Delivery Route Optimization using Genetic Algorithms

Houtari, Abdulaziz
houtaria@msu.edu

Shakir, Samir
shakirsk@msu.edu

December 11, 2024

1 Introduction

The Vehicle Routing Problem (VRP) is a fundamental optimization challenge in logistics, where the goal is to determine the optimal set of routes for a fleet of vehicles to deliver goods to customers while minimizing costs and meeting constraints such as vehicle capacity and time windows. VRP is classified as an NP-hard problem, meaning the solution space grows exponentially with the number of locations, making exact methods impractical for large instances. This project leverages Genetic Algorithms (GAs) to address the VRP due to their flexibility in handling large search spaces and complex constraints.

In this work, we experimented with various configurations of GAs, including multi-objective and single-objective setups, with and without local optimization (2-opt) to evaluate their performance on solving VRP instances. Our goal was to minimize the total distance traveled while adhering to capacity constraints and optimizing vehicle usage.

2 Related Works

Several studies have applied Genetic Algorithms (GAs) to the Vehicle Routing Problem (VRP), highlighting their adaptability and efficiency in solving complex logistics problems. Xin et al. [4] demonstrated GA's effectiveness in logistics distribution route optimization. Ibrahim et al. [2] proposed optimized GA operators for VRPs with time windows, emphasizing improvements in crossover and mutation. Surekha and Sumathi [3] addressed multi-depot VRP challenges using GA, while Agrawal et al. [1] applied GA to perishable product routing under quality and time constraints. Zhang et al. [5] explored cold chain logistics using an improved GA to meet time-critical delivery requirements. These works collectively underscore GA's suitability for solving VRP variants, inspiring our methodology.

3 Methodology

3.1 Problem Setup

The VRP is set up in a simulated environment, where:

1. Customer locations are clustered, ensuring spatial distribution that mimics real-world delivery locations.
2. Each customer has a demand, and the fleet of vehicles has defined capacity limits.
3. The depot is centrally located, and routes are optimized to minimize travel distance and satisfy capacity and vehicle constraints.

The algorithm divides the VRP into several clusters, each representing a region with unique customer locations. Figure 1 demonstrates the initial setup, showing clusters with customer locations, cluster centers, and the depot.

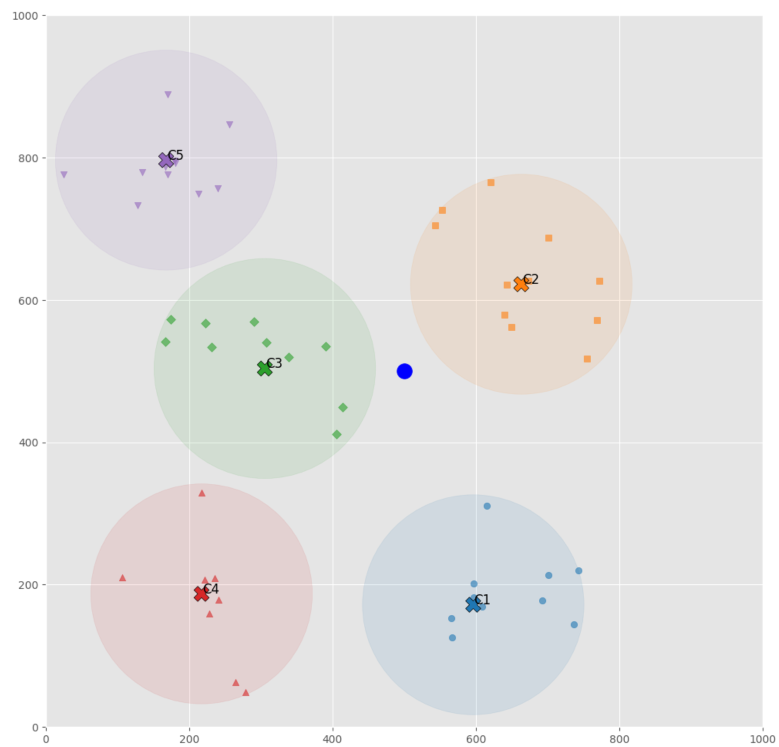

Figure 1: Problem setup example with Depot in the center and clustered customers.

3.2 GA Implementation

3.2.1 Chromosome Representation

Each chromosome encodes customer visits and route delimiters to distinguish different vehicle routes. This representation allows flexibility in assigning customers to vehicles.

