




Final Report-

Route planning for optimised last mile delivery

Primary Team ID - 45

Report Structure

VOLUME ESTIMATION


- Objective
- Methodology
- Pipeline
- Shape Handling
- Filters
- Results

ROUTE OPTIMISATION

- Introduction
- Modeling the problem
- Flexibility
- Dynamic Pickups
- Optimization Model Performance
- Scalability
- Bin Packing

DEV


- Software Development Life Cycle
- Tech Stack Used
- Database
- API Calls
- Screens



Volume Estimation

Objective

Our primary aim is to accurately calculate the volumetric weight of objects at the last-mile hub using the most reliable and efficient method available. To this end, we propose a **stereoscopic depth-based multi-modal volume approximation** approach that is based on shape-specific features, with filter enhancements and **zero manual intervention** in the pipeline.



Methodology

The Hardware Setup

We have prepared a proof-of-concept **conveyor belt setup** with the **camera mounting stand** to get the volume of objects in a warehouse setting like in last mile delivery hub. As shown in the diagram alongside, the setup is a **mock warehouse setup** with a platform arching over the top of the belt, containing the camera module and lights for volume estimation. Also, there is another camera on the side of the belt to see the lateral view of the object. The conveyor plate is supported by castor wheels, so the object can move in and out just like it would in a **scaled-up factory setup**.

The top panel and side panels are made of acrylic sheets which have been cut using laser cutter, using grooves to create a detachable setup. A acrylic board is used to move the object in and out of the field of view.

Acrylic Sheets



Object




Acrylic Tray (Conveyor)




Intel RealSense Camera

Flowchart

Main Pipeline




Shape-specific pipeline



Shape Handling

Object Recognition

We use a Deep Learning Model named **U2-Net for the Object Recognition purpose**. This enables us to separate the foreground object from the background, so that we can calculate the cross sectional area and also aid in shape classification.



Cross Sectional Area — Similarity Scaling And Pixel Enumeration

To calculate the cross-sectional area of the foreground object, we calculate the **Pixel Per Metric ratio** by first calculating the area of the actual surface visible in the camera. Then, we count the pixels in the foreground mask. Also since as the object moves closer to the camera, it occupies more and more area on the screen, while its actual area remains same. So, to correct for this, we employ the **concept of Similarity Scaling**. The following equations explain the algorithm well:-

$$\text{Similarity Factor} = \frac{(\text{Minimum Depth})}{(\text{Average Background Depth})}$$

$$\text{Cross Sectional Area} = (\text{Pixel Per Metric}) * (\text{Similarity Factor})^2 * (\text{Number of Pixels in image})$$

Shape Classification

We use a custom trained Deep Learning Model for Shape Classification on the foreground mask. The model was trained on a custom generated dataset of 2D shapes — rectangle, triangle and square. The model was able to attain near perfect accuracy on the validation dataset. The model architecture and dataset are attached below:




Object Height Calculation

From the Intel RealSense camera, we get two types of frames — color frame and depth frame. In this way, we can get the depth values of each pixel.


$$\text{Average Background Depth} = \frac{1}{n} \sum_{(i,j) \in \text{image}}^n \text{Depth}(i,j)$$

$$\text{Object Height} = \text{Average Background Depth} - \text{Minimum Depth}$$

Shape Handling



CUBOID




Cuboid is the simplest case in our example. The volume is simply the product of the cross-sectional area and the height that we obtained using the methods described above.

$$\text{Volume of cuboid} = \text{Cross-sectional Area} * \text{Height}$$



CONE




Cone happens to be the most challenging object to calculate the volume especially using one camera. To tackle this problem, we have installed a second camera which a RGB webcam to detect the lateral shape of the object in question. If the lateral shape is a triangle, then we can directly infer that the object is an upright cone. Now from the lateral view, we can get the height of the cone and from the top view, we can get the cross-sectional area of the base of the cone. Thus, we can get the radius of the base (just like in the case of sphere).

$$\text{Volume of cone} = \frac{1}{3} * \pi * r^2 * h$$

Shape Handling



CYLINDER




Cylinder is very similar to a cuboid except that it has a circular cross-sectional area. The method for finding the area of the circular cross-section is the same pixel enumeration approach as explained earlier. As for obtaining the height of the cylinder, we find out the difference between the average surface depth of the background and the depth of the highest point.

$$\text{Volume of cylinder} = \text{Circular cross-sectional area} * \text{Height}$$



SPHERE



Calculating the volume of a sphere has two challenges, the first being identifying if the object is a sphere or not and then calculating its radius. Now for the former, we use a mathematical comparison to classify the object into a sphere or otherwise.

