

Delivering Vaccines, The Last Mile: Operations Research Optimizes Cold Chain Distribution in Mozambique and Norway

Zeel Shah¹, Vedant Laijawala², Vivaan Rajani³, Yash Rathi⁴, Yashvardhan Mohanty⁵, Yug Jain⁶
^{1,2,3,4,5,6} Anil Surendra Modi School of Commerce (ASMSOC), SVKM's NMIMS, Mumbai, Maharashtra, India

I. INTRODUCTION TO OPERATIONS RESEARCH AND THE PROBLEM TYPE

1.1 Operations Research

Operations Research (OR) is the discipline of using mathematical models, optimisation, data and analytical techniques to help decision-makers choose the best course of action when resources are limited and constraints must be respected (Taha, 2017). An OR study typically follows a fixed sequence: the real situation is described, the decision to be taken is identified, the problem is translated into a mathematical model with an objective function and constraints, the model is solved with an exact or approximate method, and the solution is tested before it is used in practice.

The value of OR lies in this translation. A messy, real-world question such as “how should vaccines reach every clinic in a district before they spoil?” becomes a structured problem whose alternatives can be compared objectively.

Vaccine distribution is a natural application of OR. Health authorities must decide where vaccines should be stored, which vehicle should carry them, which facilities each vehicle should visit and in what order, and how much each facility should receive.

Each of these decisions is constrained by vehicle capacity, storage capacity, working hours, road conditions and, above all, the cold chain, because vaccines lose potency if they are exposed to temperatures outside a narrow range (Zabinsky et al., 2021). These are exactly the conditions under which OR techniques outperform intuition or hand-drawn plans.

1.2 The Transportation Problem

This report studies vaccine logistics through the lens of the transportation problem, one of the classical models of OR, first formulated by Hitchcock (1941). In its basic form, a single homogeneous product is shipped from m sources (for example, vaccine stores) to n destinations (for example, district centres). Source i can supply at most a_i units, destination j requires b_j units, and shipping one unit from i to j costs c_{ij} . The decision variable x_{ij} is the quantity shipped from source i to destination j . The model is a linear programme (Taha, 2017):

$$\text{Minimise } Z = \sum_{i=1}^m \sum_{j=1}^n c_{ij} x_{ij}$$

$$\text{subject to } \sum_{j=1}^n x_{ij} \leq a_i \quad \forall i \quad (\text{supply constraints})$$

$$\sum_{i=1}^m x_{ij} \geq b_j \quad \forall j \quad (\text{demand constraints})$$

$$x_{ij} \geq 0 \quad \forall i, j \quad (\text{non-negativity})$$

When total supply equals total demand, the problem is balanced; otherwise, a dummy source or destination with zero cost is added. The standard solution procedure first finds an initial basic feasible solution using the North-West Corner Rule, the Least Cost Method or Vogel's Approximation Method (VAM), and then tests and improves it for optimality using the Modified Distribution (MODI) or stepping-stone method (Taha, 2017).

A short illustration shows how the model works in a vaccine setting. Suppose two regional stores, A and B, hold 300 and 400 vaccine carriers, and three district

centres, D1, D2 and D3, need 200, 250 and 250 carriers. The cost of moving one carrier (in rupees) is shown below.

Store	D1	D2	D3	Supply
A	4	6	8	300
B	5	3	7	400
Demand	200	250	250	700

Because supply and demand both total 700, the problem is balanced. VAM first allocates 250 carriers from B to D2 (the column with the largest penalty), then 200 from A to D1, and the remaining D3 demand is met with 100 carriers from A and 150 from B. The total cost is $4(200) + 3(250) + 8(100) + 7(150) = ₹3,400$. The MODI test gives non-negative opportunity costs (+2) for both empty cells (A–D2 and B–D1), so this allocation is optimal, and the result was confirmed by solving the same linear programme in software.

1.3 From the Classical Model to Vaccine Logistics

The classical transportation problem assumes direct shipments, linear costs, unlimited vehicles and certain

demand. Real vaccine distribution breaks all four assumptions. Vaccines usually pass through intermediate stores, which makes the problem a transshipment problem. Deciding which stores to open adds fixed costs and binary variables, which turns it into a network design or hub location problem (Campbell & O’Kelly, 2012; Farahani et al., 2013). A single vehicle usually serves several facilities on one trip, so the order of stops matters, which makes it a Vehicle Routing Problem (VRP).

The VRP was first introduced by Dantzig and Ramser (1959) as the “truck dispatching problem” and is a generalisation of the travelling salesman problem (Laporte, 1992; Toth & Vigo, 2014; Rodrigues & Lima, 2021). Finally, demand, supply and road conditions are uncertain. The five papers reviewed in this report each take the transportation problem as their foundation and extend it in one or more of these directions. Taken together, they move from strategic network design (where to store vaccines and how much should flow along each link) to operational routing (which vehicle drives which route on a given day).

No.	Paper	Setting	Transportation-family model
P1	Yang, Bidkhor & Rajgopal (2021)	Four sub-Saharan African countries	Capacitated transshipment / hub network design (MIP)
P2	Mohammadi, Dehghan, Pirayesh & Dolgui (2022)	Metropolitan France	Multi-level transshipment network design under uncertainty (bi-objective MINLP)
P3	Andoh & Yu (2023)	Inland County, Norway	Capacitated VRP with heterogeneous vehicles plus simulation
P4	Zabinsky, Zameer, Petroianu, Muteia & Coelho (2021)	Mozambique	Bi-objective VRP (transit time and risk) in an Excel tool
P5	Rodrigues & Lima (2021)	Pernambuco, Brazil	Multi-depot VRP solved by a GRASP-VND metaheuristic

Two broader reviews confirm this picture. Duijzer, Van Jaarsveld and Dekker (2018) survey OR models across the whole vaccine supply chain, from product design to distribution, and Dey, Kurbanzade, Gel, Mihaljevic and Mehrotra (2023) focus on models developed during COVID-19.

It groups the OR models developed during COVID-19 into site selection, inventory management, allocation, distribution logistics and route optimisation, which are the same families represented by the five papers studied here.

II. UNDERSTANDING THE REAL-LIFE CONTEXT

2.1 Industry Description: The Vaccine Cold Chain

The vaccine supply chain is a specialised part of pharmaceutical and healthcare logistics. In most countries it is organised as a hierarchy. Vaccines are bought in bulk by a national authority, flown into a national store, moved to regional and district stores, and finally delivered to clinics or vaccination centres where people are immunised (Yang et al., 2021; Mohammadi et al., 2022). In low- and middle-income countries (LMICs), routine childhood vaccines are

managed under the World Health Organization's Expanded Programme on Immunization (WHO-EPI), with financial support from Gavi, the Vaccine Alliance (Yang et al., 2021). In high-income countries, transport is frequently outsourced to third-party logistics (3PL) companies that operate refrigerated vehicles (Andoh & Yu, 2023).