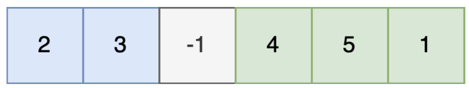

Figure 2: This represent one vehicle serving customers 2 and 3 and then coming back to the depot, and another vehicle serving 4, 5, and 1 and coming back to the depot.

3.2.2 Fitness Evaluation

For the multi-objective GA, the fitness function aims to minimize travel distance, capacity per vehicle, and number of vehicles used. Whereas in the single-objective GA, the fitness functions mainly aims to minimize travel distance and penalizing if constraints are violated.

3.2.3 Genetic Operators

- **Selection**: NSGA-II is used for multi-objective optimization, while tournament selection (size = 3) is used for single-objective optimization.
- **Crossover**: Routes are exchanged between parents while ensuring all customers are served.
- **Mutation**: Customers are swapped within or across routes to explore new solutions.

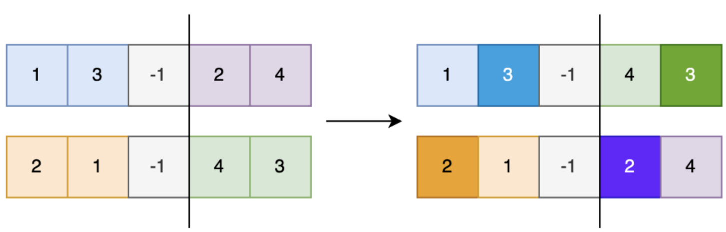

Figure 3: Crossover operator (before repairing).

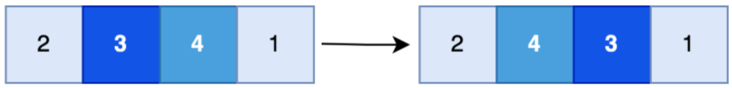

(a) Intra-route mutation.

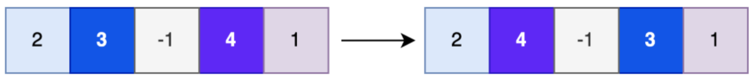

(b) Inter-router mutation.

Figure 4: Mutation operator.

3.2.4 Repair Mechanism

A repair function redistributes excess demand among vehicles to adhere to capacity constraints. If a route becomes infeasible, customers are reassigned to new or less utilized routes.

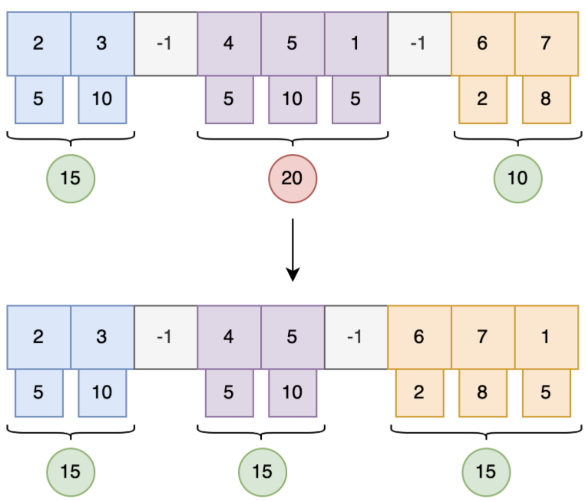

Figure 5: Repair mechanism example with vehicle capacity=15 and maximum number of vehicles=3

3.2.5 Local Optimization

Every 10 generations, the top 5 individuals undergo the 2-opt algorithm to refine their routes. This step reduces overlapping routes and improves overall efficiency.