The steps are:-

1. The difference in height of background and object is h .
2. Now from the circular cross section visible from the top view, we find out the area of the cross-section. From the obtained area, we calculate the estimated radius of the sphere by the formula, $r = \sqrt{\text{area}/\pi}$.
3. Now we compare the values $h/2$ and r and if the absolute difference between them falls inside a pre-defined threshold, then we declare that the object is a sphere.

$$\text{Volume of sphere} = \frac{4}{3}\pi r^3$$

Filters

We can drastically improve the performance of our camera output with the help of some post processing filters using deep learning methods. These methods are as follows :-




RGB Image


This is the RGB Image corresponding to the Depth images below.

No Filters

This is the depth map we obtain without using any filters to fill the holes that are present due to lack of exposure or noise.



Filters




Decimation Filter


Effectively reduces the complexity of the frame by reducing resolution. This could come in handy if our solution took too long to process each frame.

Spatial Filter

It helps improve the clarity and sharpness of object edges, making it easier to perceive depth. This filter operates in a linear-time computational manner, and its output is not dependent on the selection of arbitrary parameters, making it highly efficient and reliable.



Filters







Hole Filling Filter

It is the most essential filter, as it implements several methods to rectify missing data in the resulting image. The filter obtains the four immediate pixel "neighbors" (up, down, left, right), and selects one of them.

We finally use the Hole Filling filter as it offers visibly better hole filling as compared to Decimation, while giving a higher resolution than the Spatial Filter.

Results

Shape	Depth Image	RGB Image	Point Cloud	Volume (Stereoscopic Depth)	Volume (RGB to Point Cloud)	Ground Truth Volume
Cuboid				2097	2316	1953
Sphere				183	212	180
Cylinder				302	374	314
Cone				1578	1789	1526

As we can see from the table above, we have obtained significant improvement by using the current **stereoscopic depth method** instead of the RGB -to-point-cloud approach. This was made possible by the availability of the Intel RealSense Camera.

Advantages of Stereoscopic Multi-Modal Approach:

- This setup can be **scaled upwards** by **using a conveyor belt** instead of the acrylic tray that we use, and the same RealSense camera can be used in an industrial setting.
- The pipeline can handle moving objects and needs **no manual intervention** when scanning the objects.
- The pipeline uses very light Deep Learning models, and hence there are very less frame drops and the process is **real-time in nature**.
- Very accurate results are obtained as compared to pure Deep Learning based methods, and can compete against LiDAR based approaches **without even generating a point cloud**.

Optimization of Last Mile Deliveries

Introduction

The Capacitated Vehicle Routing Problem with Mixed Pickup and Delivery (CVRPMPD) is a optimization problem that involves finding the most effective routes while maintaining the various constraints such as capacity, route lengths and route times. It is an NP-Hard problem whose time complexity increases exponentially with increase in problem size and hence cannot be solved using exact methods in a reasonable amount of time.

To address this issue, the OR-Tools library is employed. OR-Tools is an open-source optimization library that provides tools which offers an efficient modelling framework to model the given problem while providing us with the flexibility for including contextual constraints and variables and gives results with great bounds in a reasonable time frame.

Modeling the problem

Step 1 : Problem Set-up

The decision variables in this problem are the routes that the vehicles take which minimized the total distance travelled while ensuring that all the customer demands are met without violating any constraints.

The distance matrix and the time matrix of the orders are calculated using OSRM and subsequently fed to the solver.

Step 2: Creating Constraints

Next, constraints must be defined that ensure the problem is solved correctly.

- Since both deliveries and pickups can be included in the same routes, we model this by giving **Positive** weights to Pickups and **Negative** weights to Deliveries. Then we set up constraints to ensure that the cumulative load at any customer location does not exceed the maximum capacity of a vehicle.
- The second set of constraints ensure that the time duration of a route does not exceed 5 hours. This is done by ensuring that the cumulative value of the time taken along a formulated route does not exceed the specified limit.

