

INSTITUTO TECNOLÓGICO DE AERONÁUTICA



João Paulo Penna

**OPERATIONAL AND ECONOMIC VIABILITY OF A
DRONE-AS-A-SERVICE FOR FOOD DELIVERY: CASE
STUDY WITH REAL ORDER DATA FROM BRAZIL**

Final Paper
2025

Course of Aerospace Engineering

João Paulo Penna

**OPERATIONAL AND ECONOMIC VIABILITY OF A
DRONE-AS-A-SERVICE FOR FOOD DELIVERY: CASE
STUDY WITH REAL ORDER DATA FROM BRAZIL**

Advisor

Maj. Lucas Oliveira Barbacovi (ITA)

AEROSPACE ENGINEERING

SÃO JOSÉ DOS CAMPOS
INSTITUTO TECNOLÓGICO DE AERONÁUTICA

Cataloging-in Publication Data
Documentation and Information Division

Penna, João Paulo
Operational and Economic Viability of a Drone-as-a-Service for Food Delivery: Case Study with Real Order Data from Brazil / João Paulo Penna.
São José dos Campos, 2025.
133f.

Final paper (Undergraduation study) – Course of Aerospace Engineering– Instituto Tecnológico de Aeronáutica, 2025. Advisor: Maj. Lucas Oliveira Barbacovi.

1. Drone Delivery. 2. Last-Mile Logistics. 3. Economic viability. I. Instituto Tecnológico de Aeronáutica. II. Operational and Economic Viability of a Drone-as-a-Service for Food Delivery: Case Study with Real Order Data from Brazil.

BIBLIOGRAPHIC REFERENCE

PENNA, João Paulo. **Operational and Economic Viability of a Drone-as-a-Service for Food Delivery: Case Study with Real Order Data from Brazil**. 2025. 133f. Final paper (Undergraduation study) – Instituto Tecnológico de Aeronáutica, São José dos Campos.

CESSION OF RIGHTS

AUTHOR'S NAME: João Paulo Penna

PUBLICATION TITLE: Operational and Economic Viability of a Drone-as-a-Service for Food Delivery: Case Study with Real Order Data from Brazil.

PUBLICATION KIND/YEAR: Final paper (Undergraduation study) / 2025

It is granted to Instituto Tecnológico de Aeronáutica permission to reproduce copies of this final paper and to only loan or to sell copies for academic and scientific purposes. The author reserves other publication rights and no part of this final paper can be reproduced without the authorization of the author.

João Paulo Penna
Rua do H8A, 139
12.228-460 – São José dos Campos–SP

OPERATIONAL AND ECONOMIC VIABILITY OF A DRONE-AS-A-SERVICE FOR FOOD DELIVERY: CASE STUDY WITH REAL ORDER DATA FROM BRAZIL

This publication was accepted like Final Work of Undergraduation Study

João Paulo Penna

Author

Lucas Oliveira Barbacovi (ITA)

Advisor

Prof. Dr. Maisa de Oliveria Terra
Course Coordinator of Aerospace Engineering

São José dos Campos: June 8, 2025.

Acknowledgments

Esse trabalho concretiza o fim de um ciclo de sete anos, entre descoberta, estudos e graduação. Durante essa jornada construí uma clareza de que as motivações que me fizeram entrar no ITA se tornaram um tanto quanto diferente daquelas que me fizeram permanecer. E são a esses motivos que me inspiraram a quem dedico aqui meus agradecimentos:

Aos meus pais, por me ensinarem na prática e pelo exemplo o que realmente importa importa na vida.

Aos meus avós e em especial ao meu avô Hugo, engenheiro que se interessava a cada pequena conquista minha e que comemora conosco hoje lá de cima!

Aos meus amigos de BH, em especial ao Pedrinho com quem iniciei junto a jornada dos estudos para o ITA.

À Atlética por marcar de azul e amarelo algumas das melhores experiências nesses últimos 5 anos, em especial ao time de futsal, departamento de esportes de 2022 e diretorias e presidência de 2023 e 2024.

Aos meus queridos professores, em especial o Prof. Christopher, Profa. Máisa, Prof. Carlos Alberto, Prof. Leo, Prof. Márcio e a todos que com carinho e amor ajudam a manter a essência e a chama acesa.

E por fim, a todos os irmãos que compartilharam sua jornada comigo e que talvez sem eles não teria chegado até aqui, ou com certeza, teria sido menos divertido.

"Don't settle" - Steve Jobs

Abstract

Brazil's rapidly expanding food-delivery market, with over \$12,8 billion in 2024 and expected to grow to \$27,9 billion in 2030, with a annual growth rate of 22%(FORUM, 2024), exposes the drawbacks of a motorcycle-centric logistics model that intensifies congestion, emissions and accidents. Recent regulatory advances — notably amendments to RBAC-E 94 and Brazil's first BVLOS airworthiness certificate — now permit commercial drone operations, motivating a comprehensive assessment of a Drone-as-a-Service (DaaS) solution for urban last-mile delivery. This study draws on 1.05 million real orders from the Brendi platform and deploys an end-to-end pipeline: data geocoding and cleansing, computation of air (Haversine) and street distances, derivation of a Distance-Shortening Rate, placement of droneports through K-means and eigenvector centrality, fleet sizing with an M/M/c queue, and segment-level energy modelling that factors in wind and rainfall from ERA5 and INMET. These operational outputs feed a five-year discounted-cash-flow framework whose CapEx and OpEx explicitly incorporate Brazilian compliance costs.

Resumo

O crescimento acelerado do mercado brasileiro de entrega de refeições — estimado em \$12,8 bilhões de GMV, projetado para crescer até \$27,9 bilhões em 2030, com uma taxa anual de 22% (FORUM, 2024) — expõe as limitações do modelo logístico baseado em motocicletas, responsável por agravamento do trânsito, emissões de poluentes e elevada sinistralidade. A atualização do RBAC-E 94 e a recente emissão do primeiro CAER para voos BVLOS tornaram tecnicamente e legalmente viável a adoção de aeronaves remotamente pilotadas na última milha. Este trabalho propõe avaliar a viabilidade logística, energética, regulatória e econômico-financeira de um serviço Drone-as-a-Service (DaaS) aplicado ao food-delivery urbano. A pesquisa utilizará 1,05 milhão de pedidos reais da plataforma Brendi e seguirá um pipeline que compreende: geocodificação e higienização dos dados; cálculo de distâncias aérea (Haversine) e viária (Dijkstra); obtenção do índice de encurtamento de distância (DSR); localização de droneports por K-means combinado a centralidade eigenvector; dimensionamento de frota via fila M/M/c; e modelagem energética trecho-a-trecho com ajustes para vento e precipitação derivados de séries ERA5 e INMET. Esses módulos alimentarão um modelo de fluxo de caixa descontado a cinco anos, no qual CapEx e OpEx incluem custos regulatórios nacionais.

List of Figures

FIGURE 1.1 – Evolution of the global <i>online food delivery</i> market (2020–2030). Grand View Research’s estimate projects growth from US\$ 12.8 billion in 2024 to US\$ 27.9 billion in 2030, corresponding to a CAGR of 22% in the 2025–2030 period(FORUM, 2024).	19
FIGURE 2.1 – Operational modes for meal delivery. (a) Ground Delivery : the entire route is made on the road network. (b) Pure Drone Delivery : the package flies directly from the restaurant to the customer’s home and returns. (c) Mixed – last-mile by drone : a ground vehicle takes orders to a <i>depot</i> , from where drones cover the last mile. (d) Mixed – last-mile by motorcycle : the drone acts as an ”air bridge” between restaurant and <i>droneport</i> ; motorcycles complete the final route for addresses outside the battery radius. Red segments indicate package movement, while the blue band symbolizes the BVLOS air corridor.	23
FIGURE 2.2 – Example of mixed motorcycle–drone routing with time windows, adapted from Han <i>et al.</i> (HAN <i>et al.</i> , 2023). Black customers are served by the motorcycle (solid arrows), while red customers receive from the drone (dashed arrows). The depot (<i>node 0</i>) functions as the aerial launch and recovery point. The image visually illustrates the logic that the VRP-D formulation seeks to optimize: reduce ground distance and meet time windows by correctly allocating each customer to the most efficient mode.	24
FIGURE 2.3 – Intuitive view of the K-means algorithm: <i>before</i> (<i>left</i>) points are distributed without labels; <i>after</i> (<i>right</i>) the method partitions the space into k groups minimizing internal variance (LLOYD, 1982; ARTHUR; VASSILVITSKII, 2007).	29

- FIGURE 2.4 – Three-step procedure to locate *droneports*, according to Bine *et al.* (BINE *et al.*, 2023). **(a)** Aggregated vehicle flow (0–24 h); lighter roads concentrate greater traffic. **(b)** K-means result with $k = 3$, revealing relatively homogeneous zones. **(c)** Final position of droneports after applying the flow filter and selecting the node with highest eigenvector centrality in each cluster. 30
- FIGURE 4.1 – Geographic area selected for analysis, delimited by latitude and longitude coordinates. The region encompasses the main urban area of São José dos Campos and adjacent parts, including municipalities such as Jacareí, where there is significant presence of Brendi restaurants. 56
- FIGURE 4.2 – Boxplot of delivery distance distribution before filter application. The concentration of data near zero and the presence of extreme outliers evidence the need for statistical validation. 57
- FIGURE 4.3 – Comparison of distance distributions before and after outlier filter application. The left histogram shows the original distribution with extreme values up to 400 km, while the right one presents the filtered distribution concentrated in distances up to 14 km, more representative of the urban context. 58
- FIGURE 4.4 – Elbow method applied to Approach A for determining the optimal number of clusters. Analysis of inertia $J(k)$ as a function of the number of clusters k reveals an inflection point (elbow) at $K = 4$, indicating that adding more clusters beyond this point offers diminishing returns in inertia reduction. 60
- FIGURE 4.5 – Order coverage as a function of the number of *droneports* (Approach A). The curve shows diminishing returns as the number of *droneports* increases, with significant marginal gains until $K=4$, after which adding more infrastructure offers progressively smaller improvements in coverage. The dashed horizontal line indicates 100% coverage as reference. 61
- FIGURE 4.6 – Final location of the 4 *droneports* determined by the K-means algorithm in the São José dos Campos region. The centroids (marked with red 'X') represent optimal positions of the *droneports*, while colored points represent the distribution of order demand grouped by cluster. The map evidences the strategic distribution of *droneports* to maximize demand coverage within the 5 km radius. 62

- FIGURE 4.7 – Location of the 4 *droneports* (marked with red 'X') and the main centroid (marked with blue star) used as reference point for meteorological data collection in the São José dos Campos region. 63
- FIGURE 4.8 – Meteorological data (wind speed and direction, and precipitation) for the first 7 days of analysis (June 2024). Data illustrate favorable conditions for operation, with wind speeds consistently below the 25 km/h limit and absence of significant precipitation in the period. 64
- FIGURE 4.9 – Flight viability matrix (flyable) for the conservative scenario, showing the first 30 days of analysis. Each cell represents a specific hour of a day, with dark green indicating viable flight (1) and red indicating restriction (0). The conservative scenario presents restriction periods concentrated mainly at the end of June, especially in afternoon hours. 65
- FIGURE 4.10 – Flight viability matrix (flyable) for the moderate scenario, showing the first 30 days of analysis. Compared to the conservative scenario, it presents almost continuous availability, with only some isolated restriction periods. 65
- FIGURE 4.11 – Flight viability matrix (flyable) for the aggressive scenario, showing the first 30 days of analysis. It presents almost total availability, with only a minimal restriction period at the beginning of the first day. 66
- FIGURE 4.12 – Flight restriction reasons by meteorological scenario, detailing contribution of wind, precipitation, or both. The conservative scenario presents the highest number of restrictions, predominantly due to wind speeds above the limit, while moderate and aggressive scenarios present reduced number of restrictions. 66
- FIGURE 4.13 – Influence zones of *droneports* in the São José dos Campos region. Markers with 'X' represent the 4 *droneports* (DP1 in red, DP2 in blue, DP3 in green, DP4 in yellow), while colored squares represent restaurants associated with each *droneport* according to shortest road distance calculated by Google Routes API. Spatial distribution evidences natural partition of the region into coverage zones based on real road proximity, allowing logistical optimization of terrestrial transport to takeoff points. 69

- FIGURE 4.14 – Conceptual triangular diagram of delivery modes. The diagram illustrates the two possible operational strategies for an order: (1) direct mode, where the motoboy transports the order directly from restaurant to customer, covering distance $d_{\text{rest-consumer}}$; (2) mixed mode, where the motoboy transports the order from restaurant to *droneport* (distance $d_{\text{rest-droneport}}$), and then the drone completes delivery from *droneport* to customer (distance $d_{\text{droneport-consumer}}$). The decision between modes depends on comparison of total times of each strategy, considering meteorological and operational range restrictions. 71
- FIGURE 4.15 – Comparative analysis of delivery modes: (a) distribution of chosen modes, showing predominance of motorcycle mode (88,78%); (b) distribution by decision reason, evidencing that shorter motorcycle time is the main factor; (c) comparison of delivery time distributions through boxplots, indicating that motorcycle mode presents shorter median times; (d) time savings distribution, showing that when mixed mode is chosen, there is positive savings relative to motorcycle mode. 75
- FIGURE 4.16 – Detailed analysis of orders with mixed delivery: (a) time comparison through scatter plot, where each point represents an order and color indicates time savings obtained; points below equality line indicate positive savings of mixed mode; (b) percentage time savings distribution, showing that most orders present savings between 5% and 15%, with mean of 11,41% and median of 8,24%. Distribution is right-skewed, indicating that some orders present significantly larger percentage savings. 76
- FIGURE 4.17 – Percentage distribution of delivery modes by distance range between *droneport* and customer. The stacked bar chart shows that mixed mode (in turquoise blue) presents highest proportion in the 2–4 km range (16,67%), decreasing for smaller distances (9,22% in 0–2 km) and larger distances (11,86% in 4–6 km). For distances superior to 6 km, mixed mode is not used due to operational range restriction, with 100% of orders served by motorcycle. 77

- FIGURE 4.18 –Geographic distribution of orders by delivery mode in the São José dos Campos region: (a) general view showing all orders (red = motorcycle, cyan = mixed, yellow stars = *dronports*); (b) areas served exclusively by mixed delivery (2,618 orders), evidencing concentrated clusters near *dronports*; (c) areas served by motorcycle (20,717 orders), showing much broader coverage; (d) overlaid comparison with transparency indicating density, evidencing that motorcycle has dominant spatial distribution while mixed mode operates in specific zones. 78
- FIGURE 4.19 –Order volume distribution among the 20 restaurants with highest demand. The chart evidences high concentration of orders in top restaurants: the leading restaurant concentrates 6,205 orders, while following ones show significantly smaller volumes. The top 6 restaurants concentrate a substantial proportion of total demand, justifying *dronport* positioning strategies based on demand instead of uniform spatial distribution. 80
- FIGURE 4.20 –Geographic distribution of restaurants coded by order volume. Size and color of circles indicate order volume of each restaurant, varying from light yellow (low volume) to dark red (high volume). *Dronports* positioned via K-means are marked with cyan stars. Visualization evidences that highest-volume restaurants (large red circles) do not necessarily coincide with calculated centroids, demonstrating that strategies based on spatial mean may not adequately capture real demand concentration. 81
- FIGURE 4.21 –Geographic distribution of orders by delivery mode for configuration with 4 *dronports* positioned at top 4 restaurants with highest order volume: (a) areas served by mixed delivery (9,351 orders), showing concentration near *dronports*; (b) areas served by motorcycle (13,984 orders), with broader distribution; (c) overlaid geographic comparison, evidencing that mixed mode operates mainly within *dronport* coverage radius; (d) order distribution by *dronport*, showing that each *dronport* serves both orders from the restaurant where it is located and from nearby restaurants in the region. This configuration presents 40,07% mixed mode penetration and will be used as the basis for subsequent analyses. 83