The defining feature of this industry is the cold chain. Most vaccines must be stored and transported between 2°C and 8°C (Yang et al., 2021; Zabinsky et al., 2021), and some COVID-19 vaccines, such as Pfizer-BioNTech's, required ultra-cold storage at around -70°C (Andoh & Yu, 2023). A vaccine that arrives late, or is exposed to heat or freezing on a long journey, can lose its potency and become useless even though it physically reaches the clinic (Zabinsky et al., 2021). The stakes are high. Vaccines avert an estimated 2 to 3 million deaths each year (Yang et al., 2021), yet global coverage has stagnated at about 85% since 2010, and in 2017 an estimated 20 million children still lacked access to routine immunisation, around 60% of them in resource-constrained countries, and inefficiency in the supply chain is one of the reasons for this gap (Zabinsky et al., 2021). The COVID-19 pandemic raised the stakes further: supply was scarce and uncertain in the early months, several vaccines with different storage needs were in use at once, and entire national populations had to be reached quickly (Kartoglu et al., 2020; Andoh & Yu, 2023; Dey et al., 2023). The last mile, which is the final movement from a district or regional store to the vaccination point, is the most difficult and expensive stretch (Lee & Chen, 2021).

A district may contain dozens of facilities spread over a wide area but only a handful of vehicles. Roads may be unpaved or seasonally flooded, vehicles may break down, and each facility needs a different quantity. Poor planning at this stage leads to spoiled doses, stockouts, higher transport costs and lost vaccination opportunities (Zabinsky et al., 2021; Andoh & Yu, 2023).

2.2 The Problem Addressed in Each Paper

P1: Network design in sub-Saharan Africa (Yang et al., 2021). In many LMICs, EPI vaccines move through a four-tier legacy network (national, regional, district and clinic) whose structure was set by political boundaries and history rather than by efficiency. The

authors ask whether a leaner network, using only a subset of existing facilities as hubs and following WHO operating rules, could deliver vaccines to every child at lower cost. The problem is to choose which hubs to open, which clinics each hub serves, which storage device and vehicle each location should use, and whether each hub should be replenished monthly or quarterly. Earlier work by Chen et al. (2014) had modelled how to operate an existing WHO-EPI network; Yang et al. instead redesigns the network itself.

P2: A resilient vaccine network for France (Mohammadi et al., 2022). France had to distribute four COVID-19 vaccines with limited supply through national, regional and departmental warehouses to vaccination centres serving a population of about 65 million. The network was exposed to disruptions such as storage capacity loss, destroyed inventory, supplier delays and the emergence of new virus variants. The question is how to design and operate this network so that it minimises deaths while keeping costs under control, and remains robust under several possible futures.

P3: Sustainable last-mile delivery in Norway (Andoh & Yu, 2023). Norway's early COVID-19 supply was small, uncertain and had to be delivered fairly to 46 municipalities in Inland County. Planners had to balance timely delivery, transportation cost and carbon dioxide (CO₂) emissions, three goals that pull against each other.

The question is how network structure (where the depot is), fleet size and fleet composition (electric versus fuel vehicles) affect these three dimensions of sustainability.

P4: Route planning in Mozambique (Zabinsky et al., 2021). In Mozambique, delivery routes were drawn by hand, and logisticians often learned which vehicle was available only when it arrived at the warehouse. Seasonal flooding and Cyclone Idai in 2019 destroyed roads and damaged health facilities, making some areas completely inaccessible. Government users needed a free, fast and simple tool to create and re-plan routes that minimise both travel time and the risk of cold chain failure, an idea the authors draw from risk-based routing in humanitarian logistics (Hamed et al., 2012).

P5: Vaccine routing in Pernambuco, Brazil (Rodrigues & Lima, 2021). In the state of Pernambuco, vaccines arrive by air in Recife and pass to 12 regional health offices (GERES). Officials from each of the 172 surrounding municipalities then travel to their GERES and back to collect doses, so many routes are driven twice. The authors ask whether a centrally planned routing system, with one or more depots, could deliver the same doses while travelling far less distance.

2.3 Why These Particular OR Techniques Were Used

The choice of technique in each paper follows directly from the nature of the decision and the people who must use the result. Where the decision is strategic and made infrequently, as in network design (P1 and P2), the authors use mixed integer programming because the model must capture yes-or-no choices such as opening a hub, alongside continuous flows of vaccines. These are extensions of the transportation and transshipment problems with fixed costs. Where the decision is operational and repeated weekly or daily, as in routing (P3, P4 and P5), the authors use vehicle routing models because the order in which a vehicle visits facilities determines travel time, cost and cold chain exposure.

Three further considerations shape the choice of method. First, size: routing and network design problems are NP-hard, meaning that computation time grows explosively with the number of facilities (Laporte, 2009; Andoh & Yu, 2023; Rodrigues & Lima, 2021), so P1 develops a decomposition algorithm, P4 a fast branch-and-bound heuristic and P5 a metaheuristic. Second, uncertainty: P2 uses scenario-based robust optimisation and P3 adds discrete event and Monte Carlo simulation, because a plan that is optimal on paper may fail when travel speeds, supply or demand change (Ivanov, 2020; Tordecilla et al., 2021). Third, multiple goals: P2 weighs deaths against cost, P3 weighs service level, cost and emissions, and P4 lets the user weigh transit time against risk. A single cost-minimising transportation model could not represent these trade-offs, so each paper adapts the basic model to its setting.

III. MODELLING AND PROBLEM FORMULATION

This section sets out the mathematical model used in each paper, identifying the type of problem, its

decision variables and parameters, the objective function, the constraints and the main assumptions. Where a paper describes its model only in words, this is stated, and the formulation shown is a translation of that description into standard notation.

3.1 P1: Vaccine Network Design as a Mixed Integer Programme (Yang et al., 2021)

Type of problem. A capacitated, single-allocation hub location and transshipment problem (Campbell & O'Kelly, 2012) formulated as a mixed integer linear programme (MIP-1). Vaccines flow from a national centre (source) to clinics (sinks) through hubs (transshipment nodes).

Sets. $N = \{0\}$ is the national distribution centre, H is the set of potential hubs (existing regional and district stores), C is the set of clinics, T is the set of vehicle types, R is the set of WHO-approved storage devices, and $F = \{\text{quarterly, monthly}\}$ is the set of replenishment frequencies.

Parameters. c_{ijt}^P is the transport cost per km for vehicle type t ; c_{jr}^S is the annual cost of operating facility j with storage device r ; p_t^P and p_r^S are vehicle and storage capacities; g_f is the number of replenishments per year (4 for quarterly, 12 for monthly); d_{ij} is the road distance; and b_j is the annual demand volume of clinic j .

Decision variables. $U_{ijtf} \in \{0,1\}$ equals 1 if vaccines flow from i to j by vehicle t at frequency f ; $W_{jrf} \in \{0,1\}$ equals 1 if hub j is open with device r and frequency f ; and $X_{ij} \geq 0$ is the annual volume shipped from i to j .