Action	Bag Fill Level	
Bag Load	0	
Warehouse	16	
Delivery	12 -4	
Pickup	12 +8	

- The third set of constraints ensure that the total route length of a route does not exceed the given maximum limit. This is done by ensuring that the cumulative value of the route length along a formulated route does not exceed the specified limit.
- The final constraint is the start and end point of the driver, which is initially the depot for both the start and end points.

Step 3: Adding Disjunction Penalties

A disjunction is a variable that the solver uses to decide whether to include a given location in the solution. If the solver decides not to include a order in today's dispatch, the the objective function is penalized. Higher the drop penalty, less are the chances of the node to be dropped and carry forwarded to the next day.


Additionally, instead of solely scheduling deliveries based on their expected delivery date, we consider the order density as a factor in the solution.

The main objective of the optimization problem would thus become minimizing the distance of the routes and at the same time minimizing the dropping penalties of orders which were not scheduled. The dropping penalty has two factors:

1. EDD

Including this factor in dropping penalty will enable us to prioritize orders with imminent delivery dates. Thus, an exponentially decreasing EDD penalty function will allow us to penalise the algorithm for missing out on deliveries scheduled immediately while, at the same time, allowing it to drop the orders scheduled a few days from now without major retribution. Mathematically,

$$P_{EDD} = 1000 * e^{-\frac{\text{Due In (days)}}{2}}$$



2. Location

An order in nearby areas with a high order density has a high priority value, as drivers delivering in nearby areas can piggyback this order, saving fuel and time and reducing trips. To include this factor, the city is divided into tiles based on Bangalore's ward data.

But, there might be cases where an order lies in a ward with a low order density but the neighbouring tiles have a very high order density. To overcome this problem each tile's priority is calculated by taking into consideration the priority of its neighbours' as well.

Suppose we have n tiles. We define the priority vector $\mathbf{v}_{\text{initial}} (n \times 1)$ where

$v_{\text{initial},i}$ = fraction of orders in tile i

Then, we set $\mathbf{A}$ as the n by n weighted adjacency matrix, considering tiles as nodes and weights as:

$$A_{ij} = \begin{cases} 1 & i = j \\ \frac{c}{\text{time}(i,j)} & i \neq j \end{cases}$$

Here, c is a constant derived experimentally, such that $\mathbf{A}_{ij} \leq 1$ and $\text{time}(i,j)$ returns the average time of traversing from tile i to tile j .


We then calculate the final priority vector of tiles, by :

$$\mathbf{v}_{\text{final}} = \mathbf{A} \cdot \mathbf{v}_{\text{initial}}$$


This will allow the order densities to propagate from tiles with significant order densities to nearby tiles thus allowing us to get a likeliness index of a tile being visited by a rider.

The location priority of an order: $\mathbf{p}$, thus can be defined as $\mathbf{v}_{\text{final},i}$ where i is the tile corresponding to the order location. Then we define the penalty for location as

$$P_{\text{Location}} = e^{7p}$$



Initial Priority



Final Priority

Combining the penalty values obtained

$$P_{\text{Dropping}} = P_{\text{EDD}} + P_{\text{Location}}$$

Example

Example	Case 1		Case 2	
	Order 1	Order 2	Order 1	Order 2
EDD - TODAY	2	3	2	3
Location Priority	0.4	0.6	0.2	0.8
Dropping Penalty	383	290	371	494

Step 4: Setting up the parameters

In OR-Tools, the parameters of the optimization model can be set to control the behavior of the solver and influence the solution obtained. The parameters include various elements such as the method used to find the initial solution, local search options (metaheuristics), the time limit for finding a solution and the limit to the number of solutions searched.

Step 5: Running the solver

The solver is then run to solve the problem. Upon completion the solver returns a list of nodes for each vehicle in the order in which they are supposed to visit them.

Flexibility

Balance Vehicle Utilization

Economically, it makes sense to minimize the number of vehicles used, but there are cases where timely delivery of the orders has to be prioritized in which case our solution has the ability to balance the vehicle loads among all the available vehicles to minimize the time of delivery.

A two step approach was employed to achieve this:

1. Minimize the maximum number of orders carried out by a vehicle.
2. Penalize the solver if the number of orders for a vehicle exceed

Time Windows

Incorporating time windows in a Capacitated Vehicle Routing Problem (CVRP) can add a temporal dimension to the routing problem, making it more realistic and allowing us to incorporate EDT (Estimated Delivery Time) for the orders. The solver is encouraged to prioritize making routes which enable the deliveries to be completed in the time window to avoid the penalty.