- FIGURE 4.22 –Temporal distribution of mixed orders: (a) distribution by hour of day, evidencing demand concentration in evening period (17h–23h) with peak at 20h; (b) distribution by day of week, showing progressive increase throughout the week with larger volumes on weekends. Temporal analysis is fundamental to identify peak periods that require greater capacity and system idleness periods. 85
- FIGURE 4.23 –Heatmap of mixed order distribution by day of week and hour of day. Color intensity indicates order volume, varying from light yellow (low volume) to dark red (high volume). The heatmap clearly evidences that periods of highest demand concentrate in hours 18h–22h, especially on weekends (Friday, Saturday, and Sunday). Early morning and morning periods show very low demand, indicating potential significant system idleness during these hours. 86
- FIGURE 4.24 –M/M/c sizing analysis by *droneport*: (a) minimum number of drones necessary to meet 20-minute SLA, showing that Droneports 1 and 2 require 4 drones each, while Droneports 3 and 4 require 2 drones each; (b) utilization rate (ρ) by *droneport*, evidencing that all operate below 100% (instability limit), with Droneports 1 and 2 showing higher utilization; (c) mean total time in system (W) compared to SLA, confirming that all *droneports* meet SLA; (d) relationship between arrival rate (λ) and total capacity ($c \cdot \mu$), showing that capacity exceeds demand in all cases. 88
- FIGURE 4.25 –SLA exceedance analysis by *droneport*: (a) probability of exceeding 20-minute SLA, showing that Droneport 1 presents highest probability (5,73%), while Droneport 4 presents lowest probability (1,80%); (b) expected number of orders that would exceed SLA, evidencing that approximately 415 orders (4,4% of total) could exceed SLA; (c) mean time in system vs SLA, confirming that all meet SLA on average; (d) relationship between mean time and exceedance probability, showing positive correlation between these metrics. 89
- FIGURE 4.26 –Distribution of headwind components in cruise segments: (a) outbound distribution, showing slight tendency for tailwind (negative values); (b) return distribution, showing slight tendency for headwind (positive values). The dashed red line indicates no-wind condition ($v_w = 0$). 92

- FIGURE 4.27 –Relationship between headwind component and energy consumption in cruise segments: (a) outbound cruise, showing positive correlation (0,346) between headwind and energy; (b) return cruise, showing similar correlation (0,361). Consumption in return cruise is slightly smaller than outbound due to lower mass (empty drone). 93
- FIGURE 4.28 –Energy distribution by flight segment: (a) bar chart with mean contribution in Wh; (b) pie chart with percentage distribution. Cruise segments (outbound and return) are responsible for approximately 77% of total consumption. 93
- FIGURE 4.29 –Distribution of energy consumed per order. The dashed green line indicates mean (157.2 Wh) and the dashed red line indicates usable battery limit (360 Wh). All orders are within safety limit. 94
- FIGURE 4.30 –Relationship between energy consumed and flight distance. The dashed red line indicates usable battery limit (360 Wh). Strong positive correlation between distance and energy is observed, with all orders within safety limit. 95
- FIGURE 4.31 –CapEx distribution by component: (a) horizontal bar chart with absolute values in thousand USD; (b) pie chart with percentage distribution. Drone acquisition represents approximately 57% of total investment. 100
- FIGURE 4.32 –CapEx by droneport with breakdown by component. Droneports 1 and 2 show higher investment due to greater number of drones (4 drones each), while Droneports 3 and 4 operate with 2 drones each. Shared costs (regulatory, software) were allocated proportionally. . . 100
- FIGURE 4.33 –OpEx distribution by component: (a) horizontal bar chart with absolute values in thousand USD; (b) pie chart with percentage distribution. Salaries and operation represent approximately 44% of total operational cost. 103
- FIGURE 4.34 –Demand distribution by hour of day in the analyzed period. Concentration of orders in evening hours (18h–23h) evidences significant infrastructure idleness during daytime period, especially between 5h and 17h. Evening period (21h–4h) is highlighted in light purple. . . 105

- FIGURE 4.35 –Demand distribution before and after mirroring (17h–22h → 10h–15h). Blue bars represent original demand, while red bars represent demand after mirroring. Green arrows visually indicate mapping between evening hours and lunch hours. Shaded areas highlight periods involved in mirroring: light purple for source evening period (17h–22h) and light red for target lunch period (10h–15h). 108
- FIGURE 4.36 –OpEx per order comparison before and after demand mirroring. Reduction of USD 2,39 per order (from USD 6,86 to USD 4,45) represents significant improvement in operational efficiency through better utilization of existing infrastructure. 109
- FIGURE 4.37 –Market share expansion modeling: infrastructure and cost dilution analysis. Charts show evolution of (a) total drones and *droneports*, (b) necessary batteries, (c) fixed, variable, and total CapEx, (d) fixed, variable, and total OpEx, (e) CapEx dilution per order, (f) OpEx dilution per order, and (g) mean and maximum system utilization. Dashed vertical lines indicate current market share (0,49%, blue) and target market share (20%, green). 117
- FIGURE 4.38 –Financial and operational evolution over 5 years: (a) market share evolution, (b) demand growth, (c) annual revenue evolution, (d) OpEx and incremental CapEx, (e) annual free cash flow, (f) accumulated free cash flow (with break-even line), (g) operating margin, and (h) infrastructure evolution (drones and *droneports*). Dashed line in chart (f) indicates break-even point (Year 2). 122

Contents

1	INTRODUCTION	18
1.1	Contextualization	18
1.2	Regulatory and technological motivation	19
1.3	Research problem	21
1.4	Hypothesis	21
1.5	Objectives	21
1.6	Work structure	22
2	LITERATURE REVIEW AND FUNDAMENTALS	23
2.1	Logistical Fundamentals	23
2.1.1	Vehicle Routing Problem with Drones (VRP-D)	24
2.1.2	Distance-Shortening Rate (DSR)	27
2.1.3	K-means clustering and centrality	28
2.1.4	M/M/c queue — fundamentals and sizing	30
2.2	Drone energy modeling	32
2.2.1	Power in stationary flight (hover)	33
2.2.2	Power in translation	33
2.2.3	Energy per order	33
2.3	Economic–financial evaluation	34
2.4	Brazilian regulatory framework	36
2.5	Related work and positioning	38
3	MATERIALS AND METHODS	41
3.1	Data and inputs	42

3.1.1	Brendi order dataset	42
3.1.2	Google Routes API	43
3.1.3	ERA5 and INMET climate series	44
3.1.4	Cost structure	44
3.1.5	Drone specification	45
3.2	City and analysis period selection	45
3.3	Operational pipeline	46
3.4	Economic-financial model	48
3.5	Tools and reproducibility	50
4	DISCUSSIONS	52
4.1	Discussion of Logistical Results	52
4.1.1	Preprocessing and Data Cleaning	52
4.2	Discussion of Economic-Financial Results	97
4.2.1	CapEx Analysis (Capital Expenditure)	97
4.2.2	OpEx Analysis (Operational Expenditure)	101
4.3	Model Validation and Limitations	110
4.3.1	Validation with Reference Model (iFood Delivery Plan)	110
4.3.2	Horizontal Scale Analysis	114
4.3.3	Temporal Growth Modeling: Financial Evolution over 5 Years	118
5	CONCLUSION	124
5.1	Synthesis of Main Results	124
5.2	Response to Hypothesis and Research Problem	125
5.3	Contributions of the Work	126
5.4	Limitations and Future Work	128
	BIBLIOGRAPHY	131

1 Introduction

The introduction aims to provide the reader with a clear and self-contained overview of the problem that will be investigated throughout this work. To this end, the chapter is structured in five moments. First, a *contextualization* of the food delivery market and the negative impacts observed in the traditional logistics model based on motorcycles is presented. Next, the *regulatory and technological motivation* that makes plausible — and unprecedented in Brazil — the commercial use of remotely piloted aircraft in last-mile routes is discussed. With this background established, the *research problem* that guides the study is formulated, the *hypothesis* that will be tested is presented and, finally, the specific *objectives* that will be pursued to answer it are made explicit. The logical sequence of these sections introduces, gradually, the economic, social and academic relevance of the theme, while delimiting the scope and evaluation criteria adopted in subsequent stages.

1.1 Contextualization

In the last two decades, *food delivery* has ceased to be a niche to become one of the most dynamic segments of digital commerce. Estimates from the *World Economic Forum* indicate that the Brazilian market moved approximately US\$ 12.8 billion in 2024 and should grow to US\$ 27.9 billion in 2030, with an annual rate of 22 % (FORUM, 2024). This accelerated growth, however, exacerbates limitations of the predominant logistics model — motorcycle delivery — which was conceived for much smaller volumes.

The massive presence of delivery motorcycles aggravates urban congestion: a survey by *CurbFlow* points out that motorcycle deliveries are responsible for more than 20 % of double-parking stops in downtown São Paulo (CURBFLOW, 2023). The social impact is also reflected in accident rates: only between January and September 2024, 1,925 motorcyclists lost their lives on São Paulo roads (SP, 2024). From an environmental perspective, dynamometer tests show that motorcycles can emit up to ten times more carbon monoxide and hydrocarbons than standard Euro 6 light vehicles (SILVA *et al.*, 2013). These factors evidence a saturation point in the current model and justify the search for safer and more sustainable alternatives.

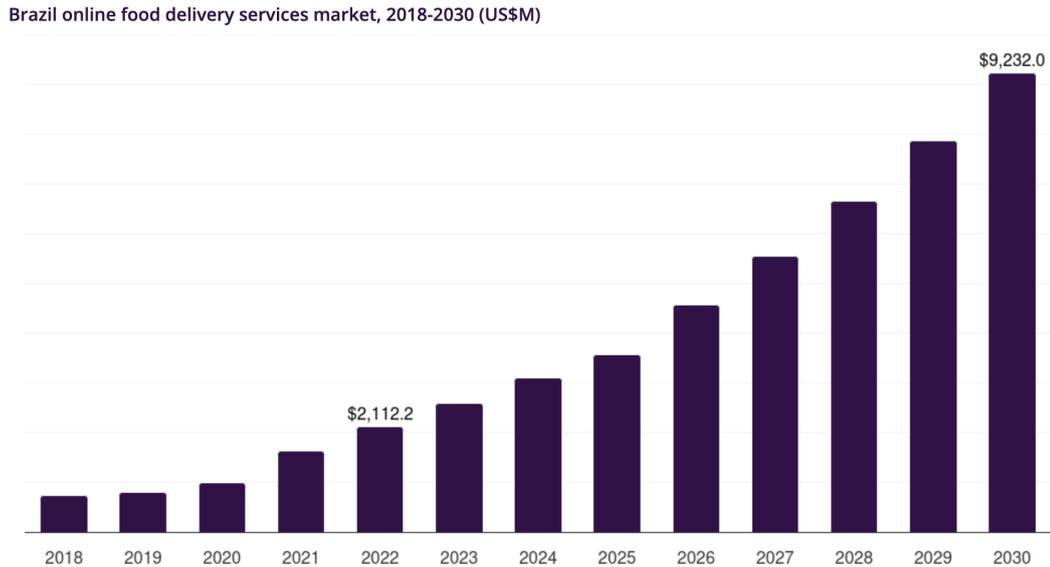


FIGURE 1.1 – Evolution of the global *online food delivery* market (2020–2030). Grand View Research’s estimate projects growth from US\$ 12.8 billion in 2024 to US\$ 27.9 billion in 2030, corresponding to a CAGR of 22 % in the 2025–2030 period(FORUM, 2024).

1.2 Regulatory and technological motivation

Brazil has, since 2017, a regulatory framework that allows commercial operations with drones. The *Brazilian Civil Aviation Regulation* RBAC-E 94, updated in 2022, and instruction ICA 100-40 from DECEA authorize *Beyond Visual Line-of-Sight* (BVLOS) flights with aircraft up to 25 kg, provided they are equipped with redundant safety systems and properly certified. The practical legitimacy of this framework was confirmed in 2022, when Speedbird Aero received the first *Special Airworthiness Certificate* (CAER) to perform deliveries in partnership with iFood(CIVIL, 2023a).

Furthermore, practical results obtained by commercial aerial delivery operators reinforce the relevance of this study’s proposal. Three experiences, in particular, serve as benchmarks to calibrate adoption parameters, operational radius and revenue model: Speedbird Aero in Brazil, Wing in the United States and Meituan in China.

Speedbird Aero & iFood (Brazil)

Speedbird Aero was the first Latin American company to receive from ANAC the *Special Airworthiness Certificate* (CAER) for BVLOS operations with multirotor drones up to 15 kg(CIVIL, 2023a). In partnership with iFood, it began regular flights in Aracaju and Campinas, using the DLV-1 model to cover segments of up to 3 km in approximately eight minutes. The process adopts a hybrid format: the drone takes off from a shopping center, lands at a *droneport* in a residential neighborhood and a ground delivery person

completes the last 400–600 m. By the end of 2024 approximately 30 000 aerial deliveries had been recorded, with an average reduction of 60 % in total time compared to motorcycle delivery.¹ Flights are charged in the *pay-per-flight* model (fixed fee passed to the restaurant), which approximates the DaaS concept evaluated in this work.

Wing & Walmart (United States)

Wing, an Alphabet subsidiary, operates last-mile services in Dallas–Fort Worth and Phoenix. In June 2025 the company announced expansion of the partnership with Walmart to one hundred stores, with capacity of ~1 000 daily deliveries and service radius of 10 km (VINCENT, 2025). Orders are dispatched from small modular stations installed in store parking lots; the flight is fully autonomous, from takeoff to landing on retractable cable at the customer’s residence. For *Walmart+* program members, aerial delivery was offered at no additional cost, a strategy that increased adoption rate by 38 % in the first month of operation. By April 2025 Wing reported more than 350 thousand global deliveries, which demonstrates the technical maturity and economic scalability of the system.

Meituan (China)

Since 2021 Meituan has maintained ten fixed aerial routes in the city of Shenzhen, approved by the Civil Aviation Administration of China for corridors of up to 6 km in urban areas. In 2022 the company surpassed the milestone of 100,000 delivered orders (GROUP, 2022), using multirotors with 8 kg maximum takeoff weight and standardized boxes of 2 kg payload. The end customer pays a premium of approximately ¥5 (\approx R\$ 3.50) over the conventional delivery fee and receives the order on average 15 minutes after confirmation. Internal reports indicate operational cost 27 % lower than motorcycle delivery, due to higher route density and elimination of traffic risk.

Together, these three examples demonstrate that drone delivery has already surpassed the technological demonstration phase and begins to consolidate varied business models — hybrid (Brazil), fully autonomous with subsidy (USA) and fixed tariffed corridor (China). They provide performance metrics that will be used to validate the simulation developed in this work, especially regarding total delivery time, operational radius and cost structure per flight.

¹Data disclosed in iFood press conference, Oct. 2024.

1.3 Research problem

Given this context, the question is whether it would be logistically and economically viable to implement, in a Brazilian city, a *Drone-as-a-Service* (DaaS) service dedicated to meal delivery, taking as a pilot phase the order portfolio already intermediated by the Brendi platform.

1.4 Hypothesis

It is hypothesized that a DaaS service dedicated to meal delivery is logistically viable — capable of meeting a minimum Service Level Agreement (SLA) of 20 minutes and reducing total delivery time compared to the traditional motorcycle-based model — and economically viable — with operational costs comparable to or lower than motorcycle delivery, considering the necessary infrastructure, regulatory costs and the operation scale of the Brendi platform.

1.5 Objectives

This work intends to **simulate** a drone delivery operation in a chosen city, using real orders as a basis for modeling. Based on this simulation, it seeks to **size** the necessary infrastructure — including the drone fleet size and location of support points (*droneports*) — capable of meeting a minimum Service Level Agreement (SLA) of 20 minutes, ensuring that delivery time is lower than the traditional motorcycle-based model. From this sizing, it is proposed to identify which deliveries could be migrated from the exclusive motorcycle model to include the drone as a participating modality.