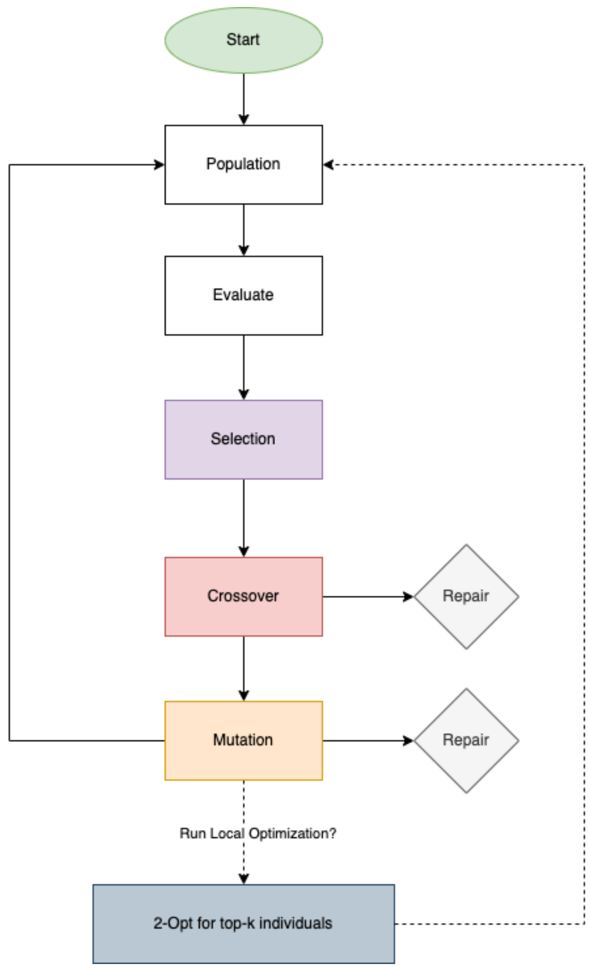

Figure 6: Lifecycle of a GA with local optimization.

4 Experiments

We implemented and evaluated the following configurations:

1. Multi-objective GA.
2. Constrained single-objective GA.
3. Pure 2-opt optimization.
4. Single-objective GA combined with 2-opt optimization.

Key GA parameters:

- Population size: 300.

- Generations: 500.
- Crossover rate: 70%.
- Mutation rate: 10%.

VRP parameters and setup:

- Grid size: 1000x1000.
- Number of locations: 50.
- Maximum vehicles: 10.
- Vehicle capacity: 50.
- Location demands: Gaussian distribution.
- Distance calculation: Euclidean.

5 Results and Analysis

Approach	Min Distance	Vehicles Used
Multi-Objective GA	17,729	6
Constrained Single-Objective GA	9,182	5
2-Opt	13,463	5
Constrained Single-Objective GA + 2-Opt	**8,401**	**5**

Table 1: Results of Different Approaches for Vehicle Routing Problem

5.1 Multi-Objective GA

5.1.1 Trade-offs and Pareto Front

The multi-objective GA demonstrates the trade-offs between minimizing total distance, reducing capacity violations, and minimizing the number of vehicles used. The Pareto front visualization reveals that while the algorithm successfully generates diverse solutions, many of these solutions emphasize reducing distance but at the cost of increasing vehicle usage which can be seen in Figure 7

5.1.2 Distance vs. Vehicles Used

The evolution of fitness over generations indicates a gradual decrease in total distance. However, the multi-objective focus creates a conflict between optimizing distance and minimizing the number of vehicles, leading to an uneven utilization of the fleet as we can see in Figure 8.

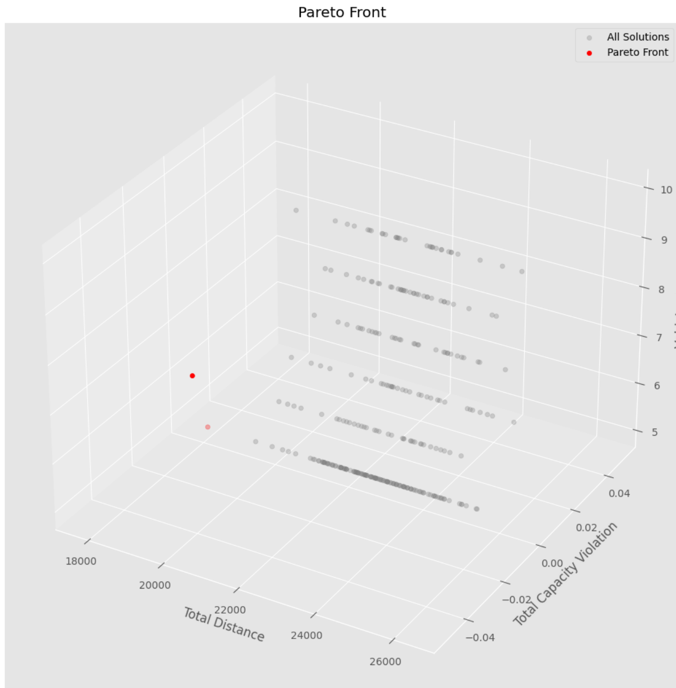

Figure 7: Pareto front of total distance, vehicles used, and capacity violations.