Objective function. Minimise the annual cost of operating hubs plus the annual cost of round trips:

$$\text{Min} \sum_{j \in H} \sum_{r \in R} \sum_{f \in F} c_{jr}^S W_{jrf} + \sum_{(i,j) \in A} \sum_{t \in T} \sum_{f \in F} 2 c_{ijt}^P g_f d_{ij} U_{ijtf}$$

Constraints.

$$\sum_{r \in R} \sum_{f \in F} W_{jrf} \leq 1 \quad \forall j \in H \quad (\text{one device and frequency per hub})$$

$$\sum_{i \in NUH} \sum_{t \in T} U_{ijt1} = 1 \quad \forall j \in C \quad (\text{each clinic has one supplier, monthly})$$

$$\sum_{i \in NUH} \sum_{t \in T} \sum_{f \in F} U_{ijtf} \leq 1 \quad \forall j \in H \quad (\text{at most one inflow per hub})$$

$$\sum_{r \in R} W_{jrf} - \sum_{i \in NUH} \sum_{t \in T} U_{ijtf} = 0 \quad \forall j \in H, f \in F \quad (\text{no flow to a closed hub})$$

$$\sum_i X_{ij} - \sum_k X_{jk} = 0 \quad \forall j \in H \quad (\text{flow balance at hubs})$$

$$\sum_i X_{ij} = b_j \quad \forall j \in C \quad (\text{demand satisfaction})$$

$$\sum_{r \in R} \sum_{f \in F} p_r^S W_{jrf} - \frac{1}{g_f} \sum_i X_{ij} \geq 0 \quad \forall j \in H \quad (\text{storage capacity})$$

$$\sum_{t \in T} \sum_{f \in F} p_t^P U_{ijtf} - \frac{1}{g_f} X_{ij} \geq 0 \quad \forall (i, j) \quad (\text{vehicle capacity})$$

$$X_{ij} \geq 0, \quad W_{jrf}, U_{ijtf} \in \{0,1\}$$

The demand and flow balance constraints are the transportation and transshipment constraints of Section 1.2. What the model adds are binary choices of hubs, equipment and frequency, and the rule that every location has exactly one supplier.

Assumptions. All demand at each clinic must be met and is estimated from the population served; clinic and national centre locations are fixed, while hubs can be chosen only from existing stores; every node except the national centre has exactly one inbound link; hubs are replenished monthly or quarterly as per WHO guidelines; each open facility uses one device from the WHO pre-qualified list; and clinic demand is inflated by a 25% safety buffer. Only total vaccine volume is modelled, not individual vaccines.

3.2 P2: A Bi-objective Robust Network under Disruption (Mohammadi et al., 2022)

Type of problem. A multi-period, multi-vaccine, five-level distribution network design problem formulated as a bi-objective stochastic mixed integer non-linear

programme (MINLP). It extends the transshipment problem by allowing horizontal transshipment between warehouses at the same level. The full model contains many index sets; a simplified representation is given here.

Sets. Vaccines V , regions R , departments D , candidate vaccination centres P , demand points (cities) I , risk classes K , periods (weeks) T , capacity levels L and scenarios S .

Decision variables. Binary variables Y for opening or equipping national, regional and departmental warehouses and vaccination centres at a given capacity level; continuous flows X_{vt} along each level of the network and transshipments $\bar{X}_{vt}$ between regions or between departments; inventories I_{vt} at each storage point; order quantities O_{vt} placed with manufacturers; and U_{kipvt} , the number of people of risk class k from city i vaccinated with a first or second dose of vaccine v at centre p in week t .

Objective functions. Under each scenario s , the first objective is total cost, and the second is the expected number of deaths:

$$\text{Min } Z_1^s = C_{order}^s + C_{fixed}^s + C_{transport}^s + C_{hold}^s + C_{vacc}^s + C_{hosp}^s$$

$$\text{Min } Z_2^s = \sum_{t \in T} \sum_{i \in I} \sum_{k \in K} \eta_{ks} Inf_{kit}^s$$

The six cost terms are, in order, the fixed and variable cost of ordering vaccines; the fixed cost of equipping warehouses and opening vaccination centres; the cost of transferring and transshipping vaccines; inventory holding cost; the cost of administering doses; and the expected hospitalisation cost of unvaccinated or single-dose individuals who become infected. In the second objective, η_{ks} is the death rate of class k and Inf_{kit}^s is the number of infections in class k , city i and week t , calculated separately for the unvaccinated, single-dose and double-dose groups. Because infections depend on contacts between these three groups, Z_2 contains products of decision variables and is non-linear.

Main constraints. (i) Inventory balance at every warehouse and centre, with inventory and capacity degraded by disruption coefficients in disrupted

scenarios; (ii) population balance, tracking unvaccinated, single-dose and double-dose individuals each week, with a second dose only after the required gap of l_v weeks; (iii) capacity limits on orders, storage and flows; and (iv) a congestion constraint in which each vaccination centre is modelled as a GI/G/c queue and the mean waiting time must not exceed a threshold:

$$W_{pts} \leq \bar{W}_p \quad \forall p \in P, t \in T, s \in S$$

Robust counterpart. Uncertainty is represented by eight scenarios formed from three yes-or-no factors: quarantine (Q), a new variant (V) and network disruption (D). The cost objective is converted into a p -robustness constraint (Snyder & Daskin, 2006), so that the cost in each scenario is no more than $100(1 + p_s)\%$ of that scenario's optimal cost, and the deaths' objective is minimised, following the robust optimisation approach of Mulvey, Vanderbei and Zenios (1995), as a weighted sum of its expected value and its variability across scenarios:

$$Z_1^s \leq (1 + p_s) Z_1^{s*} \quad \forall s \in S$$

$$\text{Min} \sum_{s \in S} \pi_s Z_2^s + \omega \sigma(Z_2^1, \dots, Z_2^{|S|})$$

where π_s is the probability of scenario s and ω is the weight placed on variability.

Assumptions. The population is grouped into four age-based risk classes; decisions are taken weekly; completely protected (double-dose) individuals do not need hospitalisation; scenario probabilities are known; and the 300 most populous cities and 500 candidate centres adequately represent the country.

3.3 P3: Capacitated Vehicle Routing with Heterogeneous Vehicles (Andoh & Yu, 2023)

Type of problem. A capacitated vehicle routing problem (CVRP) with a heterogeneous fleet, formulated as an integer linear programme based on Kilby and Shaw (2006). Unlike the transportation problem, which decides only how much flows between each pair of points, the CVRP decides the sequence in which each vehicle visits customers.

Sets and parameters. $V = \{0, 1, \dots, v\}$ is the set of nodes, where node 0 is the vaccine warehouse and the rest are municipal vaccination centres; P is the set of vehicles indexed by k ; c_k is the unit transport cost of

vehicle k ; d_{ij} is the distance between nodes i and j ; q_i is the demand of centre i ; Q_k is the capacity of vehicle k ; and D_k is its maximum travel distance per tour.

Decision variable. $x_{ijk} = 1$ if vehicle k travels directly from node i to node j , and 0 otherwise.

Objective function.

$$\text{Minimise} \sum_{i \in V} \sum_{j \in V} \sum_{k \in P} c_k d_{ij} x_{ijk}$$

Constraints.

$$\sum_{i \in V} \sum_{k \in P} x_{ijk} = 1 \quad \forall j \in V \setminus \{0\} \quad (\text{each centre entered once})$$

$$\sum_{j \in V} \sum_{k \in P} x_{ijk} = 1 \quad \forall i \in V \setminus \{0\} \quad (\text{each centre left once})$$

$$\sum_{i \in V} x_{ihk} - \sum_{j \in V} x_{hjk} = 0 \quad \forall h \in V, k \in P \quad (\text{route continuity})$$

$$\sum_{i \in V} \sum_{j \in V} q_i x_{ijk} \leq Q_k \quad \forall k \in P \quad (\text{vehicle capacity})$$

$$\sum_{i \in V} \sum_{j \in V} d_{ij} x_{ijk} \leq D_k \quad \forall k \in P \quad (\text{maximum distance per tour})$$

$$\sum_{i \in W} \sum_{j \in W} \sum_{k \in P} x_{ijk} \leq |W| - 1 \quad \forall W \subseteq V \setminus \{0\} \quad (\text{subtour elimination})$$

$$x_{ijk} \in \{0, 1\}$$

The authors note that the capacity constraint can be adapted by setting $q_i = 1$, so that the right-hand side becomes the maximum number of centres a vehicle may visit per tour; in the case study this was set to 10 centres per trip to keep deliveries responsive.