With the simulated and sized operation, the work aims to **calculate capital costs** (Capex) and **operational costs** (Opex) necessary to implement this operation, covering aircraft acquisition, *droneport* installation, safety equipment, regulatory certifications, investments in management systems, maintenance, insurance, energy costs, remote pilot compensation and administrative expenses, already incorporating the regulatory costs provided for in RBAC-E94 and ICA 100-40. Finally, these costs are **consolidated** in a financial structure that allows evaluating the economic viability of implementing this model in the chosen city, comparing them with the costs of the traditional motorcycle-based model and providing a quantitative basis for decision-making.

1.6 Work structure

The text is organized in five chapters. Chapter 1 presents the contextualization, regulatory and technological motivation, research problem, hypothesis and study objectives. Chapter 2 compiles the literature review on last-mile logistics with drones, energy modeling, queueing theory, financial evaluation and the national regulatory framework. Chapter 3 describes in detail the data used, the sample selection criteria, the operational simulation pipeline and the economic-financial model. Logistical and profitability results are discussed in Chapter 4. Finally, Chapter 5 presents the work's conclusion, synthesizing the main results, evaluating the research hypothesis, highlighting contributions and pointing out limitations and future work.

2 Literature Review and Fundamentals

This chapter establishes, in detail, the theoretical concepts necessary for the logistical, energetic and economic-financial modeling of the Drone-as-a-Service service studied in this work. Whenever possible, complete formulas, underlying assumptions and limitations of each model are presented, so that the reader can reproduce them without resorting to other sources.

2.1 Logistical Fundamentals

The first step in modeling deliveries with drones is to understand *how* these aircraft can fit into the existing distribution chain. Bine et al. (BINE *et al.*, 2023) summarize the possibilities in four configurations, shown in Figure 2.1. They range from *Ground Delivery*, which relies 100% on motorcycle routes, to hybrid scenarios where drones execute only part of the route.

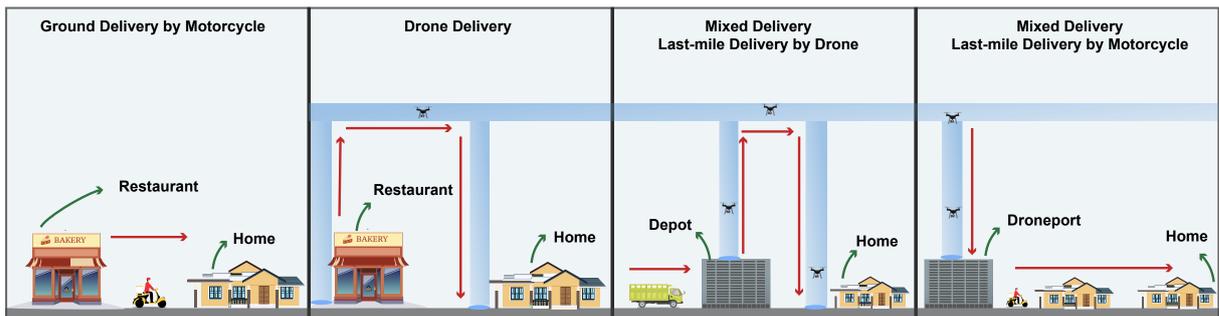


FIGURE 2.1 – Operational modes for meal delivery. (a) **Ground Delivery**: the entire route is made on the road network. (b) **Pure Drone Delivery**: the package flies directly from the restaurant to the customer’s home and returns. (c) **Mixed – last-mile by drone**: a ground vehicle takes orders to a *depot*, from where drones cover the last mile. (d) **Mixed – last-mile by motorcycle**: the drone acts as an “air bridge” between restaurant and *droneport*; motorcycles complete the final route for addresses outside the battery radius. Red segments indicate package movement, while the blue band symbolizes the BVLOS air corridor.

Figure 2.1 illustrates that the choice of transport mode is not binary; it is a spectrum that depends on customer location, drone autonomy and relative cost between air and ground modes. To automatically decide where to employ each option, we resort to the **Vehicle Routing Problem with Drones (VRP-D)**, which extends the classic VRP by

allowing aerial launches along the route. The next subsection presents the mathematical formulation of VRP-D, its variables, constraints and solution heuristics.

2.1.1 Vehicle Routing Problem with Drones (VRP-D)

The classic *Vehicle Routing Problem* (VRP) seeks to determine the set of lowest-cost routes that depart from a depot, visit each customer exactly once and return to the depot. In the *Vehicle Routing Problem with Drones* (VRP-D), a ground vehicle (in this study, a motorcycle) is combined with one or more drones, so that some customers can be served by air, shortening total time or distance. The formulation below follows Li et al. (LI *et al.*, 2021) and includes time windows (*Time-Window VRP-D*).

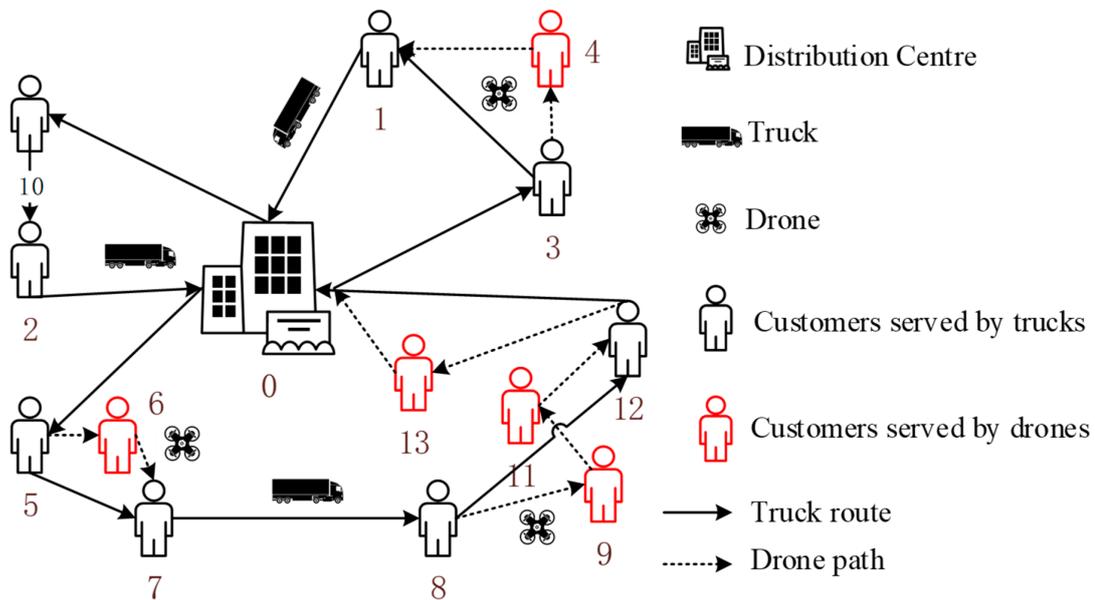


FIGURE 2.2 – Example of mixed motorcycle–drone routing with time windows, adapted from Han *et al.* (HAN *et al.*, 2023). Black customers are served by the motorcycle (solid arrows), while red customers receive from the drone (dashed arrows). The depot (*node 0*) functions as the aerial launch and recovery point. The image visually illustrates the logic that the VRP-D formulation seeks to optimize: reduce ground distance and meet time windows by correctly allocating each customer to the most efficient mode.

Sets

- $\mathcal{N} = \{1, \dots, n\}$ customers;
- 0 depot (origin and return of motorcycle);
- $\mathcal{V} = \mathcal{N} \cup \{0\}$ set of all nodes.

Parameters

c_{ij}^T cost (or time) of the motorcycle on edge (i, j) ;

c_{ij}^D cost (or time) of the drone from point i to j ;

t_{ij}^T, t_{ij}^D
ground and air travel times;

a_i, b_i start and end of customer i 's time window;

s_i service time at customer i ;

B maximum autonomy (outbound + return) of the drone;

M large constant to linearize implications.

Decision variables

x_{ij} binary variable (1 if motorcycle travels (i, j));

y_{ij} binary variable (1 if drone departs from i to serve j);

u_i instant when motorcycle service starts at node i .

Mathematical model Objective function

$$\min \underbrace{\sum_{i,j \in \mathcal{V}} c_{ij}^T x_{ij}}_{\text{ground cost}} + \underbrace{\sum_{i,j \in \mathcal{N}} c_{ij}^D y_{ij}}_{\text{air cost}} + \underbrace{\sum_{i \in \mathcal{N}} \alpha (u_i - a_i)^+}_{\text{delay penalty}}. \quad (2.1)$$

- *First term*: distance (or time) traveled by the motorcycle.
- *Second term*: accumulated cost of drone flights.
- *Third term*: penalty for each minute that service starts after the beginning of window a_i ; $(\cdot)^+$ denotes *maximum between zero and the delay*.

Constraints

$$\sum_{j \in \mathcal{V}} x_{ij} = 1, \quad \sum_{j \in \mathcal{V}} x_{ji} = 1 \quad \forall i \in \mathcal{N} \quad (2.2a)$$

$$\sum_{j \in \mathcal{N}} y_{ij} \leq 1 \quad \forall i \in \mathcal{V} \quad (2.2b)$$

$$u_j \geq u_i + s_i + t_{ij}^T - M(1 - x_{ij}) \quad \forall (i, j) \in \mathcal{V}^2 \quad (2.2c)$$

$$u_k \geq u_i + s_i + t_{ij}^D + t_{jk}^D - M(1 - y_{ij}) \quad \forall (i, j, k) \in \mathcal{V}^3 \quad (2.2d)$$

$$t_{ij}^D + t_{jk}^D \leq B \quad \forall (i, j, k) \quad (2.2e)$$

$$a_i \leq u_i \leq b_i \quad \forall i \in \mathcal{N} \quad (2.2f)$$

(2.2a)

Each customer has exactly one arrival and one departure from the motorcycle.

(2.2b)

From that node, the truck can launch at most one flight.

(2.2c)

If (i, j) is on the route ($x_{ij} = 1$), the clock must respect service and travel time; otherwise, the inequality is relaxed by M .

(2.2d)

Ensures that the drone departs from i , visits j and returns to the truck at k before it departs; linearized via M .

(2.2e)

Outbound and return air distance limited to autonomy.

(2.2f)

Service can only start within window $[a_i, b_i]$.

Solution heuristics Solving the exact MILP is NP-HARD; instances above 30 customers become intractable in academic execution time. The most efficient approaches combine:

- *Savings + Split*: uses the Clarke–Wright algorithm to construct the ground route and then "breaks" long edges that meet battery criteria to replace with drones (WANG *et al.*, 2019).
- *Large Neighborhood Search* (LNS-PD): randomly removes blocks of customers, reintroducing them via drone and motorcycle insertions iteratively.

- *Matheuristic Branch-and-Price + VND*: solves exact sub-routes by *column generation* and applies *Variable Neighborhood Descent* to the global solution (see (CONCEIÇÃO *et al.*, 2022) for Brazilian instances).

In practice, a drone operational radius of ≤ 5 km (Class 3 BVLOS) is also imposed, which drastically reduces the search space without loss of realism for meal deliveries.

2.1.2 Distance-Shortening Rate (DSR)

Delivering by drone is attractive especially when the air route is significantly shorter than the ground route required by the road network. To quantify this geometric gain, Bine *et al.* (BINE *et al.*, 2023) proposed the **Distance-Shortening Rate (DSR)**, a metric that explains a relevant portion of the time savings obtained in air delivery simulations.

Definition Given an origin–destination pair (o, d) , let $d_{\text{air}}(o, d)$ be the Euclidean distance and $d_{\text{street}}(o, d)$ be the shortest distance over the road graph (obtained via Google Maps API). DSR is defined as

$$\text{DSR}(o, d) = 1 - \frac{d_{\text{air}}(o, d)}{d_{\text{street}}(o, d)}, \quad 0 \leq \text{DSR} < 1. \quad (2.3)$$

- $\text{DSR} \approx 0 \Rightarrow$ nearly straight route; little to gain with flight.
- $\text{DSR} \rightarrow 1 \Rightarrow$ streets go around blocks/rivers; great potential for savings.

Why is it important?

- Mode selection.** High DSR indicates that the drone can reduce time substantially; low DSR does not justify the extra cost of flight.
- Droneport sizing.** Bine *et al.* show a correlation of 0.82 between the 75th percentile of DSR in a zone and the optimal number of air bases needed.
- Urban benchmark.** Orthogonal layout cities (e.g., Manhattan) exhibit $\overline{\text{DSR}} \sim 0,30$, while historic centers with organic streets (e.g., Lisbon) reach $\overline{\text{DSR}} \geq 0,55$.

Distance calculation To estimate the direct air distance between restaurant o and customer d , the **Haversine formula** is used, established in position astronomy by Sinnott (SINNOTT, 1984). Assuming Earth as a sphere of mean radius $R=6\,371$ km, the geodetic coordinates (φ, λ) in radians produce

$$d_{\text{air}} = 2R \arcsin\left(\sqrt{\sin^2\left(\frac{\Delta\varphi}{2}\right) + \cos\varphi_o \cos\varphi_d \sin^2\left(\frac{\Delta\lambda}{2}\right)}\right), \quad (2.4)$$

where $\Delta\varphi = \varphi_d - \varphi_o$ and $\Delta\lambda = \lambda_d - \lambda_o$. For pairs at a few tens of kilometers, the relative error is $\leq 0.1\%$, adequate for urban routes.

The road distance d_{street} is obtained by the shortest path (*shortest-path*) on the city's road graph. We employ the **Google Maps API** service, which:

- a) provides optimized routes based on updated road network in real time;
- b) considers current traffic conditions, allowing more accurate travel time calculations that account for congestion;
- c) offers greater precision in estimating distances and travel times, essential for correct sizing of the delivery service.

Thus, the pair $(d_{\text{air}}, d_{\text{street}})$ feeds the DSR Equation (2.3), while the routing data provided by Google Maps API will also be reused later in the motorcycle–drone route simulation, ensuring greater realism and temporal precision in route calculation.

Calculation algorithm

1. For each order, geocode restaurant o and customer d (`lat`, `lon`).
2. Obtain d_{air} via Haversine formula.
3. Obtain d_{street} and travel time via Google Maps API, considering real-time traffic conditions.
4. Apply (2.3); store DSR_i .
5. Summary statistics: mean, median and 25/75 percentiles.

2.1.3 K-means clustering and centrality

Implementing a Drone-as-a-Service requires defining where to install operation bases (*droneports*) to cover most restaurants without multiplying infrastructure. The problem translates into two sub-steps:

- a) Group origin points using **K-means**, minimizing intra-cluster variance;
- b) Choose, within each cluster, the most "influential" road node, measured by **eigenvector centrality**, but filtered by traffic flow to avoid already saturated zones.

1. K-means algorithm Let $\{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset \mathbb{R}^d$ be the restaurant coordinates. Lloyd's classic objective ((LLOYD, 1982)) is

$$\min_{\mathcal{C}_1, \dots, \mathcal{C}_k} J = \sum_{j=1}^k \sum_{\mathbf{x}_i \in \mathcal{C}_j} \|\mathbf{x}_i - \boldsymbol{\mu}_j\|_2^2, \quad \boldsymbol{\mu}_j = \frac{1}{|\mathcal{C}_j|} \sum_{\mathbf{x}_h \in \mathcal{C}_j} \mathbf{x}_h. \quad (2.5)$$

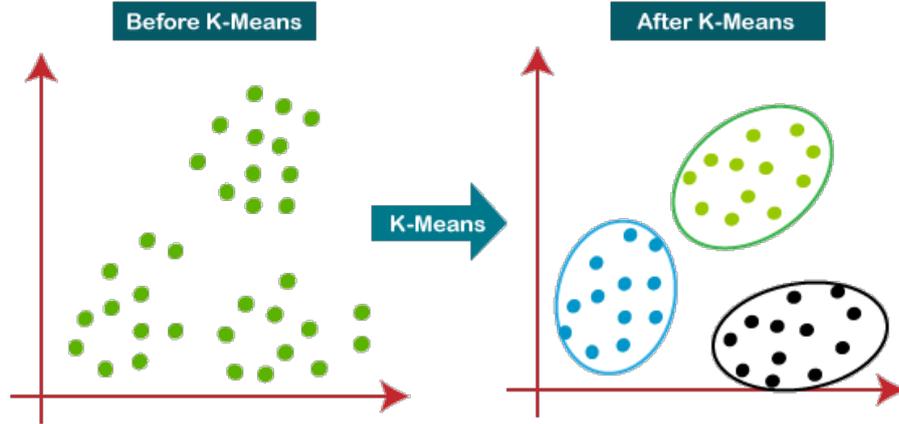


FIGURE 2.3 – Intuitive view of the K-means algorithm: *before* (left) points are distributed without labels; *after* (right) the method partitions the space into k groups minimizing internal variance (LLOYD, 1982; ARTHUR; VASSILVITSKII, 2007).