Incorporate Weather Conditions

Incorporating area-wise weather data into a Vehicle Routing Problem (VRP) can help make the routing solution more robust and efficient. The idea is to incorporate the hyper-local area effects of weather conditions on the routing distances, so that the VRP model considers the weather data when optimizing the routes. One way to do this is to increase the distance matrix cost for points that lie in areas with bad weather conditions. In this way, incorporating weather data into a VRP can help ensure that the routes are optimized for both distance and safety, taking into account the impact of weather on the road conditions.

Service Time

In addition to travelling time, we can include the time of servicing orders as well. Orders may take more time in certain locations such as societies due to time taken to do entry which may be included in the solver while forming routes which will maximise the percentage of successful deliveries in the predicted time.

Maximum Route Distance and Travel Time

While calculating the routes, the solutions allows us to restrict the maximum travel length of a particular route as well as the time of travel of the route. This is achieved by putting a constraint on the cumulative distance and cumulative travel time for routes to be less than the maximum limit.

Dynamic Pickups

Initial Condition

After the initial routes are created, the drivers set-off on their routes. As the driver goes on completing the orders, the status of the deliveries are constantly being updated in the solver.

When dynamic points are added and a rerouting request is created, we have the following information:

- Remaining orders
- Initial routes for each vehicle
- Current position of each vehicle
- Current capacity of each vehicle

Additional Constraints

- We add constraints on the start point of each driver as the current location and the end point being the depot.
- Another constraint that has to be applied is that since the rider already has the delivery orders in his bag, these orders cannot be assigned to different vehicles. Hence we add a constraint that delivery orders have to be assigned to the previously assigned vehicles.
- Since we already have the routes from the initial assignment, we provide the initial solution to the solver as the initial route minus the orders already delivered by the rider to improve performance.

Note: Pickup orders that have not been completed can now be assigned to a different driver to give more optimal results.

Volume Edge Cases

1. Imagine a scenario where the rider has attempted the delivery of an order and it failed, since the delivery can't be reattempted today, the space occupied by the order is technically 'blocked' until the orders are unloaded at the depot. Similarly, once a pickup order is completed, the space in the bag is 'blocked' until the orders are unloaded at the depot.
2. To overcome this challenge, two variables are maintained for each vehicle, namely *maximum_capacity* and *actual_capacity*.
3. To get the initial routes, we consider the *maximum_capacity* of the vehicles. But during rerouting, we directly consider the *actual_capacity* of the vehicle, hence simplifying the problem. Once the rider reaches the depot, the *actual_capacity* is reset to *maximum_capacity*.

Optimization Model Performance

Testing Against Benchmark Datasets

Benchmark Results

The CVRP solver was run on the various problem sets from the Uchoa et al. (2014) benchmark. The time taken to obtain a solution ranged from 4 seconds to 200 seconds with the optimality gap ranging from 4% to 12%.



Benchmark Instance	Nodes	Best Known	Result	Gap (%)	Time Taken(s)
X-n125-k30	125	55539.0	58198.4	4.78	4.45
X-n266-k58	266	75478.0	80286.0	6.37	30.25
X-n313-k71	313	94043	99075.7	5.35	66.05
X-n536-k96	536	94846	105773.5	11.52	59.78
X-n749-k98	749	77269	83716.2	8.34	102.81
X-n979-k58	979	118976	124390.9	4.55	197.98

The results are expected to approach the Best Known value if allowed to run for more time.

Result Plots



Various distributions were simulated and the routes were calculated as follows:

1. Uniform Distribution





Optimization Model Performance

2. Dense and Sparse Distribution



3. Clustered & Corner Depot Distribution



Scalability

Introduction

The CVRPMPD is an NP-hard problem, where even state-of-the-art solvers like OR tools struggle with large-scale datasets, around 5000 nodes or so. As the company scales up in cities, 5000+ orders are common, which will hamper scalability significantly. Clustering the nodes to represent virtual customers reduces the complexity by a large amount and allows the CVRP solvers to handle very large instances with low optimality gaps.

Advantages of Clustering

- Reduced complexity: Clustering allows for dividing a large problem into smaller, more manageable sub-problems, making it easier to solve the VRP.
- Improved scalability: As the size of the VRP increases, clustering can be used to efficiently handle larger problem instances by dividing the problem into smaller sub-problems.
- Flexibility: Clustering algorithms can be adapted and modified to incorporate specific constraints and requirements of the VRP, making it a highly flexible solution approach.