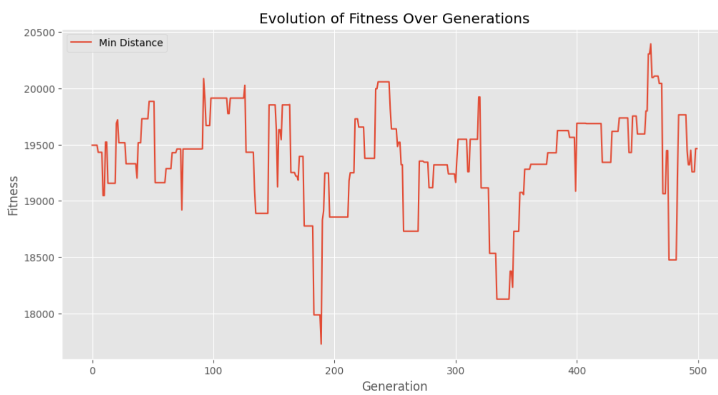

Figure 8: Evolution of fitness over generations in the multi-objective GA.

5.1.3 Route Clustering

The resulting routes show poor geographical clustering. Vehicles are assigned to customers scattered across clusters, resulting in overlapping paths. This inefficiency highlights the algorithm's struggle to balance multiple objectives effectively. Furthermore, the repair function, while maintaining capacity constraints, may limit the algorithm's ability to explore better configurations by overly distributing customers. Such result can be seen in Figure 9.

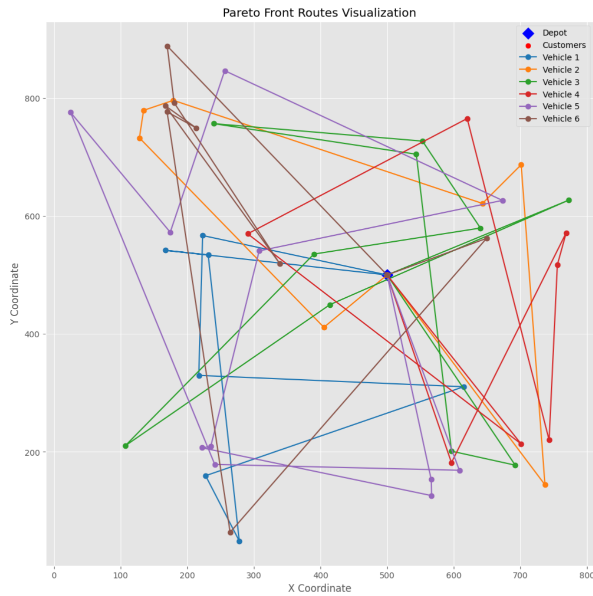

Figure 9: Multi-Objective GA best solution for minimal distance traveled.

5.2 Constrained Single-Objective GA

5.2.1 Focused Optimization

The single-objective GA significantly reduces travel distance by focusing solely on distance minimization and treating capacity and fleet size as strict constraints. This approach eliminates the conflict between objectives, enabling the algorithm to converge faster to optimal solutions.

5.2.2 Geographical Clustering and Routes

Routes are more localized and tightly clustered, which reduces overlap and unnecessary travel. The use of a repair function ensures that capacity constraints are always met without sacrificing the primary objective. This clustering behavior not only improves route efficiency but also aligns with real-world delivery expectations where vehicles typically serve well-defined areas. This can be seen in Figure 11.

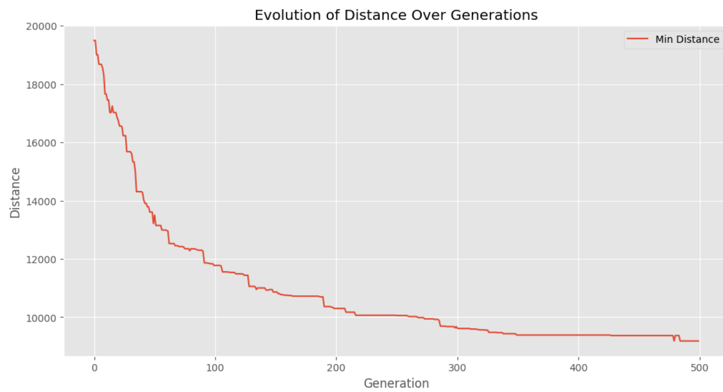

Figure 10: Single-Objective GA fitness evolution over generations.