Figure 2.3 evidences the role of K-means: transforming an amorphous set of coordinates into cohesive clusters, each summarized by its centroid. It is this centroid that will serve as a "raw candidate" for the initial location of *droneports* before applying the eigenvector centrality filter.

Lloyd algorithm steps

1. Initialize centroids $\{\boldsymbol{\mu}_1^{(0)}, \dots, \boldsymbol{\mu}_k^{(0)}\}$ with **k-means++** (ARTHUR; VASSILVITSKII, 2007).
2. *Assignment*: for each point, choose the nearest centroid (L_2 metric).
3. *Recalculation*: recompute $\boldsymbol{\mu}_j$ by (2.5).
4. Repeat 2–3 until convergence ($\Delta J < \varepsilon$).

Complexity: $O(nkdI)$, where I is the number of iterations (generally < 10).

Choice of k . We use the "elbow" criterion: plot $J(k)$ vs. k and choose the smallest k beyond which the marginal reduction of J becomes negligible; equivalent to the point where the curve forms an elbow.

2. Eigenvector centrality Once clusters are defined, we map each street in the cluster to a directed graph $G = (V, E)$. **Eigenvector centrality** measures how well connected a node is to other also central nodes (BONACICH, 1987):

$$\lambda c(v) = \sum_{u \in V} A_{vu} c(u), \quad c(v) \geq 0, \quad (2.6)$$

where A is the adjacency matrix (weighted by length or flow) and λ is the largest eigenvalue of A . We interpret $c(v)$ as "topological influence".

3. Droneport selection Inspired by Bine et al. (BINE *et al.*, 2023), we filter only vertices with moderate vehicle flow — above the mean, but not so high as to create a bottleneck:

$$T = \{v \in V : \mu_{\text{flow}} < f(v) < \mu_{\text{flow}} + \sigma_{\text{flow}}\}, \quad v^* = \arg \max_{v \in T} c(v).$$

This vertex v^* becomes the *droneport* of the cluster.

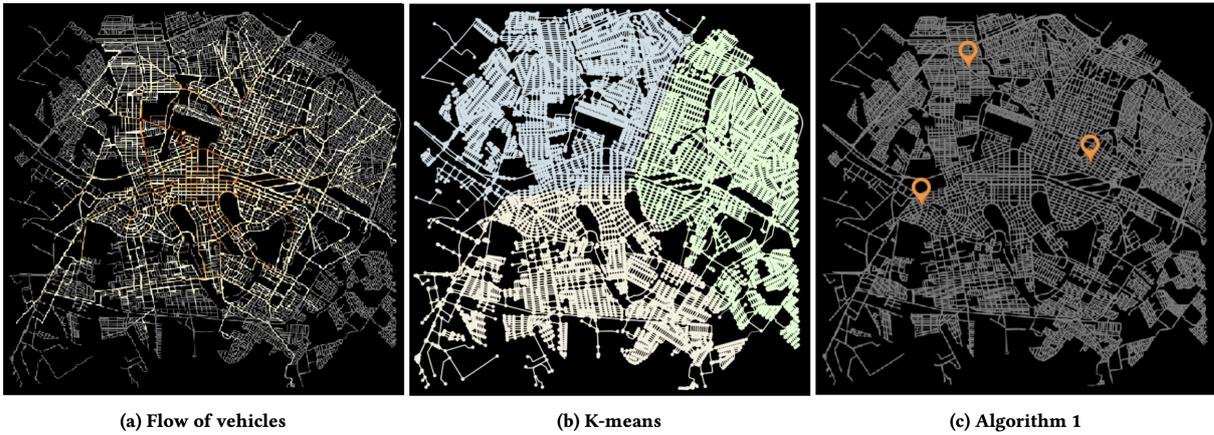


FIGURE 2.4 – Three-step procedure to locate *droneports*, according to Bine *et al.* (BINE *et al.*, 2023). (a) Aggregated vehicle flow (0–24 h); lighter roads concentrate greater traffic. (b) K-means result with $k = 3$, revealing relatively homogeneous zones. (c) Final position of droneports after applying the flow filter and selecting the node with highest eigenvector centrality in each cluster.

Revisited example — synthesis of Bine *et al.*'s method Figure 2.4 reproduces, on a reduced scale, the procedure detailed by Bine et al. (BINE *et al.*, 2023). In panel (a) the road graph is colored by vehicle flow obtained via GPS data and traffic counts; panel (b) shows the K-means partition into three clusters (k selected by the elbow of $J(k)$); finally, panel (c) highlights, in orange, the vertex with highest eigenvector centrality within the flow interval $[\mu_{\text{flow}}, \mu_{\text{flow}} + \sigma_{\text{flow}})$ — exactly the algorithm presented above. The authors report an average reduction of 18 % in droneport–customer distance compared to a merely geometric choice, empirically validating the criterion adopted in this work.

2.1.4 M/M/c queue — fundamentals and sizing

When several drones depart from the same *droneport*, the system can be seen as c parallel "servers" that serve independent orders. The classic assumption — widely validated in meal delivery (FIGLIOZZI *et al.*, 2021) — is *Poisson* arrivals with rate λ

(orders/h) and service times $\text{Exp}(\mu)$ with mean $1/\mu$ hours per order (flight + battery swap). The resulting model is a Kendall M/M/c queue.

Concept and notation In queueing theory, Kendall’s notation $A/B/c$ describes a system with *three* fundamental characteristics (KENDALL, 1953):

- A — **arrival** distribution. The letter M (“Markovian”) indicates a Poisson process, that is, the inter-arrival time is exponential with rate λ .
- B — **service** distribution. Another M means that service duration is also exponential, with mean $1/\mu$.
- c — number of parallel **servers** (simultaneous drones in our case).

A M/M/c queue also assumes:

- Infinite queue capacity (no order is lost);
- *First-Come, First-Served* service discipline;
- Infinite customer population — arrival of an order does not alter the global rate λ .

The model is analytical because the system state ($N =$ number of orders) forms a continuous-time Markov chain; the closed formulas for $P(N = n)$, W_q and W (seen below) derive from this property.

Why M/M/c for drones? Food order arrivals usually follow Poisson in intervals of 5–15 min (FIGLIOZZI *et al.*, 2021), and the natural variability of flight time (and battery swap) approximates an exponential. Furthermore, multiple drones departing from the same *dronoport* fit the parameter c . Thus, M/M/c offers a balance between realism and analytical simplicity to size the minimum fleet that respects the SLA. The M/M/c model is adopted as a long-term equilibrium approximation, recognizing that SARPAS-NG operational windows and hourly peaks may introduce non-Poissonian behavior.

Parameters and notation

$$\rho = \frac{\lambda}{c\mu} \quad (\text{utilization rate, must be } < 1); \quad \alpha = \frac{\lambda}{\mu}.$$

State probability

$$P(N = n) = \begin{cases} \frac{\alpha^n}{n!} P_0, & 0 \leq n < c, \\ \frac{\alpha^n}{c! c^{n-c}} P_0, & n \geq c, \end{cases} \quad P_0 = \left[\sum_{n=0}^{c-1} \frac{\alpha^n}{n!} + \frac{\alpha^c}{c!(1-\rho)} \right]^{-1}.$$

Average waiting time (Erlang-C)

$$P_W = \frac{\alpha^c}{c!(1-\rho)} P_0, \quad W_q = \frac{P_W}{c\mu - \lambda}, \quad W = W_q + \frac{1}{\mu}. \quad (2.7)$$

Symbol legend

- λ Average **arrival** rate of orders (orders per hour).
- μ Average **service** rate of each drone (orders completed per hour); $1/\mu$ is the average cycle time.
- c Number of parallel **servers** (drones operating simultaneously).
- $\rho = \frac{\lambda}{c\mu}$
System **utilization** (fraction of time drones are busy).
- $\alpha = \frac{\lambda}{\mu}$
Aggregated traffic rate; auxiliary notation to simplify powers in terms α^n .
- P_0 Probability of the system being **empty** (no order in service or queue).
- P_W Probability of an order finding a **queue** (all drones busy at arrival instant).
- W_q Average **queue waiting** time (hours), calculated by Erlang-C formula.
- $W = W_q + 1/\mu$
Average **total** time in system (waiting + service); must be less than or equal to SLA.
- N Random variable "number of orders" in system; $P(N = n)$ is given by the piecewise distribution (customers $< c$ and $\geq c$).

As shown by Gross & Shortle (GROSS *et al.*, 2018), (2.7) minimizes W for a given SLA by choosing the smallest integer c that satisfies $W \leq \text{SLA}$.

2.2 Drone energy modeling

The energy consumption of a multirotor impacts (i) **OpEx**, since electricity is proportional to Wh consumed, and (ii) **CapEx**: the larger the battery, the greater the weight and number of packs needed to maintain flight cadence. Therefore we use a *segment-by-segment* model that decomposes power into three components, according to momentum and blade-element theory (LEISHMAN, 2006; DORLING *et al.*, 2017).

Symbol table

m	total mass (drone + payload) [kg]
g	gravitational acceleration (9.81 m/s ²)
ρ	air density (1.225 kg/m ³ at sea level)
A	rotor disk area (πR^2) [m ²]
S	frontal body area [m ²]
C_D	body drag coefficient
V	drone forward velocity [m/s]
v_w	headwind component [m/s]
η	overall electrical–mechanical efficiency

2.2.1 Power in stationary flight (hover)

Froude’s theory gives the induced velocity at the disk $v_i = \sqrt{mg/(2\rho A)}$. The ideal power would be $P_i = mg v_i$. Correcting for electrical and propeller losses ($\eta \simeq 0.72$) yields

$$P_{\text{hover}} = \frac{(mg)^{3/2}}{\eta\sqrt{2\rho A}}. \quad (2.8)$$

2.2.2 Power in translation

During cruise, the *parasitic* power of the body is added (SIVRIOGLU; TEMIZ, 2020):

$$P_{\text{parasite}} = \frac{1}{2} \rho S C_D (V + v_w)^3. \quad (2.9)$$

The cubic term shows strong sensitivity to headwind: a 20

2.2.3 Energy per order

We divide the route into N segments (climb, outbound cruise, descent, return cruise, landing hover):

$$E_{\text{pedido}} = \sum_{k=1}^N \left[P_{\text{hover},k} + P_{\text{parasite},k} \right] \Delta t_k, \quad \Delta t_k = \frac{d_k}{V_k + v_w \cos \theta_k}. \quad (2.10)$$

Blade profile power ($< 5\%$ of total at low Reynolds) was absorbed into P_{hover} via factor η , following Filiopoulou et al. (FILIOPOULOU; AL., 2025).

Calculation procedure — step by step

1. **Segment the flight:** divide each delivery into climb, outbound cruise, descent, return cruise and landing hover.
2. **Compute powers:** for each segment, apply (2.8) (hover) and (2.9) (parasitic), considering segment mass (loaded or empty) and projected wind vector.
3. **Time estimate:** $\Delta t_k = d_k / (V_k + v_w \cos \theta_k)$, where V_k is cruise velocity and θ_k is the relative wind angle.
4. **Energy summation:** $E_{\text{pedido}} = \sum_k (P_{\text{hover},k} + P_{\text{parasite},k}) \Delta t_k$.
5. **Validate battery:** ensure that $E_{\text{pedido}} \leq 0.9 E_{\text{bat}}$; otherwise, classify the order as "not serviceable by drone".

Integration into cash flow For a 12-month horizon, total energy is $E_{\text{ano}} = \sum_{p=1}^{N_{\text{ped}}} E_{\text{pedido},p}$. Multiplying by the average residential tariff of R\$0,73/kWh (ANEEL/2025) yields the annual electricity cost, a direct OpEx item. The model also tracks the number of charge cycles per battery; dividing the pack acquisition cost (C_{bat}) by $\text{Cycles}_{\text{vida}} \approx 400$ yields unit depreciation, which enters as amortized CapEx.

2.3 Economic–financial evaluation

The viability of an aerial drone delivery system must be evaluated from two complementary perspectives: (i) **logistical efficiency** (as modeled in previous sections) and (ii) **economic-financial sustainability**, which determines the operational capacity of the service compared to the current motorcycle delivery model.

Before evaluating investment attractiveness for stakeholders — which would require long-term return analyses, net present value and risk metrics — it is fundamental to establish whether the DaaS model can operate with competitive **OpEx** relative to the ground mode. This is a critical prerequisite: if even in a simplified analysis the drone's OpEx significantly exceeds that of the motorcycle, there is no point in advancing to investor return analyses, as the model would not be operationally sustainable.

The approach adopted in this work consists of calculating the CapEx and OpEx of the drone service and comparing them directly with the costs of the motorcycle delivery

model, establishing an **acceptable OpEx target** that allows competing in the current market. Once this operational viability is achieved, the next steps for an investor return analysis in long-term horizons can be envisioned.

Cost structure

CapEx (Capital Expenditure). Represents the capital investment made in year zero. For DaaS, this includes: drone acquisition (or long-term leasing contract), extra batteries (sized based on energy per delivery from Section 2.2), parachutes, embedded sensors, *droneport* structure, control software, communication infrastructure, regulatory licensing and operator training.

This value immobilizes resources and, in general, does not have liquidity in the short term. However, the analysis proposed here initially focuses on operational cost comparison, leaving the impact of CapEx on long-term profitability for later stages of evaluation.

OpEx (Operational Expenditure). These are recurring operational costs per unit of time (generally monthly or annual). In the case of DaaS, the main components include:

- Electricity consumption (estimated in Section 2.2, with wind correction);
- Preventive and corrective maintenance (replacement of rotors, sensors and batteries);
- Aircraft and civil liability insurance;
- Remote operator and support staff salaries;
- Navigation and communication software amortization;
- *Leasing* cost (if drones are rented instead of purchased);
- Operational licenses and regulatory fees.

OpEx directly affects profitability per delivery and determines service competitiveness relative to the motorcycle model. Therefore, it is the central indicator of this preliminary viability analysis.

Comparison with the motorcycle model

To establish DaaS operational viability, we compare the OpEx per delivery of the drone service with the equivalent OpEx of the current motorcycle delivery model. This comparison allows identifying:

- a) **Immediate competitiveness:** whether the drone's OpEx is within an acceptable range relative to the motorcycle (possibly tolerating a premium justified by logistical advantages, such as shorter delivery time);
- b) **Target OpEx goal:** what should be the maximum operational cost per delivery for the model to be viable in the current market;
- c) **Component sensitivity:** which OpEx items (energy, maintenance, insurance, etc.) most impact competitiveness and where optimization is possible.

If the drone's OpEx significantly exceeds that of the motorcycle — even considering eventual efficiency gains — the model would not be operationally sustainable, regardless of investor return potential. On the other hand, competitive or even lower OpEx than the ground mode opens the way for deeper financial return and investor attractiveness analyses in medium and long-term horizons.

Next steps

Once operational viability is established through CapEx/OpEx comparison, the analysis can evolve to:

- Investor return evaluation through metrics such as Net Present Value (NPV), Return on Investment (ROI) and payback time (*pay-back*);
- Sizing of the necessary initial investment and structuring of financing or *leasing* models.

However, this more advanced stage only makes sense after confirming that the model operates with acceptable OpEx compared to the ground alternative — which constitutes the main focus of this section.