Assumptions. Time windows are deliberately left out of the optimisation because they would make the NP-hard CVRP much slower to solve (Baldacci et al., 2012); they are instead tested in the simulation stage. In the simulation, outbound handling at the depot follows a uniform distribution between 20 and 30

minutes, inbound handling at centres between 10 and 20 minutes, vehicle speed follows a triangular distribution (minimum 40, mode 50, maximum 70 km/h), shipping takes place Monday to Friday between 6:00 am and 3:00 pm, electric vehicles have zero tailpipe emissions, and fleet maintenance and purchase costs are ignored because deliveries occur only once or twice a week.

3.4 P4: A Bi-objective Vehicle Routing Model for Risk and Time (Zabinsky et al., 2021)

Type of problem. A vehicle routing and scheduling problem with two objectives, transit time and risk, combined through user-chosen weights. The paper describes its model in words (its Box 2 and Box 3) rather than equations; the formulation below translates those stated objectives and constraints into standard VRP notation.

Sets and parameters. Node 0 is the distribution warehouse and N is the set of health facilities; K is the set of vehicle trips, each assigned to an available vehicle; t_{ij} is the travel time from i to j ; π_{ij} is a penalty for the road condition between i and j (fully paved, partially paved, good dirt road or rough dirt road), following the penalty approach of Hamedi et al. (2012); ϕ_k is a penalty for the condition of the vehicle used on trip k ; q_j^c and q_j^a are the cold-chain and ambient demand of facility j ; Q_k^c and Q_k^a are the matching vehicle capacities; and $T_{\max}$ is the maximum route duration (at most 24 hours).

Decision variable. $x_{ijk} = 1$ if trip k travels from i to j . Objective function. The user chooses a time weight W_t between 0 and 10, and the risk weight is $W_p = 10 - W_t$:

$$\begin{aligned} \text{Minimise } W_t \sum_k \sum_i \sum_j t_{ij} x_{ijk} \\ + W_p \sum_k \sum_i \sum_j (\pi_{ij} + \phi_k) x_{ijk} \end{aligned}$$

Setting $W_t = 10$ minimises transit time only, $W_t = 0$ minimises risk only, and any value in between balances the two.

Constraints.

$$\begin{aligned} \sum_i \sum_k x_{ijk} = 1 \quad \forall j \\ \in N \quad (\text{visit each facility once}) \end{aligned}$$

$$\begin{aligned} \sum_j x_{0jk} = \sum_i x_{i0k} \\ \leq 1 \quad \forall k \quad (\text{start and end at warehouse}) \end{aligned}$$

$$\sum_i \sum_j q_j^c x_{ijk} \leq Q_k^c \quad \forall k \quad (\text{cold-chain capacity})$$

$$\sum_i \sum_j q_j^a x_{ijk} \leq Q_k^a \quad \forall k \quad (\text{ambient capacity})$$

$$\begin{aligned} \sum_i \sum_j t_{ij} x_{ijk} \\ \leq T_{\max} \quad \forall k \quad (\text{single-day route limit}) \end{aligned}$$

$$x_{ijk} = 0 \quad \text{if road } (i, j) \text{ is marked not accessible}$$

In addition, all demand must be met, trips are spread across available vehicles as evenly as possible, and the usual route continuity and subtour elimination conditions apply.

Assumptions. Routes are completed within one day (at most 24 hours); vehicles start and return to the same warehouse; the warehouse has enough stock to meet all requested demand; facilities can store whatever they request (the tool warns, but does not stop, when demand exceeds storage); and routes use roads only, so water or foot transport is not modelled. All products are aggregated into two groups, cold-chain and ambient, which keeps computation fast.

3.5 P5: A Multi-Depot Vehicle Routing Model (Rodrigues & Lima, 2021)

Type of problem. A multi-depot vehicle routing problem (MDVRP), formulated as an integer programme. It extends the CVRP by allowing more than one starting depot (Renaud et al., 1996; Montoya-Torres et al., 2015) (for example, Recife and Petrolina).

Sets and parameters. $V_c = \{v_1, \dots, v_N\}$ is the set of cities, $V_d = \{v_{N+1}, \dots, v_{N+M}\}$ is the set of depots, K is the number of vehicles, c_{ij} is the distance between nodes, P_k is the capacity of vehicle k , and for each city p_i , γ_i and σ_i are its population, active COVID-19 cases and COVID-19 deaths, with state totals p_{total} , γ_{total} and σ_{total} . ϕ_i is the number of available doses of each vaccine type in the set Φ .

Decision variable. $x_{ijk} = 1$ if vehicle k travels from node i to node j .

Objective function. Minimise the total distance travelled:

$$\text{Minimise } \sum_{i=1}^{N+M} \sum_{j=1}^{N+M} \sum_{k=1}^K c_{ij} x_{ijk}$$

Constraints. Each city is served by exactly one vehicle (assignment constraints, as in P3) and route continuity is preserved. The authors' main contribution is the demand constraint, which allocates the available doses to each city in proportion to the average of its shares of population, active cases and deaths:

$$Q_i = \frac{1}{3} \left(\frac{p_i}{p_{total}} + \frac{\gamma_i}{\gamma_{total}} + \frac{\sigma_i}{\sigma_{total}} \right) \sum_{\phi} \phi_i \quad \forall i = 1, \dots, N$$

$$\sum_{i=1}^{N+M} \sum_{j=1}^{N+M} Q_i x_{ijk} \leq P_k \quad \forall k \quad (\text{vehicle capacity})$$

$$\sum_{i=N+1}^{N+M} \sum_{j=1}^N x_{ijk} \leq 1 \quad \forall k \quad (\text{leaves a depot at most once})$$

$$\sum_{j=N+1}^{N+M} \sum_{i=1}^N x_{ijk} \leq 1 \quad \forall k \quad (\text{returns to a depot at most once})$$

$$y_i - y_j + (M + N) \sum_k x_{ijk} \leq N + M - 1 \quad (\text{subtour elimination})$$

$$x_{ijk} \in \{0,1\}$$

Assumptions. Distances between cities are straight-line (Euclidean) distances computed from map coordinates; all vehicles are identical refrigerated trucks carrying 24 coolers, which works out to 144,000 doses each; a 5% margin is added to demand for possible damage in transit; all doses are treated as identical items; and the second dose is handled by re-running the model with new demand.

