k25 (100 customers, vehicle capacity 200, best known solution 27,591 using 25 vehicles) and X-n524-k153 (524 customers, capacity 125, best known 154,593 using 153 vehicles). X-n101-k25 is a well-studied benchmark of moderate size, while X-n524-k153 is an extremely large and difficult CVRP example with over 500 customers. For each problem instance, both algorithms were run multiple times (10 independent runs for the medium instances d2103 and X-n101-k25, and 5 runs for the larger rl11849 and X-n524-k153) to assess solution quality and consistency. Each run was limited by the same iteration budget (as described above), and all runs were conducted on the same hardware environment.

Comparison of Results: Table 1 summarizes the performance of basic QPSO versus MoLT-QPSO on the four benchmark instances. We report the best solution distance found in any run and the average solution distance over all runs for each algorithm. For fairness, we also list the gap (%) of those distances relative to the best known solution for the instance (with a lower gap indicating a better solution). In the CVRP cases, we note the number of vehicles used in the best solution as well, since minimizing vehicles is also desirable.

Instance (Best Known)	Algorithm	Best (Gap)	GB update	Vehicles	Avg (Gap)
d2103 (80450)	basic QPSO	81031 (+0.722%)	30	—	81083 (+0.787%)
	MoLT-QPSO	80669 (+0.272%)	49	—	80954 (+0.626%)
rl11849 (923288)	basic QPSO	949115 (+2.797%)	87	—	950686 (+2.967%)
	MoLT-QPSO	944146 (+2.259%)	122	—	945493 (+2.405%)
x-n101-k25 (27591)	basic QPSO	27843 (+0.913%)	16	27	27964 (+1.350%)
	MoLT-QPSO	27629 (+0.138%)	31	26	27679 (+0.321%)
x-n524-k153 (154593)	basic QPSO	157774 (+2.058%)	21	160	159110 (+2.145%)
	MoLT-QPSO	157139 (+1.647%)	32	158	157306 (+1.755%)

Table 1: Performance comparison of standard QPSO vs. MoLT-QPSO on TSP and CVRP instances. Best and average tour lengths over multiple runs are shown, with percentage gap from the best known solution in parentheses. (For CVRP, the number of vehicles in the best solution is also indicated in italics.)

As seen in Table 1, MoLT-QPSO outperforms the basic QPSO on all test instances in terms of solution quality. MoLT-QPSO consistently finds a shorter route (lower distance) than basic QPSO's best, and also yields better average results across runs, indicating improved reliability. For example, on the medium TSP d2103, basic QPSO's best tour length was 81,031 (gap +0.72%), whereas MoLT-QPSO found a tour of length 80,669 (+0.27%), effectively cutting the gap to the optimum by about 0.5 percentage points. MoLT-QPSO also achieved a much lower average gap (+0.63% vs +0.79%), demonstrating it can consistently attain high-quality solutions, not just occasional good runs. Similarly, on the very large TSP rl11849, MoLT-QPSO improved the best solution gap from +2.80% to +2.26%, and the average from +2.97% to +2.40%, indicating a substantial improvement on this extremely challenging instance.

In the CVRP instances, the advantage of MoLT-QPSO is also clear. For X-n101-k25, basic QPSO's best solution was 27,843 distance (0.91% above the known best) using 27 vehicles, whereas MoLT-QPSO found 27,629 (+0.14%) using only 26 vehicles. Not only did MoLT-QPSO get within 0.14% of the optimum distance, but it also managed to service all customers with one fewer vehicle, which is a meaningful improvement in a routing context. The average performance on X-n101-k25 was dramatically better: MoLT-QPSO's average gap was just +0.32%, compared to +1.35% for basic QPSO, nearly a full percentage point improvement. This indicates MoLT-QPSO is far more stable and consistently near-optimal on this instance, highlighting the benefit of its diversity mechanisms and local search integration in avoiding poor runs. On the largest CVRP X-n524-k153 (524 customers), neither algorithm could find the optimal solution within the limited iterations (which is expected for such a large problem), but MoLT-QPSO's best result (157,139, +1.65%) was notably better than basic QPSO's best (157,774, +2.06%). MoLT-QPSO also used fewer vehicles (158 vs 160), again showing better route optimization. The average gap was improved from ~2.15% to ~1.76%. Although a 1.65% gap remains from the best known solution, this gap is significantly smaller than that of basic QPSO, showing that MoLT-QPSO can reach much closer to the best known territory under the same compute budget.