2.4 Brazilian regulatory framework

For commercial drone delivery operations in urban areas in Brazil, the regulatory framework follows mainly **RBAC-E No. 94** (CIVIL, 2023b), complemented by regulations from ANAC, DECEA (AÉREO, 2022) and ANATEL (OES, 2019). In this study, we consider the scenario of Class 3 drones (maximum takeoff weight up to 25 kg), operating **BVLOS** (Beyond Visual Line of Sight), with partial flight autonomy.

Technical and operational requirements

Below, we list the main legal requirements and their direct impacts on system modeling:

- **Maximum takeoff weight (MTOW):** limited to 25 kg.
 - Imposes limit on drone mass, battery and payload.
 - Already considered in energy modeling and determines viability of multiple deliveries per flight (CapEx \times routes trade-off).
- **Maximum speed:** 120 knots (\approx 222 km/h).
 - No direct practical impact on DLV-1, whose cruise speed is lower.
- **Operational ceiling:** 400 ft AGL (\approx 120 m), except specific NOTAM.
 - Limits cruise altitude, affecting aerodynamic performance and energy consumption modeling in cruise.
- **Mandatory safety equipment:**
 - Remote *kill-switch* and certified emergency parachute.
 - Add costs per drone in CapEx, according to current market quotation (BRASIL, 2023).
 - The additional parachute mass must be added to the aircraft empty weight (impacting energy per delivery).
- **Mandatory flight plan (SARPAS-NG):**
 - Request must be sent with \geq 24 h advance (ICA 100-40) (AÉREO, 2022).
 - Restricts operational flexibility, requiring scheduled takeoffs in waves (batch scheduling).
 - Reflected in the M/M/c queue model as arrival concentration and demand peaks in fixed windows.
- **Mandatory insurance (RC-RETA):**
 - Civil liability with minimum coverage of R\$ 500 thousand (PRIVADOS, 2021).
 - Added to annual fixed OpEx.
- **Radio station licensing (ANATEL):**
 - Drones with control link above 1 W power must be certified (OES, 2019).
 - Added to CapEx, in addition to certification time (\approx 3 months).

Integration into the technical–economic model

All requirements above were integrated into the model as follows:

- **CapEx** — costs with parachutes, certifications, kill-switch and ANATEL fee.
- **OpEx** — insurance and regulatory audits compose the annual fixed cost and directly impact cash flow.
- **Energy model** — altitude and weight restrictions alter flight efficiency and autonomy (Section 2.2).
- **M/M/c queue** — the need for operations in hourly windows (SARPAS-NG) introduces arrival peaks (non-Poisson) and may generate higher waiting times in critical periods.

Furthermore, the operational team will need to be trained and certified to operate drones in BVLOS according to RBAC-E 94 guidelines (CIVIL, 2023b), which adds training costs and may represent entry barriers in regions with shortage of certified operators.

2.5 Related work and positioning

This section positions the present work in the context of existing literature, highlighting specific contributions and characteristics of the Brazilian context that differentiate it from previous studies.

Similar work in the literature

Several studies address the problem of routing with drones in last-mile deliveries, but most focus on international contexts and do not incorporate the specificities of the Brazilian regulatory and operational market.

Routing and optimization. Li et al. (LI *et al.*, 2021) and Han et al. (HAN *et al.*, 2023) propose mathematical formulations for VRP-D with time windows, which serve as the theoretical basis for this work. However, these studies do not consider batch operation restrictions (batch scheduling) imposed by the Brazilian SARPAS-NG system, which significantly alters the arrival pattern and queue modeling. Furthermore, works such as Wang et al. (WANG *et al.*, 2019) and Conceição et al. (CONCEIÇÃO *et al.*, 2022) develop efficient heuristics for VRP-D, but do not integrate detailed energy modeling nor comparative economic-financial analysis with the ground model.

Infrastructure location. Bine et al. (BINE *et al.*, 2023) introduce the DSR index and propose a methodology for *droneport* location using K-means and eigenvector centrality, a methodology that is adopted and adapted in this work. However, Bine et al. validate their method in cities with urban layout distinct from the Brazilian context (e.g., Lisbon, European cities), where population density and road network present different characteristics. This work adapts the methodology to the Brazilian urban context, incorporating real vehicle flow data and demand patterns from the Brendi platform.

Energy modeling. Studies such as Filiopoulou et al. (FILIOPOULOU; AL., 2025) and Sivrioglu et al. (SIVRIOGLU; TEMIZ, 2020) present energy consumption models for drones, but often consider only ideal cruise scenarios, without incorporating wind variations, variable mass during flight and operational altitude restrictions (400 ft AGL from RBAC-E 94). This work integrates these variables into the energy model, resulting in more realistic estimates for the Brazilian context.

Economic-financial evaluation. Literature on economic viability of drone deliveries is still incipient, especially for the Brazilian market. International works often assume energy, insurance and regulatory costs different from those observed in Brazil. This study incorporates real residential electricity costs (R\$0,73/kWh according to ANEEL/2025), RC-RETA insurance premiums according to quotation (SEGUROS, 2024), and ANAC/DECEA certification costs, providing a CapEx/OpEx analysis specific to the national context.

Specific characteristics of this work's context

This work differentiates itself by incorporating the following specific characteristics of the Brazilian context:

- a) **Real data from the Brendi platform.** Unlike studies that use synthetic data or from other international platforms, this work is based on real orders intermediated by the Brendi platform, capturing demand patterns, geographic distribution of restaurants and customers, and time windows specific to the Brazilian market.
- b) **RBAC-E 94 regulatory framework.** All operational restrictions of the Brazilian regulation are incorporated into the model: maximum weight of 25 kg, operational ceiling of 400 ft AGL, mandatory parachutes and kill-switch, need for SARPAS-NG flight plans with 24 h advance, and RC-RETA insurance coverage of R\$500 thousand. These restrictions directly impact operational viability and costs, differentiating this study from works that assume more flexible regulations.
- c) **Direct comparison with motorcycle delivery.** The economic-financial analysis explicitly compares the OpEx of the DaaS service with the operational costs of the

current motorcycle delivery model in Brazil, establishing an operational competitiveness goal. International works often compare drones with generic ground vehicles, without considering the specificities of the motorcycle delivery model predominant in the Brazilian market.

- d) **Validation with international practical cases.** Simulation results are validated against metrics reported by real commercial operators (Speedbird Aero in Brazil, Wing in the USA, Meituan in China), ensuring that the model captures behaviors observed in commercial operations at scale, not just pilot demonstrations.

Contributions and differentiators

In summary, this work contributes to the literature by:

- Integrating logistical (VRP-D), energetic and economic-financial modeling in a unique framework for DaaS viability evaluation;
- Adapting consolidated methodologies (DSR, K-means, M/M/c) to the Brazilian regulatory and operational context;
- Using real data from a national delivery platform, capturing specific patterns of the Brazilian market;
- Providing quantitative CapEx/OpEx comparative analysis with the motorcycle delivery model, establishing objective operational viability criteria.

With these mathematical, energetic and regulatory foundations established, we proceed to the data collection and simulation methodology described in Chapter 3.

3 Materials and Methods

In this part of the work we move from theoretical foundation to practical application. Our objective is to demonstrate, end-to-end, *how* to transform raw order data into logistical, energetic and economic-financial indicators capable of answering the research question formulated in Chapter 1. This chapter describes, therefore, the complete *pipeline* — from datasets used to cash flow construction — highlighting the advantages and limitations of each technique employed.

To guide the reader, the section is organized as follows:

- **Section 3.1 – Data and inputs:** presents the origins, justifications and formats of the data used (Brendi orders, Google Maps road network, ERA5/INMET climate series and cost tables), in addition to the preprocessing applied.
- **Section 3.2 – City and period selection:** details the selection criteria for the pilot area — demand density and temporal stability.
- **Section 3.3 – Operational pipeline:** describes, step by step, the computational methods: distance calculation (*Haversine* × Google Maps API), DSR index, *droneport* location (K-means + eigenvector centrality), modal classification of orders and fleet sizing via M/M/c queue.
- **Section 3.4 – Economic-financial model:** presents the CapEx and OpEx tables for the two compared scenarios (Motoboy and DaaS).
- **Section 3.5 – Tools and reproducibility:** lists the development environment (Python 3.11, Docker, Google Routes API, libraries for geospatial and data analysis), simulation tools (SimPy) and the Git repository that ensures traceability of results.

This structure ensures that each component — data, analytical methods and financial metrics — is presented transparently and reproducibly, allowing other researchers or market actors to repeat the experiment or adapt it to their contexts.

3.1 Data and inputs

This section describes, in detail, the sources, formats and preprocessing that will be *planned* to feed the pipeline presented in Section 3.3. Table 3.1 summarizes the main characteristics of each dataset; in the following subsections, it is explained how each resource should be obtained, cleaned and employed in the model.

TABLE 3.1 – Overview of datasets employed.

Dataset	Format	Relevant fields	Function in model
Brendi Orders	CSV	Order ID, Store/Customer lat/lon, Date, Order Total Price, Delivery Type, Status	Estimation of λ , DSR, revenues
Google Routes API	JSON/REST	<code>distanceMeters</code> , <code>duration</code> , <code>travelMode</code> (DRIVE), <code>routingPreference</code> (TRAFFIC_AWARE)	d_{street} , ground travel time
ERA5 (Wind)	NetCDF	u, v at 10 m, 1 h (single point)	<i>Flyable</i> matrix
INMET (Rain)	CSV	Rain (mm), hourly; other variables available	<i>Flyable</i>
Unit costs	CSV	CapEx: Drones, Batteries, Safety, Infrastructure, Communication, Regulatory, Software, Training; OpEx: Salaries, Amortization, Rent, Insurance, Batteries, Maintenance, Licenses, Energy	CapEx, OpEx, NPV
DLV-1 Spec.	YAML	Mass, rotor, η , price	Energy modeling, CapEx

3.1.1 Brendi order dataset

The project will use an anonymized *dump* provided by Brendi, stored in CSV, containing approximately 1 048 575 records (1,048,575 orders). Access will be formalized via NDA exclusively for academic purposes, ensuring that all personal data remain protected by cryptographic hashing.

The dataset has 16 columns, organized as follows:

- **Anonymized identifiers:** Hashed Order ID, Hashed Store ID, Hashed Customer ID;
- **Delivery information:** Delivery Type, Order Status, Order Total Price, Payment Method, Platform;
- **Temporal location:** Date (request timestamp);
- **Geographic location:** Country, City, State, Lat and Lng (customer coordinates), Store Lat and Store Lng (restaurant coordinates).

From these fields, spatiotemporal demand patterns will be extracted, the hourly arrival rate λ will be estimated, DSR will be calculated using restaurant and customer coordinates, and the logistical, energetic and financial modules will be fed. The **Order Total Price** field will allow revenue estimation, while **Delivery Type** and **Order Status** will assist in filtering valid orders for analysis.

Preprocessing will be conducted in four phases. In the first, duplicates based on **Hashed Order ID** will be removed and only orders with **Order Status** indicating successful completion will be filtered. Next, coordinate validation will proceed, excluding records with null values in **Lat**, **Lng**, **Store Lat** or **Store Lng**, points outside the urban perimeter of the analyzed city, or whose total distance (calculated via Haversine between restaurant and customer) is less than 50 m, in addition, total distance outliers (calculated via Haversine between restaurant and customer) with value greater than 3 times the standard deviation will be discarded; this criterion eliminates test orders or geocoding errors. In the third phase, records will be filtered by **Store Lat** and **Store Lng** to isolate the pilot city chosen in Section 3.2. Finally, a continuous temporal window based on the **Date** field that does not include extended holidays or extreme seasonal events will be defined, according to parameters established in Section 3.2. After these steps, the resulting set will serve as the central empirical basis for simulations.

3.1.2 Google Routes API

Ground travel distances and times will be obtained through the **Google Routes API**, a service that provides optimized routes considering real-time and historical traffic conditions. For each origin–destination pair (restaurant–customer), the API returns:

- **distanceMeters:** total route distance in meters;
- **duration:** estimated travel time in seconds;
- **polyline:** encoded polyline of the route (when necessary).

Requests are configured with `travelMode: "DRIVE"` (approximate mode for motorcycle, since the API does not have a specific mode for this vehicle) and `routingPreference: "TRAFFIC_AWARE"`, which incorporates historical traffic patterns and real-time conditions when available. For orders with dates in the past, the API uses average historical traffic patterns, providing realistic estimates even without real-time traffic data from the exact moment.

Returned data are stored in the processed dataset as `Route Distance (km)` and `Route Duration (s)`, and used for:

- a) Calculate DSR through comparison between air distance (Haversine) and road distance returned by the API;
- b) Estimate ground travel times for the VRP-D model;
- c) Compare air versus ground mode efficiency.

3.1.3 ERA5 and INMET climate series

It is intended to characterize meteorological restrictions from two open and complementary sources. Wind intensity and direction will be extracted from ERA5 reanalysis, obtained for a representative point of the pilot city (coordinates extracted from the centroid of Brendi restaurants), with hourly resolution. The u and v wind components at 10 m height are stored in a `NetCDF` file with temporal dimensions (one hour) and spatial (a single point). Precipitation will be obtained from the INMET automatic station closest to the same point, also with hourly resolution. The INMET `CSV` file contains multiple meteorological variables, with the `Chuva (mm)` column used to determine flight conditions.

From these data, flight *scenarios* will be defined: *conservative*, *moderate* and *aggressive*. Each scenario will give rise to a binary matrix `flyable_day[day, hour]`, which will indicate whether an order can be allocated to the air mode or must migrate to contingency motorcycle delivery. This procedure avoids fixing an arbitrary cutoff value, while providing sufficient variability for the sensitivity analysis proposed in Section 3.4.

3.1.4 Cost structure

All prices necessary for economic-financial evaluation will be consolidated in a `costs.csv` file, structured with columns `item`, `valor_unidade`, `freq_ano` and `categoria` (CapEx or OpEx). The file contemplates eight CapEx categories: Drones, Batteries, Safety, Infrastructure, Communication, Regulatory, Software and Training. OpEx items include: Salaries and Operation, Amortization, Facility Rent, Insurance, Battery Replacement,

Maintenance, Operational Licenses and Electricity Consumption. The project plans to collect these values directly from public sources — for example, the average energy tariff from ANEEL, salary tables from unions and insurance quotations — always on the cutoff date immediately prior to simulation. When necessary, amounts will be converted to the cash flow base year through the accumulated IPCA index.

3.1.5 Drone specification

Instead of fixing the analysis to a specific commercial model, we opt at this stage to define a *reference profile* compatible with RBAC-E 94 Class 3 BVLOS (MTOW up to 25 kg) and capable of transporting, at minimum, one kilogram of payload. Definitive technical parameters — empty mass, battery capacity, rotor disk area, overall powertrain efficiency and acquisition value — will be filled once the manufacturer is selected, according to market criteria and leasing availability.

To ensure traceability, it is planned to concentrate these characteristics in a `drone_spec.yaml` file. The file will contain standardized keys, such as `mass_empty`, `battery_Wh`, `payload_max`, `rotor_area`, `eta_overall` and `capex`. During the simulation phase (Section 2.2), the Python script will read this YAML directly, allowing replacement of the reference drone with another model — or even evaluating scenarios with different hardware generations — without changing the code body.

Therefore, flexibility is maintained to test multiple configurations before converging to the final specification, while preserving the integrity of energy calculations and CapEx estimates presented in Chapter 3.4.

3.2 City and analysis period selection

Before starting any simulation, it is necessary to define *where* and *when* the experiment will be executed. The pilot city choice and temporal window will directly influence result reliability, as it determines the degree of representativeness of the order set, road network density, climate conditions and data volume to be processed. It is intended, therefore, to follow a two-stage selection protocol: first spatial criteria linked to Brendi restaurant coverage are analyzed; then temporal criteria that ensure a stable historical series are evaluated.

Spatial coverage criterion

The starting point was to analyze, for each candidate city, the spatial density of orders present in the Brendi *dataset*. The analysis considered the number of orders per city and its area in square meters, calculating order density per square meter. The objective was to ensure that the selected city presented sufficient order density to provide an adequate understanding of the real demand situation, allowing routing models to be tested in a representative context. The city was defined based on the corresponding geographic area, and those with higher order density per square meter were prioritized in selection.