3.6 Summary of the Formulations

Paper	Problem type	Objective	Key constraints
P1	Hub location and transshipment MIP	Min facility and transport cost	Single supplier, flow balance, storage and vehicle capacity, WHO frequencies
P2	Bi-objective stochastic MINLP	Min deaths and cost	Inventory and population balance, disruption, queue waiting time
P3	Heterogeneous CVRP	Min transport cost	Visit once, capacity, maximum distance, subtours
P4	Weighted bi-objective VRP	Min weighted time and risk	Visit once, cold and ambient capacity, 24-hour limit, closed roads
P5	Multi-depot VRP	Min total distance	Visit once, proportional demand, capacity, subtours

IV. SOLUTIONS DERIVED AND METHODS USED

4.1 P1: Disaggregation-and-Merging Algorithm

The authors first tried to solve MIP-1 directly with a commercial solver. For a typical 100-node problem with 15 candidate hubs, the model already has close to 10,000 binary variables, and one full-country instance with 685 nodes and 41 candidate hubs has 168,838 integer variables. As a rule of thumb, problems with more than about 200 nodes and 15 potential hubs could

not be solved directly. The authors therefore developed a divide-and-conquer heuristic based on a simple observation: the best network structure in one part of a country is largely independent of the structure in a distant part.

The algorithm works in stages. The potential hubs are divided into clusters by hierarchical clustering, each cluster is solved optimally as a small MIP, and the clusters are then merged one at a time, starting with the one nearest the national centre. At each merger, hubs are classified as critical (near the boundary between

the existing network and the new region, so left free to change), intermediate (between the national centre and the new region, so they may become transshipment points) or non-critical (far away, so their decisions are fixed). A “shrinkage” refinement replaces all clinics served by a fixed hub with a single dummy clinic at the hub, which sharply reduces the number of binary variables.

The method was tested on 43 problems drawn from four sub-Saharan African countries. Where the optimum was known (28 instances), the algorithm found the optimal solution in 22 cases and came within 0.5% of the optimum in five of the remaining six; the worst gap was 0.69%. It solved almost every problem in under about five minutes, and even the largest instance, with over 2 million variables, took about 12 minutes. Compared with the existing national networks, the redesigned networks reduced storage and transport costs by 24.84%, 5.97%, 27.10% and 9.28% in Countries A, B, C and D respectively.

4.2 P2: Linearisation and Scenario-based Robust Optimisation

The non-linear waiting-time constraint was replaced by an equivalent linear constraint using an enumeration technique, after the authors proved that the underlying function is convex, and the products of variables in the death's objective were linearised. Uncertainty was handled through two robust optimisation approaches: p -robustness, which caps the regret in each scenario, for the cost objective, and a loss-function approach, which penalises both the average and the spread of outcomes, for the death's objective. Turning cost into a constraint reflects the authors' view that a vaccine network is not a business network and that saving lives should come first.

The solution showed that, depending on the scenario, between 65% and 87% of the 500 candidate

vaccination centres should be opened, while the fully robust plan opens all of them at the higher capacity level to allow mass vaccination. Quarantine restrictions reduced network cost by between 4% and 12% across the scenario pairs. Protecting the network against disruption added between 19.5% and 44.5% to cost, with the largest premium when a new variant and disruption occurred together. When disruption was likely, the model ordered more vaccines earlier and favoured suppliers with higher reliability. Allocation to regions depended not only on population but also on contacts, risk classes, variants and quarantine rules, so the busiest region, Île-de-France, received the most doses. Some departments with spare storage acted as intermediate hubs for their neighbours, reducing transport cost. The authors report that combining population prioritisation with quarantine restrictions reduces total deaths by 8.5%.

4.3 P3: Two-Stage Optimisation and Simulation

In the first stage, the CVRP was solved for each network design to obtain optimal routes (milk runs) and vehicle assignments. The location of a candidate distribution centre at Hamar was chosen by greenfield (centre-of-gravity) analysis on real road networks. In the second stage, these routes were exported into a discrete event simulation with stochastic handling times and vehicle speeds, working-hour limits and expected lead times of one day for nearby centres and two days for distant ones. Monte Carlo simulation repeated each scenario 30 times to make the results statistically reliable. Ten scenarios were tested across three models: Model 1 (Oslo depot, mixed electric and fuel fleet), Model 2 (new Hamar depot, mixed fleet) and Model 3 (Oslo depot, fuel vehicles only, including larger VW Transporters).

Scenario	Model	Fleet	Cost (NOK)	CO ₂ (kg)	On-time service level
1	1	3	567,802	3,380	42.3%
2	1	5	605,802	3,432	94.6%
3	1	6	601,099	3,380	100%
4	2	3	448,227	227	51.1%
7	2	7	476,930	227	100%
8	3	3	705,387	8,327	65.7%
10	3	6	726,565	8,522	100%

The results show that the Hamar depot (Model 2) gives by far the lowest cost and emissions. Increasing the fleet from three to five vehicles in Model 1 raised on-time delivery from 42.3% to 94.6% and cut average lead time from 1.229 to 0.67 days, with only a modest change in cost. The all-fuel fleet in Model 3 produced about 148% more CO₂ than the mixed fleet in Model 1, because electric vans could serve the nearby centres.

4.4 P4: The VeRSA Heuristic inside an Excel Tool

Because users needed results within minutes rather than hours, the authors built their own algorithm, a variant of the Vehicle Routing and Scheduling Algorithm (VeRSA) first developed for collecting medical specimens (Zabinsky et al., 2020). VeRSA embeds an indexing rule inside a branch-and-bound framework, which lets it construct a good feasible solution in seconds; specific constraints are checked directly within its feasibility test to speed up the search. Products are aggregated into cold-chain and ambient groups during optimisation and split back into individual items in the output, so adding more products does not slow the tool. The default run time is two minutes, and the tool is recommended for up to 50 facilities and five vehicles.

The solution produced by RoOT is an Excel file listing each vehicle's route, departure times from each

facility, vehicle capacity utilisation, the quantity of each product delivered, and the fuel and per diem cost of each route. In a simulated cyclone, the user changed the depot, set demand to zero at damaged facilities, raised demand at others and marked a road as not accessible; RoOT returned three new routes for two vehicles within two minutes, none of which used the blocked road.

4.5 P5: GRASP-VND Metaheuristic

The authors combined two metaheuristics. GRASP (Greedy Randomized Adaptive Search Procedure; Feo & Resende, 1995) builds a starting set of routes by repeatedly adding a city chosen at random from a restricted candidate list of the best options, where a parameter α controls how greedy or random the choice is. VND (Variable Neighbourhood Descent; Mladenović & Hansen, 1997) then improves the routes using eleven neighbourhood operators: five within a route (or-opt1, or-opt2, or-opt3, swap and 2-opt) and six between routes (shift (1,0), shift (2,0), shift(3,0), swap(1,1), swap(2,2) and swap(3,3)). The process is repeated many times and the best solution is kept.

Three cases were compared with the current system, using 184 municipalities and a real batch of 529,980 doses received on 2 July 2021.

Case	Depots and vehicles	Best distance (km)	Average run time (s)
Current practice	Municipalities collect from 12 GERES	14,434	–
1	Recife only, 4 vehicles	3,974	2,870.11
2	Recife and Petrolina, 4 vehicles	3,513	384.49
3	12 GERES, 12 vehicles	3,655	37.35

Case 2 was best, saving almost 11,000 km, which the authors describe as routing 75.66% more efficient than current practice. Case 3 was the most stable across runs (standard deviation of only 8.48 km) because each GERES solves a much smaller problem.