Proposed Method



Each order in a cluster can be represented at the centroid of the cluster, and the set of clusters can be used as a warm-up solution for solvers to route. Within the cluster, each cluster can further be solved using VRP.

We use a divide-and-conquer strategy to solve the routing problem and account for future scalability.

We have considered the sweep algorithm, a standard for vehicle routing problems, due to its robustness.

Sweep Algorithm

The Sweep Algorithm is a clustering method used for grouping geographical coordinates based on their cosine similarity to each other with respect to the depot. The algorithm works by sorting the orders according to the polar angles and iteratively sweeping a line in an angle and grouping orders together into clusters until the capacity is satisfied. The Sweep Algorithm is efficient and scalable, making it suitable for large datasets, and it can be easily adapted to incorporate specific constraints and requirements.



Suppose we have a total of 2000 orders and we have 80 riders. Let us cluster the orders into 4 clusters of 500 orders each with 20 riders assigned to every cluster.

```
coords[Depo] = [0, 0]
```

```
 $\theta_i$  (Polar Angle) of  $i^{th}$  order =  $\tan^{-1}(\text{lat}(i) / \text{long}(i))$ 
```

```
Sort customers according to  $\theta$  in ascending order
```

```
Set cluster_id = 1
```

```
Set  $\theta_c = 0^\circ$ 
```

```
Sweep customer by increasing  $\theta_c$ , add customer to cluster.
```

```
Stop the sweep when we reach 500 nodes in the cluster.
```

```
cluster_id = cluster_id + 1
```



```
Repeat above 3 steps till we have completed allotment of all customers
```

Potential Pitfalls

Choosing the starting angle $\theta_c = 0^\circ$ is not always efficient.

Consider the following example where if we start from 0° , we need to have 5 clusters instead of the optimal 4 clusters. Also, a clustering is considered good if each cluster has a small overall spread. Hence, to overcome this issue, we check the clustering for θ_c in increments of 5° .

The overall complexity of the algorithm will be $O(n \cdot \log n)$ where n is the number of orders.



To address this we can calculate clusters by varying the start point by 5° every time and calculating the number of clusters and spread of points we can evaluate how good the starting point is.

Bin Packing

Our algorithm uses a predefined order of box insertion in the container. Efficiency is improved by sorting based on volume, dimension or area of the largest face.

In our case priority of orders place a very important role so we are sorting the items based on priority of orders considering the one that should be delivered last should be placed on top.


Methods Used:

Extreme Point Insertion:

The algorithm starts by setting a default IP of $(0, 0, 0)$ and adding it to a set of available IPs. For each piece, the set is used to find possible insertion points in the container. The algorithm loops through all pieces in a set order and tries to insert them, rotating the piece in all allowed rotations. It checks if the piece can be loaded at each insertion point and calculates a score based on specific metrics after each packing. The process repeats until a better score is obtained or a timeout occurs.

Push Insertion:




The concept relies on normalizing by pushing pieces towards the container's origin along defined axes. The order of the axes can be set as a parameter and must include all axes. The algorithm loops through the given pieces in order, inserting each one without a set of available insertion points, but generating them on-the-fly.



Initially the piece is inserted at a corner of the container and the next step is normalizing it if the insertion was successful.

Greedy Adaptive Search Procedure:

The solution generated from above heuristics is optimized by trying different permutations of the order of blocks and orientation. Pieces are assigned a score based on the initial order, which is modified and re-sorted after each iteration. If no improvement is made in a set number of iterations, the score is reset to the best solution found so far.



Evaluation Metrics:

- Minimize residual volume of container after packing.
- Minimize the base area for the items to allow for easy pickups.
- Maximize the residual space's utilization of the resulting packing: The method aims to fit new pieces best by tracking residual space at each extreme point by projecting three axes until hitting another piece in the container.
- Minimize the Euclidean distance to the origin to increase the density of blocks.


Reference paper: <https://github.com/merschformann/sardine-can/tree/main/Material/MasterThesis>

Optimization of Last Mile Deliveries

Software Development Life Cycle

The industrial standards of the Software Development Life Cycle (SDLC) were followed to develop the prototype. It consists of a detailed plan describing developing, maintaining, replacing, and altering or enhancing specific software.