We also analyzed the search dynamics of both algorithms. MoLT-QPSO recorded many more global best updates during the runs than basic QPSO, indicating it continued to find improvements throughout the search instead of stagnating early. For instance, on rl11849, MoLT-QPSO updated the global best 122 times versus only 87 updates in basic QPSO. On d2103 and X-n101-k25, the number of updates in MoLT-QPSO was 1.6–1.9× that of basic QPSO. This quantitatively confirms that the proposed multi-leader and teleportation mechanisms enabled the swarm to avoid getting stuck in local optima and keep searching longer, which translated to better final solutions. Additionally, MoLT-QPSO's results had lower variance between runs (as seen by the smaller gap between best and average), reflecting its improved robustness and repeatability. For example, on X-n101-k25 the worst run of MoLT-QPSO was still around +0.6% gap, whereas basic QPSO in some runs was above +1.3% gap, so MoLT-QPSO reliably produces near-optimal outcomes every time.

Overall, the experimental results validate that each component of MoLT-QPSO contributed to performance gains. The multiple leader strategy allowed parallel exploration of multiple basins of attraction, the teleportation resets prevented long stagnations, and the local search hybridization exploited by MoLT-QPSO yielded final refinements that pure QPSO could not achieve. These led to MoLT-QPSO finding better solutions both in terms of best tour length and consistency across runs.

4. Conclusion

We presented MoLT-QPSO, an improved QPSO-based metaheuristic that integrates a mixture-of-leaders strategy and a teleportation mechanism to maintain diversity during the search, along with tailored solution encoding and local search for discrete problems. The approach was tested on large-scale TSP and CVRP benchmarks and demonstrated significant improvements over the standard QPSO in finding high-quality solutions. MoLT-QPSO was able to reduce the optimality gap by about 0.4–0.8% on average compared to classic QPSO across all instances, and it consistently produced solutions closer to the known optima. In addition, it achieved more stable results with lower variance, indicating that the method is reliable and less prone to being trapped in poor local optima. The incorporation of multiple leaders and teleportation proved effective in prolonging the exploratory phase of the swarm and enabling continual improvements, while the hybrid local search ensured fine-grained optimization of each candidate solution.

The findings suggest that enhancing QPSO's diversity maintenance can yield performance competitive with problem-specific state-of-the-art algorithms, especially on medium-sized instances where MoLT-QPSO came within a fraction of a percent of the best known solutions. Even on very large instances, where obtaining the exact optimum is infeasible, MoLT-QPSO substantially closed the gap that was left by the basic QPSO. These results underscore the potential of MoLT-QPSO as a robust and general framework for tackling complex combinatorial optimization problems like TSP and VRP.

For future work, further improvements could be explored such as adaptive parameter control (e.g., dynamically adjusting the number of leader clusters or teleportation frequency based on the search progress). Additionally, applying MoLT-QPSO to other constrained routing problems (like VRP variants with time windows or scheduling problems) is an interesting direction, as the general diversity mechanisms and local search integration used here should be beneficial in those contexts as well. In summary, MoLT-QPSO effectively mitigates the early convergence issue of QPSO and opens up new possibilities for applying particle swarm techniques to large-scale discrete optimization challenges.

References

[1] Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. In Proceedings of ICNN'95 - International Conference on Neural Networks (Vol. 4, pp. 1942-1948). IEEE.

[2] Sun, J., Feng, B., & Xu, W. (2004). Particle swarm optimization with particles having quantum behavior. In Proceedings of 2004 Congress on Evolutionary Computation (CEC2004) (Vol. 1, pp. 325-331). IEEE.

[3] Lin, A., et al. (2018). Multi-Leader Comprehensive Learning Particle Swarm Optimization with Adaptive Mutation for Economic Load Dispatch Problems. Algorithms, 11(12), 203.

[4] Budhraja, I., et al. (2012). Exploration-enhanced particle swarm optimization using guided teleportation. Swarm and Evolutionary Computation, 7, 1-11.

[5] Uchoa, E., et al. (2017). New benchmark instances for the capacitated vehicle routing problem. European Journal of Operational Research, 257(3), 845-858.

[6] Li, J., Abdul Majeed, A. P. P., & Lefevre, P. (2020). Halfway Escape Optimization: A Quantum-Inspired Solution for General Optimization Problems. ArXiv Preprint.

[7] Sun, L., & Li, K. (2020). Adaptive operator selection based on dynamic Thompson sampling for MOEA/D. In Proceedings of the 16th International Conference on Parallel Problem Solving from Nature (PPSN XVI), Lecture Notes in Computer Science, 12270, 511–523. https://doi.org/10.1007/978-3-030-58112-1_34

[8] Scoczynski, F., Sobieraj, P., & Nalepa, J. (2021). HHTS: A hyper-heuristic guided by Thompson sampling for continuous optimization. ArXiv preprint arXiv:2102.07323. https://arxiv.org/abs/2102.07323