V. SOFTWARE USED AND VALIDATION OF THE SOLUTIONS

P1. The MIP-1 benchmark was solved with IBM ILOG CPLEX 12.6 on a computer with an Intel Core i5 processor and 8 GB of memory. The algorithm was validated in two ways. First, its solutions were compared against the true optimum on the 28 instances

that CPLEX could solve. Second, for full-country networks where no optimum is available, its cost was compared with the cost of the existing network using the same unit costs, and the network structures were inspected visually for reasonableness.

P2. The robust model was coded in Python 3.7 and solved with the Gurobi 9.1 solver using the barrier method, on a machine with eight processors and 32 GB of memory; one solution took around 20 hours. The model was validated on a real case: metropolitan France, with about 65.2 million people across 12 regions and 96 departments, 300 demand cities, 500 candidate centres, four approved vaccines (Pfizer-

BioNTech, Moderna, AstraZeneca and Johnson & Johnson) and four age-based risk classes. Results were compared across eight scenarios and against the actual cost of each scenario, followed by sensitivity analyses.

P3. The study used anyLogistix, a commercial supply chain package that combines optimisation and simulation. Its transportation optimisation module solves the CVRP with the built-in CPLEX solver, and its simulation module, built on AnyLogic, runs the discrete event and Monte Carlo simulations on real GIS road maps. Validation used real demand data from the Norwegian Institute of Public Health for 28 December 2020 to 15 March 2021, supplemented by interviews with stakeholders, and each scenario was replicated 30 times so that the results did not depend on a single random draw.

P4. RoOT's VerSA algorithm is coded in Python but reads and writes Microsoft Excel files, so the user never sees the code. It runs on 64-bit Windows with Excel 2007 or later and is free and open-source on GitHub under the GNU General Public License, in English and Portuguese. Validation was both technical and practical. Numerically, as reported in a companion paper (Petroianu et al., 2021), it matched the best

commercial software for 10 to 20 facilities and performed better within two minutes for 50 facilities, though not as well as commercial software run for several hours on 100 facilities. Practically, its interface was designed using Nielsen's usability heuristics (Nielsen & Mack, 1994), and it was tested with eight provincial and district logisticians in Mozambique, seven of whom felt confident using it to choose routes and vehicles; one described their previous planning as "working in the dark". The paper was peer reviewed and "approved with reservations" by three reviewers.

P5. The metaheuristic was implemented in Python 3 on a computer with an Intel Core i7 processor and 8 GB of memory. Municipal coordinates were obtained from the official map of Pernambuco using QGIS, an open-source geographic information system. The model was validated through a numerical application using real dose quantities and demand indicators, comparing the optimised routes with the current collection system. Each case was run 30 times and the best, average and standard deviation of results were reported. The authors could not contact the government to test the routes in practice and could not find an exact algorithm able to solve problems of the same size for comparison.

Paper	Software	How the solution was validated
P1	IBM ILOG CPLEX 12.6	Compared with the optimum on 28 instances; compared with existing network costs in four countries
P2	Python 3.7, Gurobi 9.1	Real case study of France; eight scenarios; sensitivity analysis
P3	anyLogistix (CPLEX and AnyLogic)	Real Norwegian data; 10 scenarios with 30 simulation replications each
P4	Python behind Microsoft Excel	Benchmark against commercial software; user testing with eight logisticians; cyclone use case
P5	Python 3, QGIS	Numerical application with real doses; 30 runs per case; comparison with current practice

VI. CRITICAL ANALYSIS: SUMMARY, KEY FINDINGS AND IMPACT

6.1 Summary of Key Findings

Across five very different settings, the papers reach a consistent set of conclusions. The first and strongest is that the shape of the network matters more than any single routing decision. P1 shows that redesigning legacy national networks saves between 6% and 27% of storage and transport cost; P3 finds that moving the

depot from Oslo to Hamar cuts both cost and emissions dramatically; and P5 shows that switching from municipal collection to two central depots cuts total distance by about three quarters. In transportation-problem terms, deciding where supply should sit, and therefore what the cost matrix looks like, has a larger effect than optimising flows over a fixed network.

The second finding is that vehicles are a strategic choice. In P3, adding two vehicles raised on-time

delivery from 42% to 95%, and using electric vans for short trips kept emissions low; in P4, routes depend on which vehicles are available and in what condition. The third finding is that trade-offs are unavoidable. P2 balances deaths against cost, P3 balances service, cost and carbon, and P4 balances speed against risk. None of these problems has a single best answer, only a best answer for the priorities a decision-maker chooses. The fourth finding is that uncertainty must be planned for. P2 shows that disruption combined with a new variant can raise costs by 44.5%, and P3 shows that routes that look optimal can still miss delivery windows once real-world variability is introduced. The final finding is that solution speed and usability matter as much as mathematical precision. P1, P4 and P5 all accept near-optimal answers in exchange for results in minutes, and P4 shows that a tool non-specialists can use is more valuable than a theoretically better one nobody adopts.

6.2 Impact on the Industry

For health ministries and immunisation programmes, the papers provide evidence that legacy distribution networks can be redesigned for lower cost without compromising access, and they give practical methods for doing so at national scale. The finding in P2 that departments with spare storage can serve as intermediate hubs for their neighbours suggests that planned transshipment and resource sharing between stores should be built into national plans rather than improvised during shortages.

For logistics providers, including the 3PL companies that carried COVID-19 vaccines in Norway, P3 offers evidence for fleet decisions: mixing electric and fuel vehicles according to distance cuts emissions while keeping costs reasonable, and simulating fleet size, depot location and vehicle mix before committing to them is far cheaper than learning from failures on the road. The same insights help governments select and evaluate 3PL providers.

For NGOs and governments in low-resource settings, P4 demonstrates that free, open-source tools running on familiar software such as Excel can compete with expensive commercial packages for small and medium routing problems, lowering the barrier to evidence-based planning. P5 shows that a state government with access to standard computers and open-source GIS can plan routes centrally rather than relying on each municipality to collect its own doses. The same

models are relevant to India, whose Universal Immunisation Programme moves vaccines from the national vaccine store through state warehouses and district cold chain points to health centres, and whose eVIN platform monitors stock and temperature across thousands of cold chain points (Tschan & Saxena, 2026). Such a network faces the same questions of hub selection, routing and last-mile reliability.

6.3 Benefits to Society

The ultimate purpose of these models is not cheaper logistics but better health. Shorter, safer and better-planned routes protect the cold chain, so fewer doses spoil and more of the vaccines that reach clinics are still potent. Reliable deliveries reduce stockouts, so clinics turn fewer people away. When vaccines are delivered directly instead of collected, health workers spend less time travelling and more time providing care, which is one of the motivations behind both P4 and P5. Lower emissions from smarter routing and electric vehicles, as in P3, reduce the environmental cost of mass vaccination campaigns.

The models also support fairness and resilience. P1 explicitly guarantees that every clinic is served, P5 allocates doses according to population, active cases and deaths, and P2 shows that equity is better judged by potential deaths than by population alone, adding to earlier work on equitable vaccine allocation (Enayati & Özaltın, 2020). P4 and P2 both show how planning tools can respond quickly to cyclones, supply delays or new variants, so that remote and vulnerable communities are not the first to lose access in a crisis. P2's estimate that combining prioritisation with quarantine measures could reduce deaths by 8.5% illustrates, in the most direct terms, that better OR models can save lives.