Temporal stability criterion

Once the study area is defined, the second step consists of electing the reference temporal window. Due to technical issues of the Brendi platform and strategic and commercial restrictions of the company itself, it was possible to obtain data only from the three most recent available months: June, July and August. Given this limitation, the analysis was initiated based on this reduced temporal window, requiring data annualization to estimate demand patterns over a complete cycle. From the `Date` field of the Brendi dataset, daily order counts were aggregated after date format conversion, allowing characterization of the hourly arrival rate λ and extrapolation of these values to annual estimates.

Integration with subsequent stages

The urban polygon and temporal window chosen in this section will serve as reference for all modules in Chapter 3: requests to Google Routes API will be filtered for coordinates within the defined spatial limit; the climate matrix (`flyable_day`) will be generated exactly for the days and hours in question; and economic-financial projections will inherit the same period to ensure comparability between scenarios. In this way, methodological cohesion is ensured and decisions made a posteriori disconnecting from premises established in the selection phase are avoided.

3.3 Operational pipeline

With the study area and temporal window definition completed, the next step consists of transforming raw databases into a coherent set of logistical, energetic and economic indicators. For this, an integrated processing flow will be established — the *operational pipeline* — which develops in nine chained stages. All stages will be implemented in Python 3.11, using `pandas` for tabular manipulation, `requests` for calls to Google Routes

API and SimPy for event-driven simulation.

In the first stage, each record in the Brendi *dataset* will use coordinates already present in fields `Store Lat/Store Lng` (origin) and `Lat/Lng` (destination), after validation of null values and geographic coherence as described in Section 3.1.1, performing data cleaning.

In the second stage, request analysis by city will be constructed and the study city will be defined based on order concentration identified in the *dataset*.

In the third stage, the script will calculate two independent distances for each origin–destination pair: the geodetic air distance, obtained with the Haversine formula applied to restaurant and customer coordinates, and the shortest road distance obtained via Google Maps Routes API (according to Section 2.1.2 of Chapter 2), which considers real-time traffic conditions. The ratio between these two metrics will produce, for each order, the *Distance-Shortening Rate* (DSR), an indicator that summarizes the potential geometric gain of the air mode.

The fourth stage will estimate travel times for both modes. The ground time will be exactly the value returned by Google Maps Routes API, which already incorporates traffic conditions and optimized routes. Flight time, in turn, will use the average cruise speed predicted for the reference drone, also incorporating climb, descent and operational time phases at the *droneport*. These values will be stored so they can be replaced, later, by simulated times in the queue model.

The fifth stage will advance to *droneport* location. The set of unique restaurants will be extracted from the Brendi dataset through unique `Hashed Store ID`, using coordinates `Store Lat/Store Lng`. This set will be partitioned by K-means, with the number of clusters defined by the elbow method. In each cluster, the road node with highest eigenvector centrality — filtered by a vehicle flow interval that avoids already congested points — will be marked as a *droneport* candidate. These points will anchor the M/M/c queue system described in Subsection 2.1.4.

The sixth stage will construct the *flyable* matrix for days and hours within the requested orders. At this point, it is important to mention that data provided by ERA5 and INMET present temporal lag greater than necessary to obtain data from June to August of the current year. For this reason, data from June to August of the previous year (2024) were used to maintain this coherence. Furthermore, three different scenarios were created — conservative, moderate and aggressive — regarding flight capability. The conservative scenario was chosen to ensure a more pessimistic scenario regarding the drone issue.

The seventh stage will consist of classifying each order as *Drone*, *Mixed* or *Motorcycle*. The algorithm will apply hierarchical rules, starting with meteorological conditions — checking whether the day is *flyable* or not, derived from the hourly matrix constructed

in the previous stage. Next, autonomy rules will be applied to verify whether the drone would be capable of performing the delivery at that distance. Orders with distance greater than the 5 km operational radius were automatically classified as motorcycle deliveries (motorcycle mode). Finally, the last analysis will be temporal cost: if motorcycle delivery time is less than the time needed to take the drone to the *droneport*, perform the operation and deliver with the drone, the motorcycle choice will be maintained. Otherwise, mixed delivery with drone will be chosen.

In the eighth stage, the arrival rate λ will be aggregated hour by hour and by day of week, for each *droneport* from the **Date** field of the Brendi dataset, extracting the hour of each order and associating it to the nearest *droneport* based on the restaurant's **Hashed Store ID**. Taking the average service time μ^{-1} calculated in the energy stage, the M/M/c queue model will size the minimum number of simultaneous drones capable of meeting the Service Level Agreement of twenty minutes. The M/M/c model is adopted as a long-term equilibrium approximation, recognizing that SARPAS-NG operational windows and hourly peaks may introduce non-Poissonian behavior. This sizing will be dynamically re-evaluated whenever the meteorological scenario excludes flight hours.

The ninth stage will calculate the energy consumption of each air mission through the segment-by-segment model presented in Section 2.2. Summing consumption in route, hover and transition phases, battery usage profile and number of charge cycles over an operational year will be obtained.

Finally, the pipeline will consolidate performance indicators: average delivery time reduction, motorcycle kilometers avoided, CO₂ savings and number of orders served by drone. These indicators will feed the economic-financial module described in Section 3.4, concluding the cycle of transforming raw data into managerial decision metrics.

With this processing chain, it is expected to capture the most relevant interactions between demand, infrastructure, climate and costs, allowing holistic evaluation of the technical and financial viability of the Drone-as-a-Service in the Brazilian urban context.

3.4 Economic–financial model

With the logistical and energetic quantification stage completed, the final viability evaluation requires transforming these operational results into comparable monetary costs. The economic–financial model here adopted starts from the premise that, before evaluating investment attractiveness for stakeholders – which would require long-term return analyses, net present value and risk metrics – it is fundamental to establish whether the DaaS model can operate with competitive **OpEx** relative to the ground motorcycle delivery mode.

The approach consists of calculating the CapEx and OpEx of the drone service and comparing them directly with the costs of the motorcycle delivery model, establishing an **acceptable OpEx target** that allows competing in the current market. This is a critical prerequisite: if even in a simplified analysis the drone's OpEx significantly exceeds that of the motorcycle, there is no point in advancing to investor return analyses, as the model would not be operationally sustainable. In this section, how capital (CapEx) and operational (OpEx) costs of the DaaS service will be estimated and how these values will be compared with equivalents from the motorcycle model to identify initial operational viability are described.

CapEx and OpEx

The initial investment (CapEx) will contemplate acquisition, installation and licensing items that have a non-recurring nature. It is expected to include here the physical infrastructure of *droneports*, purchase or leasing contract of aircraft, spare batteries, parachutes, *kill-switch*, ANATEL certification and RBAC-E 94 certification. The exact amount will be extracted from the `costs.csv` file – described in Section 3.1.4 – respecting values in effect on the projection base date.

Operational expenditures (OpEx) will encompass electricity, preventive and corrective maintenance, battery replacement, RC-RETA insurance premiums, remote operator salaries, SARPAS fees and, when applicable, per-flight rental in DaaS modalities. As all fixed and variable costs will be parameterized by range, the possibility of testing scenarios with energy inflation or increase in insurance premium without rewriting the model body will be maintained.

Regulatory integration

Regulatory expenditures identified in Section 2.4 – ANATEL certification fee, certified parachute acquisition, minimum RC-RETA policy and CAER renewal costs – will be posted in appropriate columns of CapEx or OpEx, according to their recurring or non-recurring nature. In this way, cash flow will fully reflect the legal compliance cost, avoiding underestimating the impact of regulation on profitability.

Reproducibility

All financial calculations are implemented in Python using NumPy and pandas for numerical operations and tabular manipulation; consolidated results are exported to a `.xlsx` spreadsheet made available in the project repository. Furthermore, by versioning

the code, it is ensured that any reader can reproduce or audit the model by changing only input values.

With the economic-financial model structured in this way, it is expected to translate the logistical gain into clear capital return metrics, providing an objective basis to recommend – or discard – adoption of the Drone-as-a-Service in the considered pilot city.

Managerial interpretation

The combination of the two techniques enables construction of operational red lines: if the minimum tariff accepted by the market falls below the point where the tornado chart crosses the zero axis, or if the restaurant adoption rate does not reach the indicated threshold, the investment loses attractiveness.

3.5 Tools and reproducibility

The scientific rigor of this work rests on the possibility of any researcher reproducing – or auditing – all results. To achieve this objective, routines described in previous sections will be implemented in a standardized computational environment, balancing portability, performance and data security.

The project base is a Python 3.11 environment containerized in Docker, built from the `python:3.11-slim` image and managed by a `requirements.txt` file that specifies exact dependency versions. The environment includes essential tools for data science: `pandas` (version ≥ 2.0) and `NumPy` (version ≥ 1.24) for tabular data manipulation, `pyarrow` for efficient processing of large data volumes, `requests` for HTTP calls to Google Routes API, `scikit-learn` for clustering and classification algorithms, `SimPy` for discrete event simulation, and `matplotlib` and `seaborn` for static visualizations. For geospatial and road network analysis, `OSMnx` and `NetworkX` are used, complemented by `contextily` and `pyproj` for integration with base maps and geographic projections. ERA5 meteorological data processing is performed through `cdsapi`, `xarray` and `netcdf4` for access and manipulation of NetCDF files. The `Dockerfile` also installs system dependencies necessary for operation of these geospatial libraries, such as `gdal-bin`, `libgdal-dev` and `libspatialindex-dev`, ensuring portability and reproducibility of the environment on different platforms.

All code will be organized in Python modules executable via command line, each accompanied by a `config.yaml` file that centralizes parameters – for example, wind limits for the “conservative” scenario. This modular design avoids dependency on interactive notebooks and facilitates scenario automation by simply swapping configuration files.

Version control of routines will be maintained in a `Git` repository, publicly available on `GitHub`. Datasets under NDA will be in a private repository, identified by integrity *hash*. Large files, such as the road network in `GraphML` or climate data `NetCDF` and `CSV`, will be tracked by `DVC` (`Data Version Control`), allowing code and metadata to remain light while heavy binaries are stored in a configurable remote backend.

Numerical results – logistical metrics, energy consumption, financial indicators – will be automatically exported to a `.xlsx` spreadsheet generated by `pandas`. Essential graphs will be created in the same process, ensuring that numbers and figures are always synchronized.

Finally, the repository will have continuous integration via `GitHub Actions`. Each `push` triggers automatic formatting with `Black`, type checking with `mypy` and execution of unit tests on synthetic samples, reducing the probability of introducing last-minute inconsistencies.

With this infrastructure – reproducible environment, code and data versioning, automated CI and external parameterization – the study preserves transparency, traceability and ease of extension for future research.

4 Discussions

This chapter presents a critical and interpretive analysis of the results obtained through operational simulation and the economic-financial model described in Chapter 3. Here, the numerical data are contextualized in light of the literature reviewed in Chapter 2, the practical experiences of commercial operators mentioned in Chapter 1, and the limitations inherent to the methods employed.

The structure of this chapter is organized into three main axes of discussion. First, the *logistical results* are analyzed — including the operational viability of the 20-minute SLA, the spatial distribution of *droneports*, fleet sizing, and the proportion of orders eligible for aerial delivery. Next, the *economic-financial results* are examined — comparing the operational costs of the DaaS model with traditional motorcycle delivery, evaluating the impact of regulatory costs, and discussing break-even points and scale viability. Finally, the *validation and limitations* of the proposed model are discussed — confronting the simulated metrics with data from real commercial operators, identifying the main simplifications assumed and their implications for generalizing the results.

This three-dimensional approach allows not only presenting the numbers obtained, but also interpreting them critically, identifying points of attention for practical implementation, and establishing paths for future research that may refine or extend the proposed model.

4.1 Discussion of Logistical Results

4.1.1 Preprocessing and Data Cleaning

The preprocessing process executed in *notebook* 01 — `data_cleaning_and_filtering.ipynb` — constituted the fundamental stage to ensure the quality and representativeness of the data that fed all subsequent simulations. This subsection details each stage of this process, discussing the methodological decisions adopted and their implications for the final results.

Stage 1: Initial Exploratory Analysis

The original dataset provided by Brendi contained 1 048 575 records distributed across 16 columns. The initial exploratory analysis revealed important characteristics about data quality. All identification fields (`Hashed Order ID`, `Hashed Store ID`, `Hashed Customer ID`) were complete, with no null values, ensuring referential integrity of the data. The geographic coordinate columns (`Lat`, `Lng`, `Store Lat`, and `Store Lng`) showed 100% completeness, allowing all records to be used for distance calculations and spatial analysis. Regarding auxiliary fields, it was identified that 277 799 records (26.5%) had null values in the `Country` field, while `City` showed only 367 null values (0.035%) and `State` 71 null values (0.007%). These null values did not impact the analysis, as geographic filtering was performed exclusively by coordinates, as discussed below. The `Payment Method` field showed only 13 records (0.001%) with null values, an insignificant percentage that does not compromise subsequent analyses.

Stage 2: Duplicate Removal and Filtering by Operational Criteria

The first cleaning phase consisted of three sequential checks:

Duplicate check: A duplicate analysis was performed based on the `Hashed Order ID` field, unique identifier of each order. The result showed that no duplicate records were present in the dataset — all 1 048 575 IDs were unique. This characteristic indicates excellent initial data quality, probably resulting from a well-structured extraction process by the Brendi platform.

Filtering by delivery type: All records already had `Delivery Type` equal to `DELIVERY`, which indicates that the provided dataset had already been pre-filtered at the source, keeping only delivery-type orders. No records needed to be removed at this stage.

Filtering by order status: Similarly, all 1 048 575 records already showed `Order Status` equal to `DELIVERED`, confirming that only successfully completed orders were present in the dataset. Again, no records were removed.

Implications: The fact that there was no need to remove records in the operational filtering stages demonstrates that the provided dataset was already aligned with the study's analysis criteria. This reduced the risk of selection bias and ensured that the analysis worked with the entirety of available data, maximizing the statistical power of the simulations.

Stage 3: Challenge of City Name Standardization

An initial analysis of city names revealed a critical problem: the dataset contained 2557 unique city names for only 1 048 208 records with city information. This ratio — approximately one different name for every 410 orders — indicated strong standardization issues in textual fields.

Concrete examples illustrate the problem. The city *São José dos Campos* appeared in at least 8 distinct variations: “São José dos Campos” (20,079 orders) as the standard form, “São José dos Campos” (191 orders) with subtle capitalization variation, “Sao Jose dos Campos” (179 orders) without accentuation, “São José Dos Campos” (32 orders) with mixed capitalization, “SAO JOSE DOS CAMPOS” (11 orders) entirely in uppercase, “sao jose dos campos” (6 orders) entirely in lowercase, “São jose dos campos” (5 orders) with inconsistent capitalization, and “São Jose dos campos” (1 order) with another variation. Similar patterns were observed for other cities, such as *Jacareí* (appearing as “Jacareí”, “Jacarei”, “JACAREI”), *Rio de Janeiro* (with variations including “Rio De Janeiro”, “rio de janeiro”, “RIO DE JANEIRO”) and many others.

Methodological decision: Given this standardization issue, filtering by city name would be extremely inefficient and error-prone. A text-based approach would require string normalization (accent removal, lowercase conversion, space handling), manual mapping of all known variations, would present risk of losing valid records due to unforeseen variations, and would complicate maintenance and updates as new data were incorporated.

The adopted solution was to use **geographic filtering based on coordinates**, specifically store coordinates (`Store Lat` and `Store Lng`). This approach offers decisive advantages: coordinates are objective and do not suffer from typing variations, allow defining precise geographic boundaries regardless of administrative boundaries, do not depend on textual data quality, and the same coordinate intervals can be applied to any future dataset, ensuring reproducibility.

Stage 4: Geographic Filtering for São José dos Campos

The choice of São José dos Campos as the pilot city was based on two main criteria:

Order density criterion: Table 4.1 presents the number of orders per city in descending order, along with populations obtained from IBGE (ESTATÍSTICA, 2024) and the respective calculated order densities. Although São Paulo, Rio de Janeiro, and Curitiba showed larger absolute volumes, the density analysis reveals important patterns for pilot city selection.