- Following were the stages of SDLC followed in the development of the product from the release of the Problem Statement (28th December 2022) -
- Stage 1 - Planning and Requirements Analysis (28th Dec - 31st Dec)
- Stage 2 - Defining Requirements (1st Jan - 6th Jan)
- Stage 3 - Designing the product. (6th Jan - 16th Jan)
- Stage 4 - Developing the product. (16th Jan - 2nd Feb)
- Stage 5 - Testing the product. (2nd Feb - 5th Feb)
- Stage 6 - Deployment and further testing. (5th Feb - 7th Feb) (7th Feb - Report Submission)




Tech Stack Used

Tech stack used in the development of the product was **React** in frontend and **Django** Rest framework in backend. Additionally, **Celery** was used to schedule the task of periodically running the optimization script. Rerouting for drivers was required after every 15 minutes, hence celery was used to schedule a task which puts the updated coordinates in the database.

When a task is ready to be run, Celery puts it on a queue, a list of tasks that are ready to be run. In order for the task to be executed, some other process, called a worker, has to be watching that queue for tasks. When it sees tasks on the queue, it'll pull off the first and execute it, then go back to wait for more.



ER Diagram of Database



To store data from the dashboard, ER diagram of the Database was created before the actual development of the product began. Order, Manager, Rider, and Bag have their tables in the database connected by the Trip table. Whenever the optimization algorithm runs and creates routes. a new trip is created for each route generated. The Trip Table contains foreign keys to the corresponding Rider, Order, and Bag records.

Tables in Database

Trip Table

It contains information about the order's life cycle and acts as a link between the Rider(one-to-one), Bag(many-to-one), and Orders(one-to-many) for delivery or pickup. It also contains start_time, end_time, created_time, and trip_status (unassigned, undelivered, delivered and failed).

Order Table

It contains trip_id(many-to-one relationship), which links it to the corresponding rider and bag. It has information on shape, volume, dimensions, delivery location (latitude, longitude, and text_address), owner details, and order SKU.

Rider Table

It contains details of rider(name, phone number, photo), manager_id(many-to-one), successful_deliveries, packages_delayed, current_trip_id(many-to-one), and earnings.

Manager Table

It contains manager details (name, phone number, photo) and the manager's address (latitude, longitude, and text address).

API Calls

rider-management/

GET - to fetch all riders' details and display them on the rider management page.

orders/all/

GET - to fetch details of all the orders along with rider details.

rider/<str:id>/

GET - to fetch details of a rider corresponding to the rider_id.

orders/rider/<str:id>/

GET - to fetch all the orders of the rider irrespective of their status.

orders/cancel/

POST - to cancel order corresponding to order_id.

orders/add/

POST - to dynamic pickup from the manager dashboard.

orders/update/

POST - to update the order corresponding to order_id

bin-packing/<str:id>/

GET - to fetch the details of bag packing for the corresponding rider's current trip.

geocode/


POST - to fetch geo-coordinates (latitude and longitude) corresponding to a text_address.

solve

POST - to solve the first instance of VRP.


Manager Dashboard

Volume Estimation Section



This part is the solution to Task 1 of the Problem Statement. Here, the Manager clicks on the start button to start the feed process. Then, the Depth Sensor Camera starts capturing the objects placed on the conveyor belt. Once the object is captured, the volume and dimensions of the object are displayed along with the RGB image and Depth Map Image of the object.


Inventory Section




This section provides details of the orders that are currently present in the depot. This comprised three parts: in-house, History, and Warehouse.

The warehouse comprises orders of type drop that are currently in cycle i.e. being delivered, or in warehouse for future deliveries.


Inhouse comprises of orders which have completed their cycle and will remain in the warehouse for a while, i.e. the orders where the type was dropped and failed, or type was a pickup and was successful.






History consists of orders dispatched from the depot which have completed their cycle.

Rider Management Section




This section helps the manager to access the details of riders assigned to them.

It provides information about the rider (current location, percentage of deliveries completed, basic details of rider, bag status, and delivery status of rider). It also allows managers to add dynamic pickup and cancel orders.



Rider Mobile App



The riders' screens were designed to accommodate the entire flow of deliveries. It starts with scanning the bag, then the rider is provided with the sequence of orders to be filled in the bag. The first order to be delivered is scanned, and the delivery status is set to out for delivery of that order. Once the delivery is reached, the rider is provided with two-step verification for confirming the delivery - Scan OR and OTP verification.



**Thank
You**