VII. REFLECTION: STRENGTHS, LIMITATIONS AND ALTERNATIVE POSSIBILITIES

7.1 Strengths of the Models

Each model is well matched to its setting. P1's main strength is scale: it turns a national network design problem that commercial solvers cannot handle into a sequence of small problems, and it delivers near-optimal results in minutes while respecting WHO operating rules that keep the network simple for staff to run. P2 is the most comprehensive of the five. It is the only model that links logistics to health outcomes

through an epidemiological component, handles several vaccines with one or two doses, models queues at vaccination centres and builds in disruption, which makes its plans robust across many possible futures. P3's strength is realism. Combining optimisation (which picks good routes) with simulation (which checks whether they survive real conditions) is a sound design (Tordecilla et al., 2021), and measuring societal, economic and environmental performance together gives a fuller picture than cost alone. P4 stands out for usability. It is built around the real constraints of government users, models risk rather than distance alone, and is free and open-source. P5 shows that a well-designed metaheuristic can produce large savings quickly, and its demand formula allocates doses using public health indicators rather than population alone.

7.2 Limitations and Assumptions

Every model rest on simplifying assumptions, and some of these limits how far the results can be relied upon. Demand is treated as known in P1, P4 and P5, although real vaccine supply and demand were highly uncertain, particularly during COVID-19. P1 also fixes a 25% safety buffer and relies on the assumption that distant parts of the network are independent, which may not hold if costs differ sharply between regions.

Some models optimise only one dimension of what matters. The routing stage in P3 minimises cost only, and temperature risk, the very thing that spoils vaccines, is not modelled directly. P5 minimises total distance, which leaves some vehicles with much longer routes than others, uses straight-line rather than road distances, ignores costs such as fuel and refrigeration, and assumes an identical fleet without time windows.

Evidence of real-world performance is limited. P5 is a numerical application that was never tested with the government, and its results were not compared with an exact method. RoOT in P4 was never fully deployed because of the COVID-19 pandemic, so there is no evidence yet that it changes behaviour or improves outcomes; its claim to outperform commercial software is supported by results reported in a separate publication (Petroianu et al., 2021); it plans single-day trips only, has no map view and depends on a hand-filled distance matrix that becomes difficult as facilities multiply. P3 is a single case study in a

wealthy, well-organised country and depends on commercial software that the authors acknowledge struggles as the number of customers grows. P2 needed around 20 hours to solve, reports results only as percentages, groups people only by age, and assumes that scenario probabilities are known. Finally, all five models depend on good data, and in many LMICs even basic facility locations, distances and road conditions are missing or inaccurate.

7.3 Alternative Possibilities for Improvement

Several improvements follow naturally from these limitations. First, the classical transportation and transshipment models could be used more explicitly as a fast first step, and simple constructive heuristics such as the Clarke and Wright (1964) savings algorithm could give planners a quick, explainable starting set of routes (Cordeau et al., 2002). Solving an allocation problem with VAM and MODI, or with linear programming, would decide how many doses each store should send to each facility before a routing model decides how to deliver them. This two-level approach is cheap to compute and easy to explain to non-specialists.

Second, location and routing decisions could be combined in a single location-routing model, since P1, P3 and P5 all show that depot location and routing interact strongly. Third, time windows and temperature exposure could be added as explicit constraints or objectives (Rodrigues & Lima, 2021), using vehicle routing with time windows (VRPTW; Baldacci et al., 2012) or a risk term similar to RoOT's, so that cold chain safety is optimised directly rather than indirectly. Fourth, uncertainty could be built into routing through stochastic or robust VRP models, or by pairing RoOT's speed with a simulation check of the kind used in P3.

Fifth, multi-objective trade-offs could be presented as a set of options rather than a single answer. Methods such as the ϵ -constraint approach or goal programming (Taha, 2017) would show decision-makers the full trade-off curve between cost, service, emissions and risk. Sixth, distance matrices could be generated automatically from open map services, and RoOT-style tools could be integrated with mapping software to reduce manual data entry and provide the route visualisation that users requested (Zabinsky et al., 2021). Seventh, other metaheuristics such as tabu search, simulated annealing, genetic algorithms or

adaptive large neighbourhood search could be benchmarked against GRASP-VND (Rodrigues & Lima, 2021), and open-source solvers could replace commercial packages to make advanced methods affordable for public health agencies. Finally, as P2's authors suggest (Mohammadi et al., 2022), epidemiological models such as SIR could be linked more closely to distribution decisions, as Gamchi, Torabi and Jolai (2021) do for vaccine routing, and digital tools such as real-time tracking could provide the visibility needed to re-plan routes when roads close or vehicles break down.

VIII. CONCLUSION

Producing a vaccine is only half the battle; a dose that spoils on a rough road, or arrives a day late, protects no one. The five papers reviewed here show how the transportation problem, one of the oldest models in Operations Research, can be extended to meet this challenge. P1 and P2 extend it into transshipment and network design, deciding where vaccines should be stored and how much should flow along each link, while P3, P4 and P5 extend it into vehicle routing, deciding which vehicle should visit which facilities and in what order. Each adapts the basic model to its setting: decomposition for very large networks in Africa, robust optimisation for an uncertain pandemic in France, simulation for sustainability in Norway, a light Excel tool for government users in Mozambique, and a metaheuristic for fast route planning in Brazil. Together, they show that network design and fleet choices have the largest effect on cost and service, that trade-offs between cost, speed, risk, emissions and lives must be made explicit, and that a model only creates value if it is fast enough and simple enough to be used. Their main weaknesses are deterministic assumptions, limited real-world testing and incomplete treatment of temperature risk. The best last-mile system would combine RoOT's simplicity with the rigour of the Norwegian and French studies, and would be judged by one outcome: whether more people receive vaccines that still work.

REFERENCES

- [1] E. A. Andoh and H. Yu, "A two-stage decision-support approach for improving sustainable last-mile cold chain logistics operations of COVID-19 vaccines," *Ann. Oper. Res.*, vol. 328, pp. 75–105, 2023, doi: 10.1007/s10479-022-04906-x.
- [2] R. Baldacci, A. Mingozzi, and R. Roberti, "Recent exact algorithms for solving the vehicle routing problem under capacity and time window constraints," *Eur. J. Oper. Res.*, vol. 218, no. 1, pp. 1–6, 2012.
- [3] J. F. Campbell and M. E. O'Kelly, "Twenty-five years of hub location research," *Transp. Sci.*, vol. 46, no. 2, pp. 153–169, 2012, doi: 10.1287/trsc.1120.0410.
- [4] S. I. Chen, B. A. Norman, J. Rajgopal, T. M. Assi, B. Y. Lee, and S. T. Brown, "A planning model for the WHO-EPI vaccine distribution network in developing countries," *IIE Trans.*, vol. 46, pp. 853–865, 2014, doi: 10.1080/0740817X.2013.813094.
- [5] G. Clarke and J. W. Wright, "Scheduling of vehicles from a central depot to a number of delivery points," *Oper. Res.*, vol. 12, no. 4, pp. 568–581, 1964.
- [6] J. F. Cordeau, M. Gendreau, G. Laporte, J. Y. Potvin, and F. Semet, "A guide to vehicle routing heuristics," *J. Oper. Res. Soc.*, vol. 53, no. 5, pp. 512–522, 2002, doi: 10.1057/palgrave.jors.2601319.
- [7] G. B. Dantzig and J. H. Ramser, "The truck dispatching problem," *Manage. Sci.*, vol. 6, no. 1, pp. 80–91, 1959.
- [8] S. Dey, A. K. Kurbanzade, E. S. Gel, J. Mihaljevic, and S. Mehrotra, "Optimization modeling for pandemic vaccine supply chain management: A review and future research opportunities," *arXiv preprint arXiv:2305.10942*, 2023.
- [9] L. E. Duijzer, W. van Jaarsveld, and R. Dekker, "Literature review: The vaccine supply chain," *Eur. J. Oper. Res.*, vol. 268, pp. 174–192, 2018.
- [10] S. Enayati and O. Y. Özaltın, "Optimal influenza vaccine distribution with equity," *Eur. J. Oper. Res.*, vol. 283, no. 2, pp. 714–725, 2020.
- [11] R. Z. Farahani, M. Hekmatfar, A. B. Arabani, and E. Nikbakhsh, "Hub location problems: A review of models, classification, solution techniques, and applications," *Comput. Ind. Eng.*, vol. 64, pp. 1096–1109, 2013, doi: 10.1016/j.cie.2013.01.012.
- [12] T. A. Feo and M. G. C. Resende, "Greedy randomized adaptive search procedures," *J.*