It is observed that, although São Paulo, Rio de Janeiro, and Curitiba show larger absolute volumes, São José dos Campos has the highest order density (28,28 orders per

TABLE 4.1 – Number of orders per city, population, and order density (descending order by volume).

City	Number of Orders	Population	Density (orders/1000 hab.)
São Paulo	57,814	12,325,232	4,69
Rio de Janeiro	39,592	6,747,815	5,87
Curitiba	37,568	1,963,726	19,16
São José dos Campos	21,213	750,000	28,28

Source: Number of orders — Brendi dataset; Population — IBGE (ESTATÍSTICA, 2024) (2024 estimates); Density — own calculation (number of orders per 1,000 inhabitants).

1,000 inhabitants) and Curitiba (19,16 orders per 1,000 inhabitants). São José dos Campos, with moderate-high density, presents favorable characteristics for routing analysis and infrastructure sizing, concentrating a significant proportion of orders relative to its population.

Operational control criterion: The higher density allows greater control over operational variables — distances between restaurants and customers, temporal demand patterns, and spatial distribution — facilitating validation of the proposed models.

Definition of geographic boundaries: The coordinate intervals were defined as:

$$\text{Latitude : } -23,33667 \leq \text{Store Lat} \leq -23,15875 \quad (4.1)$$

$$\text{Longitude : } -46,02554 \leq \text{Store Lng} \leq -45,75697 \quad (4.2)$$

These values were calibrated to encompass the main urban area of São José dos Campos and adjacent municipalities with significant presence of Brendi restaurants, such as Jacareí. Figure 4.1 visually illustrates the selected region, delimited by geographic markers that define the filtering rectangle boundaries.

Filtering result: Application of this geographic filter resulted in maintaining 23 529 orders (2.24% of the original total) and removal of 1 025 046 orders (97.76% of the original total). Quality verification confirmed that the coordinates of filtered records were within expected intervals: **Store Lat** ranging from -23,31627 to -23,17250, and **Store Lng** ranging from -45,98076 to -45,78449.

Processed Dataset Statistics

The final dataset after all preprocessing presented characteristics indicating good data quality and representativeness. The total volume comprised 23 529 unique and successfully delivered orders, distributed among 56 unique stores (distinct **Hashed Store ID**), indicating adequate diversity of establishments in the region. The customer base totaled 19 599 unique customers (distinct **Hashed Customer ID**), demonstrating concentration of

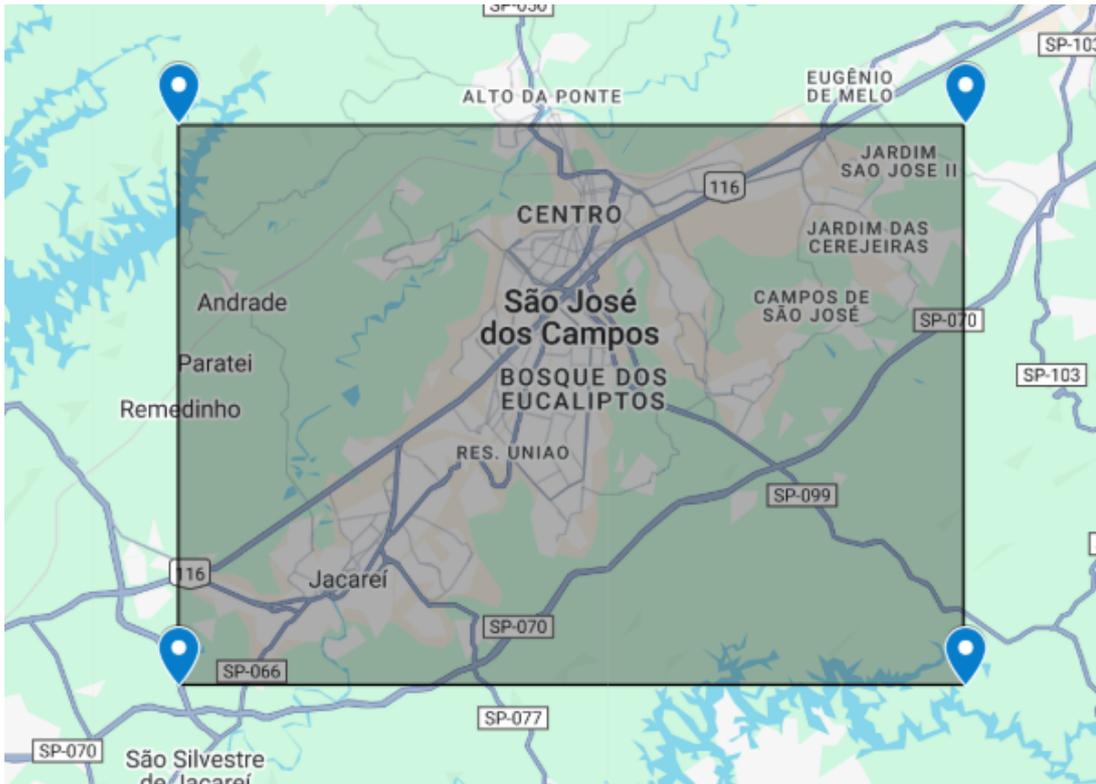


FIGURE 4.1 – Geographic area selected for analysis, delimited by latitude and longitude coordinates. The region encompasses the main urban area of São José dos Campos and adjacent parts, including municipalities such as Jacareí, where there is significant presence of Brendi restaurants.

recurring orders — the customer/order ratio of approximately 0.83 suggests that a significant portion of customers place multiple orders. The average order value was R\$ 66,71, a value consistent with the *food delivery* market in Brazil, resulting in a total processed value of R\$ 1.569.555,31 in the analyzed period. Analysis of city names in the filtered dataset confirmed the effectiveness of the geographic method: “São José dos Campos” (with variations) concentrated 20,487 orders (87.1%), “Jacareí” (with variations) totaled 3,005 orders (12.8%), while other adjacent cities summed 37 orders (0.1%).

Stage 5: Distance Calculation and Outlier Validation

The processing executed in *notebook* 02 — `02_data_haversine_distance_analysis.ipynb` — implemented validation of geodesic distances and removal of outliers as foreseen in Section 3.1.1 of Chapter 3. This stage ensured that only orders with realistic distances were used in subsequent routing and infrastructure sizing analyses.

Haversine distance calculation: For each origin–destination pair (store–customer), the geodesic aerial distance was calculated using the Haversine formula, as described in the third stage of the operational pipeline (Section 3.3). The geodesic distance is necessary to calculate the *Distance-Shortening Rate* (DSR), a metric that compares the potential gain of the aerial mode relative to the terrestrial mode.

Initial statistical analysis: The dataset of 23 529 orders showed minimum distance of 0.000 km, maximum of 462.123 km, mean of 3.060 km, median of 2.537 km, and standard deviation of 3.838 km.

The distribution showed right skewness, with 99% of orders concentrated below 9.434 km (99th percentile). The presence of extreme values — such as the maximum distance of 462.123 km — indicated need for additional validation. Figure 4.2 visually illustrates the concentration of data near zero and the presence of multiple outliers, including an extreme value near 450 km.

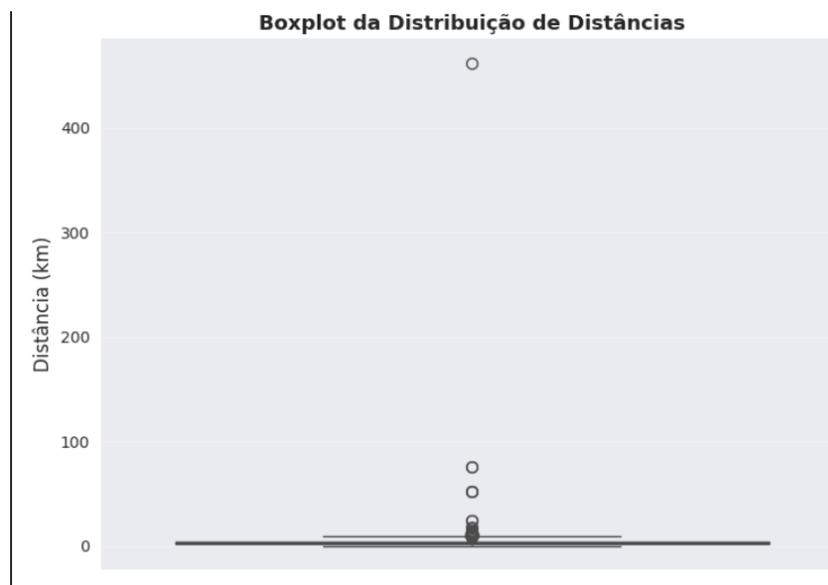


FIGURE 4.2 – Boxplot of delivery distance distribution before filter application. The concentration of data near zero and the presence of extreme outliers evidence the need for statistical validation.

Outlier removal by standard deviation: The ± 3 standard deviation method was applied, a standard statistical criterion that maintains approximately 99.7% of data in a normal distribution. The established limits were from 0.00 km to 14.57 km, resulting in removal of 10 orders (0.04% of total) with distances between 16.321 km and 462.123 km. These extreme values probably result from geocoding errors or test orders.

Additional minimum distance filter: Complementarily, orders with distance less than 50 m were removed, as foreseen in the methodology. These cases may indicate geolocation errors, inconsistent data, or in-store pickup orders (not real deliveries). 184 orders (0.78% of dataset after outlier removal) were removed, with mean distance of 21.32 m.

Final result: After applying both filters, the final dataset contains 23 335 records (99.22% of dataset after geographic filtering), showing minimum distance of 0.051 km, maximum of 14.089 km, mean of 3.048 km, median of 2.560 km, and standard deviation of 2.224 km.

The reduction of standard deviation from 3.838 km to 2.224 km indicates a more con-

centrated and realistic distribution for the urban context of São José dos Campos. Figure 4.3 visually compares distributions before and after filter application, evidencing removal of the long tail of extreme values and concentration of data in a more representative range.

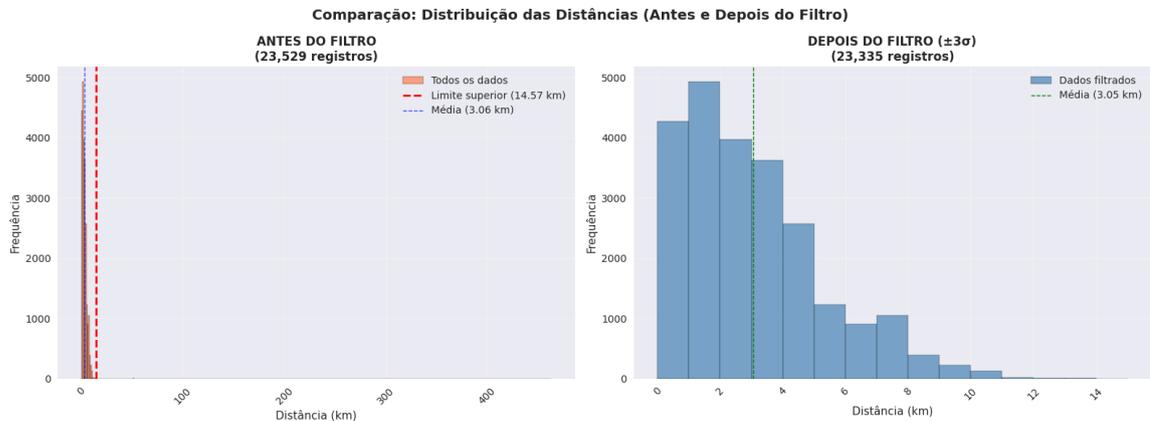


FIGURE 4.3 – Comparison of distance distributions before and after outlier filter application. The left histogram shows the original distribution with extreme values up to 400 km, while the right one presents the filtered distribution concentrated in distances up to 14 km, more representative of the urban context.

The filtered dataset was saved in `orders_sao_jose_dos_campos_filtered.csv` and used as the basis for DSR calculation and routing analyses described in subsequent stages of the operational pipeline.

Stage 6: Road Distance Calculation and DSR

The processing executed in *notebook 03 — 03_analysis.ipynb* implemented the third stage of the operational pipeline (Section 3.3), calculating road distances through the Google Routes API and computing the *Distance-Shortening Rate* (DSR) for each order, as foreseen in the methodology.

Integration with Google Routes API: For each origin–destination pair (store–customer), requests were made to the Google Routes API configured with `travelMode: "DRIVE"` (approximation for motorcycle) and `routingPreference: "TRAFFIC_AWARE"`, as described in Section 3.1.2 of Chapter 3. The API returned the optimized road distance (`Route Distance (km)`) and estimated travel time (`Route Duration (s)`), considering historical traffic patterns for orders with dates in the past.

Parallel processing: Given the need to process 23 335 orders, parallel processing with 10 simultaneous *workers* was implemented, significantly reducing execution time. The code was designed to save progress incrementally (every 50 processed orders), allowing safe interruption and resumption of processing without data loss. At the end, 23 303 orders were successfully processed (99.9% of total), with only 32 failures due to temporary API limitations.

DSR calculation: For each processed order, DSR was calculated according to the formula presented in Section 2.1.2 of Chapter 2:

$$\text{DSR} = 1 - \frac{d_{\text{haversine}}}{d_{\text{street}}} \quad (4.3)$$

where $d_{\text{haversine}}$ is the geodesic aerial distance and d_{street} is the road distance returned by the API.

Statistical results: Statistics calculated on the 23 303 processed orders revealed that road distances showed mean of 4.745 km, median of 3.887 km, ranging from 0.051 km to 22.168 km. Travel times recorded mean of 548 s (9.1 min) and median of 512 s (8.5 min). DSR showed mean of 0.3524 (35.24% potential reduction), median of 0.3363, ranging from -0.2984 to 0.9303. The amplification factor indicates that road distance is, on average, 1.67 times greater than the geodesic aerial distance.

The mean DSR of 0.3524 indicates that, on average, the aerial mode offers potential reduction of approximately 35% in distance traveled relative to the terrestrial mode, validating the hypothesis of significant geometric gain for urban deliveries. The presence of negative DSR values (1.3% of cases) indicates situations where road distance was less than geodesic, possibly due to API inaccuracies or specific characteristics of the local road network.

Final dataset: Results were consolidated in file `orders_sao_jose_dos_campos_with_routes.csv`, containing columns `Route Distance (km)`, `Route Duration (s)`, and `DSR`, ready to feed subsequent modal classification and infrastructure sizing analyses.

Stage 7: Droneport Location via K-means

The processing executed in *notebook* 04 — `04_analysis.ipynb` — implemented the fifth stage of the operational pipeline (Section 3.3), using the K-means algorithm to determine optimal locations of *droneports* in the São José dos Campos region.

Data preparation: The set of unique restaurants was extracted from the dataset through aggregation by `Hashed Store ID`, resulting in 56 distinct restaurants with coordinates `Store Lat` and `Store Lng`. The analysis revealed high demand concentration: while the mean orders per restaurant was 416.7, a single establishment concentrated 6,205 orders (26.6% of total), evidencing significant asymmetry in demand distribution.

Replication approach with jitter: To incorporate demand intensity into the clustering process, a proportional replication strategy was adopted: each restaurant generated multiple points corresponding to its order volume, limited to 100 replicas per establishment for computational optimization. To each replicated point, Gaussian noise (jitter) with

standard deviation of 0.0005 degrees (approximately 50 m) was applied, avoiding exact overlap and allowing the algorithm to capture spatial variations in demand. This approach resulted in 2763 geographic points representing the weighted distribution of order demand.

Determination of optimal number of clusters: The elbow method was applied to determine the optimal number of *droneports*, as foreseen in the methodology. Analysis of inertia $J(k)$ for values of k ranging from 1 to 20 revealed an inflection point (elbow) between $k = 3$ and $k = 4$, indicating that the marginal gain in inertia reduction decreases significantly after $k = 4$. Figure 4.4 illustrates this analysis, showing the inertia curve with the elbow point clearly identified at $K = 4$. Based on this analysis, $K = 4$ was chosen as the optimal number of *droneports*.