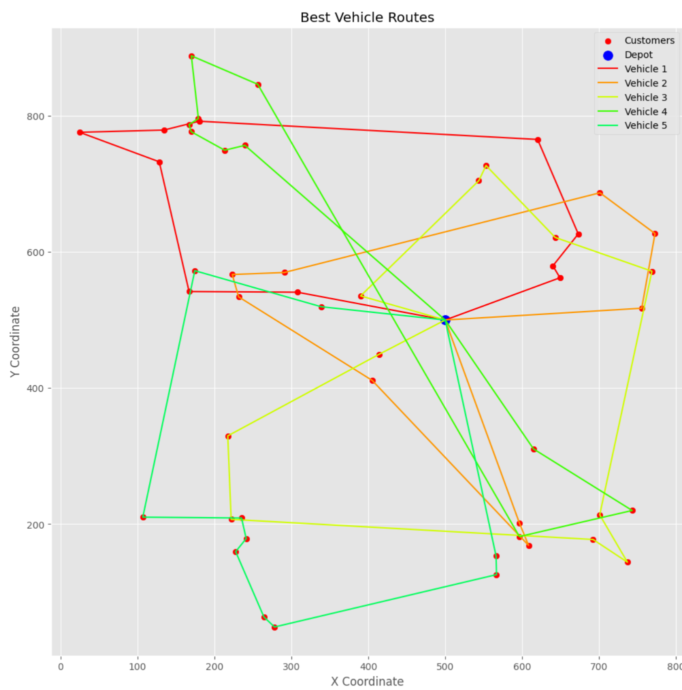

Figure 11: Single-Objective GA best solution for minimal distance traveled.

5.2.3 Performance Gains

Compared to the multi-objective GA, this approach achieves better results in terms of distance and fleet utilization, reflecting the advantage of focusing on a single well-defined goal.

5.3 2-Opt Only

5.3.1 Distance Reduction

The 2-opt local optimization algorithm successfully reduces travel distance by iteratively improving existing routes. However, it lacks the adaptability of GA in redistributing customers or generating entirely new routes, limiting its potential for large-scale optimization.

5.3.2 Route Overlap

The absence of a global search mechanism results in routes with overlapping paths and poor vehicle utilization. The improvement in travel distance is modest compared to the single-objective GA, highlighting 2-opt's dependency on good initial solutions.

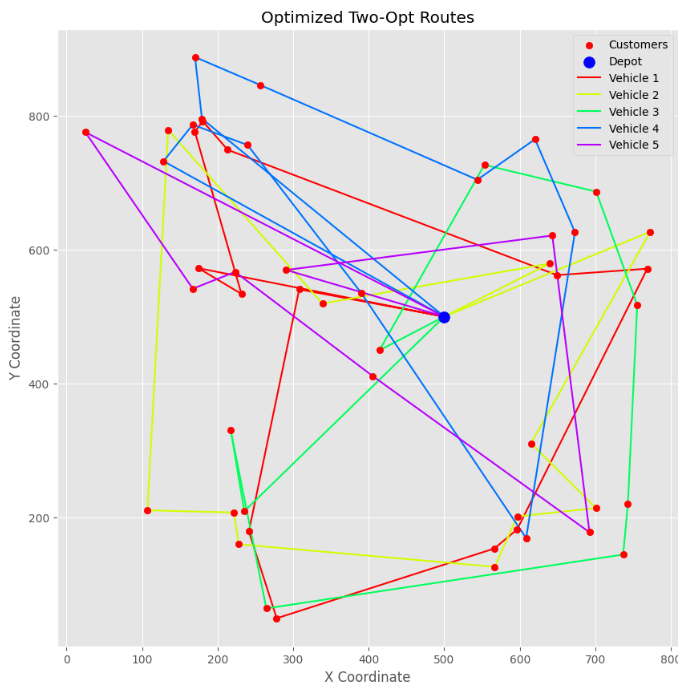

Figure 12: 2-Opt only solution.