- Global Optim.*, vol. 6, no. 2, pp. 109–133, 1995, doi: 10.1007/BF01096763.
- [13] N. S. Gamchi, S. A. Torabi, and F. Jolai, “A novel vehicle routing problem for vaccine distribution using SIR epidemic model,” *OR Spectrum*, vol. 43, no. 1, pp. 155–188, 2021.
- [14] M. Hamed, A. Haghani, and S. Yang, “Reliable transportation of humanitarian supplies in disaster response: Model and heuristic,” *Procedia-Social and Behavioral Sciences*, vol. 54, pp. 1205–1219, 2012.
- [15] F. L. Hitchcock, “The distribution of a product from several sources to numerous localities,” *J. Math. Phys.*, vol. 20, pp. 224–230, 1941.
- [16] D. Ivanov, “Predicting the impacts of epidemic outbreaks on global supply chains: A simulation-based analysis on the coronavirus outbreak (COVID-19/SARS-CoV-2) case,” *Transp. Res. E: Logistics Transp. Rev.*, vol. 136, Art. no. 101922, 2020.
- [17] U. H. Kartoglu, K. L. Moore, and J. S. Lloyd, “Logistical challenges for potential SARS-CoV-2 vaccine and a call to research institutions, developers and manufacturers,” *Vaccine*, vol. 38, pp. 5393–5395, 2020.
- [18] P. Kilby and P. Shaw, “Vehicle routing,” in F. Rossi, P. van Beek, and T. Walsh, Eds., *Handbook of Constraint Programming*. Elsevier, 2006.
- [19] G. Laporte, “The vehicle routing problem: An overview of exact and approximate algorithms,” *Eur. J. Oper. Res.*, vol. 59, no. 3, pp. 345–358, 1992, doi: 10.1016/0377-2217(92)90192-C.
- [20] G. Laporte, “Fifty years of vehicle routing,” *Transp. Sci.*, vol. 43, no. 4, pp. 408–416, 2009.
- [21] T. H. Lee and A. H. Chen, “Last-mile logistics of COVID vaccination: The role of health care organizations,” *N. Engl. J. Med.*, vol. 384, pp. 685–687, 2021.
- [22] N. Mladenović and P. Hansen, “Variable neighborhood search,” *Comput. Oper. Res.*, vol. 24, no. 11, pp. 1097–1100, 1997, doi: 10.1016/S0305-0548(97)00031-2.
- [23] M. Mohammadi, M. Dehghan, A. Pirayesh, and A. Dolgui, “Bi-objective optimization of a stochastic resilient vaccine distribution network in the context of the COVID-19 pandemic,” *Omega*, vol. 113, Art. no. 102725, 2022, doi: 10.1016/j.omega.2022.102725.
- [24] J. R. Montoya-Torres, J. L. Franco, S. N. Isaza, H. F. Jiménez, and N. Herazo-Padilla, “A literature review on the vehicle routing problem with multiple depots,” *Comput. Ind. Eng.*, vol. 79, pp. 115–129, 2015, doi: 10.1016/j.cie.2014.10.029.
- [25] J. M. Mulvey, R. J. Vanderbei, and S. A. Zenios, “Robust optimization of large-scale systems,” *Oper. Res.*, vol. 43, no. 2, pp. 264–281, 1995.
- [26] J. Nielsen and R. L. Mack, *Usability Inspection Methods*. John Wiley & Sons, 1994.
- [27] L. P. G. Petroianu, Z. B. Zabinsky, M. Zameer, et al., “A light-touch routing optimization tool (RoOT) for vaccine and medical supply distribution in Mozambique,” *Int. Trans. Oper. Res.*, vol. 28, no. 5, pp. 2334–2358, 2021.
- [28] J. Renaud, G. Laporte, and F. F. Boctor, “A tabu search heuristic for the multi-depot vehicle routing problem,” *Comput. Oper. Res.*, vol. 23, no. 3, pp. 229–235, 1996.
- [29] A. J. S. Rodrigues and G. L. Lima, “A metaheuristic to support the distribution of COVID-19 vaccines,” *Production*, vol. 31, Art. no. e20210031, 2021, doi: 10.1590/0103-6513.20210031.
- [30] L. V. Snyder and M. S. Daskin, “Stochastic p-robust location problems,” *IIE Trans.*, vol. 38, no. 11, pp. 971–985, 2006.
- [31] H. A. Taha, *Operations Research: An Introduction*, 10th ed. Pearson, 2017.
- [32] R. D. Tordecilla, A. A. Juan, J. R. Montoya-Torres, C. L. Quintero-Araujo, and J. Panadero, “Simulation-optimization methods for designing and assessing resilient supply chain networks under uncertainty scenarios: A review,” *Simul. Model. Pract. Theory*, vol. 106, Art. no. 102166, 2021.
- [33] P. Toth and D. Vigo, Eds., *Vehicle Routing: Problems, Methods, and Applications*, 2nd ed. SIAM, 2014.
- [34] Tschan and A. Saxena, “From cold chains to code to communities: How digital innovation is making India’s HPV vaccine work for every girl,” United Nations Development Programme India, Mar. 16, 2026. [Online]. Available: <https://www.undp.org/india/blog/cold-chains-code-communities-how-digital-innovation-making-indias-hpv-vaccine-work-every-girl>
- [35] Y. Yang, H. Bidkhor, and J. Rajgopal, “Optimizing vaccine distribution networks in low

and middle-income countries,” *Omega*, vol. 99, Art. no. 102197, 2021.

- [36] Z. B. Zabinsky, P. Dulyakupt, S. Zangeneh-Khamooshi, *et al.*, “Optimal collection of medical specimens and delivery to central laboratory,” *Ann. Oper. Res.*, vol. 287, pp. 537–564, 2020.
- [37] Z. B. Zabinsky, M. Zameer, L. P. G. Petroianu, M. M. Muteia, and A. L. Coelho, “Route Optimization Tool (RoOT) for distribution of vaccines and health products,” *Gates Open Res.*, vol. 5, Art. no. 34, 2021, doi: 10.12688/gatesopenres.13219.1.