5.3.3 Applicability

While effective for refining routes, 2-opt alone cannot address capacity violations or optimize customer assignments, making it less versatile than GA-based approaches.

5.4 Constrained Single-Objective GA + 2-Opt

5.4.1 Best-Performing Approach

This hybrid approach combines the strengths of GA and 2-opt, leveraging GA for global exploration and 2-opt for local optimization. The result is the lowest travel distance across all approaches (8,401 units), with efficient utilization of the available vehicles.

5.4.2 Improved Clustering

The combination of global and local optimization produces well-clustered routes, reducing overlap and unnecessary travel. The use of 2-opt at regular intervals helps refine the solutions generated by GA, ensuring that local improvements are consistently incorporated.

5.4.3 Flexibility

By treating capacity and fleet size as strict constraints, this approach avoids the trade-offs seen in multi-objective optimization. The addition of 2-opt ensures that the algorithm doesn't settle for suboptimal solutions, addressing both global and local aspects of the problem.

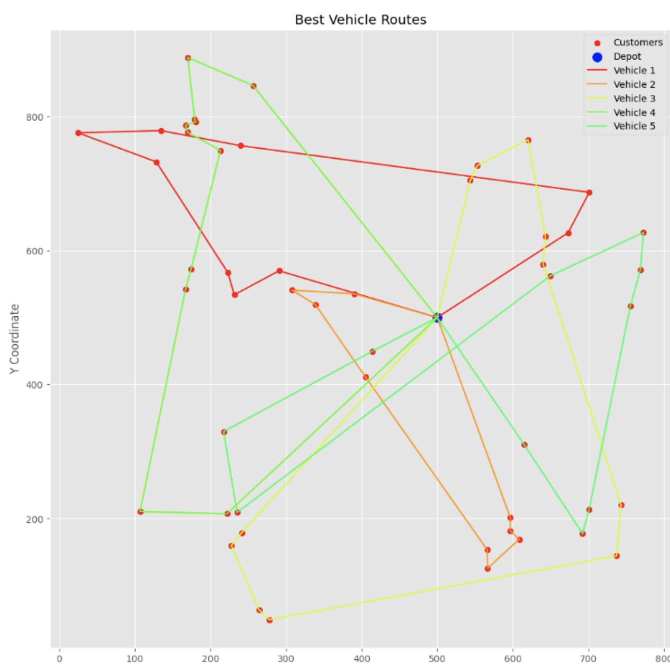

Figure 13: Single-Objective GA + 2-opt local search best solution for minimal distance traveled.

6 Conclusion

This project demonstrates the effectiveness of Genetic Algorithms in solving the Vehicle Routing Problem. Key findings include:

- Incorporating local optimization (2-opt) significantly improves solution quality.
- Treating capacity and vehicle constraints as hard constraints, rather than objectives, enhances route clustering and minimizes distance.
- Multi-objective optimization provides insights into trade-offs but may lead to suboptimal solutions for practical applications.

Future work will focus on integrating time-window constraints and expanding the VRP setup to include real-world road networks.

References

[1] A. K. Agrawal et al. "A genetic algorithm model for optimizing vehicle routing problems with perishable products under time-window and quality requirements". In: *Decision Analytics Journal* 5 (2022). DOI: 10.1016/j.dajour.2022.100139.

[2] M. F. Ibrahim et al. "Optimised genetic algorithm crossover and mutation stage for vehicle routing problem pick-up and delivery with time windows". In: *IOP Conference Series: Materials Science and Engineering*. Vol. 1071. 1. 2021, p. 012025. DOI: 10.1088/1757-899X/1071/1/012025.

[3] P. Surekha and S. Sumathi. "Solution to multi-depot vehicle routing problem using genetic algorithms". In: *World Applied Programming* 1.3 (2011), pp. 118–131.

[4] L. Xin, P. Xu, and G. Manyi. "Logistics distribution route optimization based on genetic algorithm". In: *Computational Intelligence and Neuroscience* 2022 (2022), pp. 1–9. DOI: 10.1155/2022/8468438.

[5] L. Zhang et al. "Application on cold chain logistics routing optimization based on improved genetic algorithm". In: *Automatic Control and Computer Sciences* 53.2 (2019), pp. 169–180. DOI: 10.3103/S0146411619020032.