FIGURE 4.4 – Elbow method applied to Approach A for determining the optimal number of clusters. Analysis of inertia $J(k)$ as a function of the number of clusters k reveals an inflection point (elbow) at $K = 4$, indicating that adding more clusters beyond this point offers diminishing returns in inertia reduction.

K-means application: The K-means algorithm was applied to the 2763 replicated points with $K = 4$, using 10 random initializations and `random_state=42` for reproducibility. The final inertia obtained was 1.52, generating 4 centroids that represent the geographic locations of the *droneports*.

Coverage analysis: To validate positioning effectiveness, coverage analysis was performed considering the 5 km range restriction for drone deliveries. For each customer–*droneport* pair, the geodesic distance (Haversine) was calculated and customers within the coverage radius were identified. Results indicated that the 4 positioned *droneports* can cover 19 947 orders out of a total of 23 335, representing 85.48% demand coverage. The 3388 uncovered orders (14.52%) are located beyond the 5 km radius of any *droneport*, potentially

served by traditional motorcycle delivery or requiring additional infrastructure positioning. Figure 4.5 illustrates the relationship between the number of *droneports* and achieved order coverage, evidencing diminishing returns: while increasing from 3 to 4 *droneports* provides significant coverage gain, adding more infrastructure beyond $K=4$ offers progressively smaller marginal gains, validating the choice of $K=4$ as the optimal compromise point between coverage and infrastructure investment. Figure 4.6 visually presents the final location of the 4 *droneports* (marked as centroids) distributed in the São José dos Campos region, along with aggregated demand points, evidencing positioning effectiveness to maximize spatial coverage.

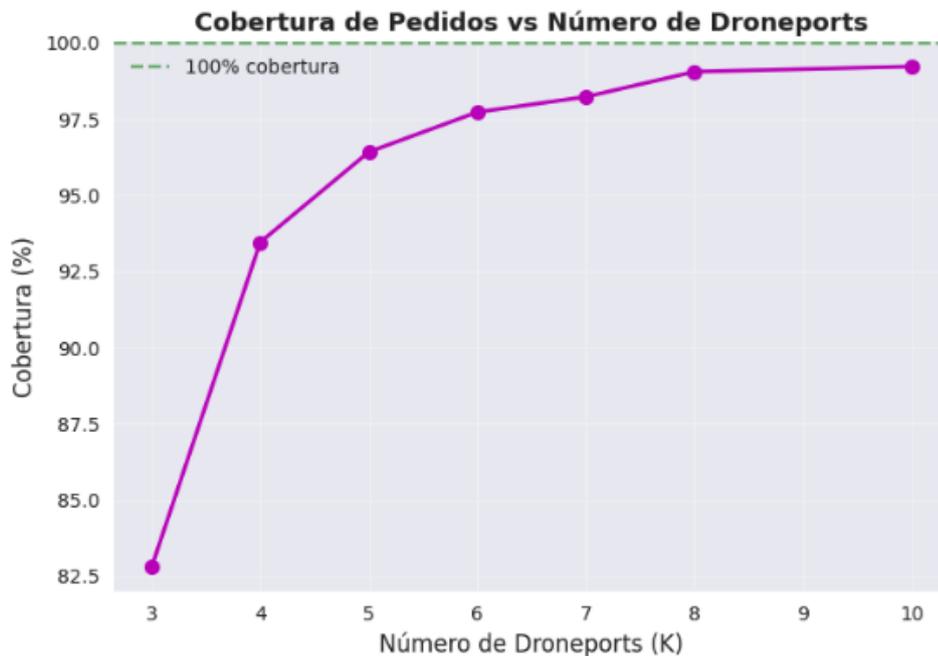


FIGURE 4.5 – Order coverage as a function of the number of *droneports* (Approach A). The curve shows diminishing returns as the number of *droneports* increases, with significant marginal gains until $K=4$, after which adding more infrastructure offers progressively smaller improvements in coverage. The dashed horizontal line indicates 100% coverage as reference.

Final results: Coordinates of the 4 centroids were saved in `centroids_A.csv` and used as the basis for fleet sizing and subsequent analyses of the $M/M/c$ queueing model described in Subsection 2.1.4. *Droneport* location was determined exclusively by spatial and demand criteria, as foreseen in the methodology; in a practical implementation, these points would be refined considering eigenvector centrality in the road network and availability of adequate infrastructure, as mentioned in the fifth stage of the operational pipeline.

Stage 8: Construction of Flyable Matrix for Meteorological Conditions

The processing executed in *notebook* 05 — `05_meteorological_analysis.ipynb` — implemented the sixth stage of the operational pipeline (Section 3.3), constructing binary

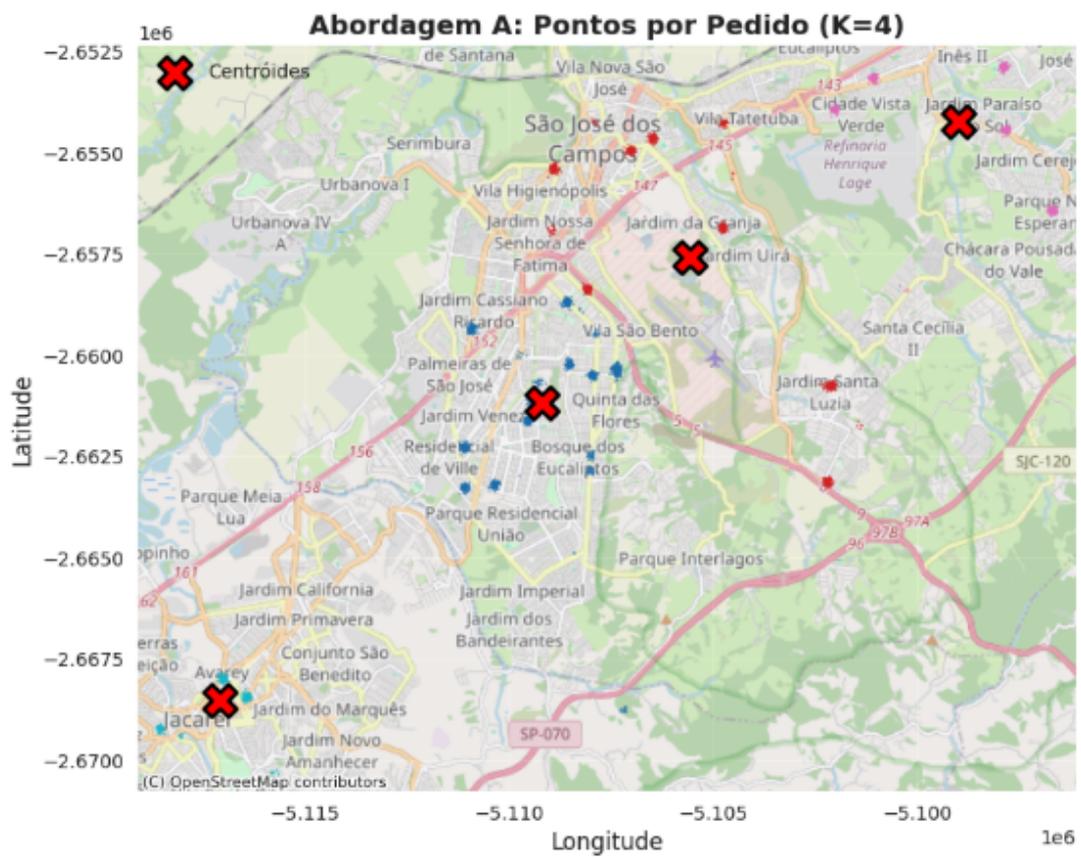


FIGURE 4.6 – Final location of the 4 *droneports* determined by the K-means algorithm in the São José dos Campos region. The centroids (marked with red 'X') represent optimal positions of the *droneports*, while colored points represent the distribution of order demand grouped by cluster. The map evidences the strategic distribution of *droneports* to maximize demand coverage within the 5 km radius.

flight viability matrices based on historical meteorological conditions, as foreseen in the methodology (Section 3.1.3).

Reference point definition: The meteorological reference point was established as the main centroid calculated from the 4 *droneports* of Approach A, resulting in coordinates latitude -23,234306 and longitude -45,883495. This choice ensures that analyzed meteorological conditions are representative of the *droneport* operation region. Figure 4.7 illustrates the location of the 4 *droneports* and the main centroid used as reference point for meteorological data collection.

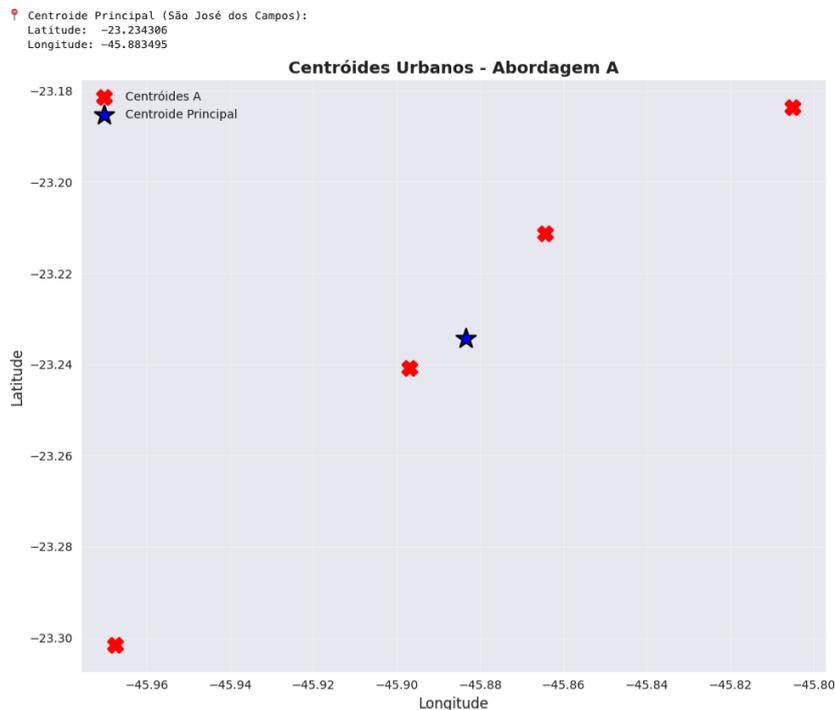


FIGURE 4.7 – Location of the 4 *droneports* (marked with red 'X') and the main centroid (marked with blue star) used as reference point for meteorological data collection in the São José dos Campos region.

Temporal data adjustment: The period of available orders comprises June to August 2025; however, meteorological data (ERA5 and INMET) present temporal lag that prevents access to future data. As foreseen in the methodology, an equivalent historical period was adopted — June to August 2024 — maintaining seasonal coherence and allowing analysis with real meteorological data.

ERA5 data acquisition and processing: Wind data were obtained from ERA5 reanalysis of the Copernicus Climate Data Store, consisting of u and v wind components at 10 m height with hourly resolution. NetCDF files were processed to extract wind speed and direction at the reference point through spatial interpolation. 2183 hourly records were processed for the three-month period, showing mean wind speed of 6.79 km/h, median of 6.65 km/h, and maximum of 21.77 km/h.

INMET data acquisition and processing: Wind and rain data were obtained from

INMET Taubaté station, 42.8 km from the urban centroid of São José dos Campos, as it is the nearest station with continuous hourly records. Due to INMET API limitations for historical data, precipitation data were obtained through manual file export from the institutional website (`rain_data_real.csv`). 1288 hourly records were processed, revealing that most hours (median and 75th percentile equal to zero) showed absence of precipitation, with maximum value of 3.8 mm/h and arithmetic mean of 0.012 mm/h. Figure 4.8 presents examples of meteorological data (wind and precipitation) for the first 7 days of analysis, illustrating condition variability and predominance of wind speeds below operational limits.



FIGURE 4.8 – Meteorological data (wind speed and direction, and precipitation) for the first 7 days of analysis (June 2024). Data illustrate favorable conditions for operation, with wind speeds consistently below the 25 km/h limit and absence of significant precipitation in the period.

Flight scenario definition: Three scenarios were defined as foreseen in the methodology, each establishing distinct limits for wind speed and precipitation. The conservative scenario establishes maximum wind of 15 km/h and maximum precipitation of 0 mm/h, representing very restrictive conditions for pessimistic analysis. The moderate scenario establishes maximum wind of 20 km/h and maximum precipitation of 1 mm/h, representing intermediate conditions. The aggressive scenario establishes maximum wind of 25 km/h and maximum precipitation of 2.5 mm/h, corresponding to the maximum limit configured in the system.

Limits were based on commercial drone specifications, scientific literature on unmanned aerial operations, and civil aviation regulations (ANAC, FAA, EASA), consid-

ering stabilization capacity, energy consumption, drift risks, and protection of electronic components.

Flyable matrix construction: For each scenario, a binary matrix `flyable_day`[day, hour] was constructed where each entry indicates flight viability (1=viaible, 0=not viable) based on simultaneous verification of both meteorological criteria. Results indicated high operational availability even in the most restrictive scenario, as detailed in Table 4.2. The high flyability rate (97.7%) results from typically favorable climatic conditions in the São José dos Campos region, and not from permissive parameters. Figures 4.9, 4.10, and 4.11 visually illustrate flight viability matrices for each scenario, showing temporal distribution of operational conditions.

TABLE 4.2 – Comparison of flight viability between meteorological scenarios.

Scenario	Total Hours	Viable Hours	Non-Viable Hours	% Viable
Conservative	2197	2146	51	97,68%
Moderate	2197	2192	5	99,77%
Aggressive	2197	2196	1	99,95%



FIGURE 4.9 – Flight viability matrix (flyable) for the conservative scenario, showing the first 30 days of analysis. Each cell represents a specific hour of a day, with dark green indicating viable flight (1) and red indicating restriction (0). The conservative scenario presents restriction periods concentrated mainly at the end of June, especially in afternoon hours.



FIGURE 4.10 – Flight viability matrix (flyable) for the moderate scenario, showing the first 30 days of analysis. Compared to the conservative scenario, it presents almost continuous availability, with only some isolated restriction periods.

Analysis of restriction reasons revealed that wind speed was the predominant limiting factor — in the conservative scenario, 51 hours were non-viable, most of which resulted from winds above 15 km/h. Precipitation had minimal impact, being responsible for only



FIGURE 4.11 – Flight viability matrix (flyable) for the aggressive scenario, showing the first 30 days of analysis. It presents almost total availability, with only a minimal restriction period at the beginning of the first day.

a small fraction of restrictions, consistent with the analyzed period (winter in the south-east region, characterized by lower precipitation). Figure 4.12 details specific restriction reasons for each scenario, evidencing wind predominance in the conservative scenario and progressive reduction of restrictions in less restrictive scenarios.



FIGURE 4.12 – Flight restriction reasons by meteorological scenario, detailing contribution of wind, precipitation, or both. The conservative scenario presents the highest number of restrictions, predominantly due to wind speeds above the limit, while moderate and aggressive scenarios present reduced number of restrictions.

Conservative scenario choice: As mentioned in the methodology, the conservative scenario was chosen for subsequent analyses, ensuring a pessimistic approach that increases result robustness and provides additional safety margin for practical operations.

Generated files: The three matrices were saved in `flyable_matrix_conservador.csv`, `flyable_matrix_moderado.csv`, and `flyable_matrix_agressivo.csv`, each containing columns `datetime`, `date`, `hour`, `wind_speed_kmh`, `precipitation_mm`, and `flyable`. These matrices feed the seventh stage of the operational pipeline, where each order is

classified as *Drone*, *Mixed*, or *Motorcycle* after verification of meteorological conditions at the time of request.

Stage 9: Centroid Analysis and Restaurant Association to Droneports

The processing executed in *notebook* 06 — `06_centroids_analysis.ipynb` — implemented strategic association of restaurants to *droneports* and calculation of operational distances necessary for modal classification and fleet sizing, connecting results from previous stages of the operational pipeline.

Centroid loading and validation: The notebook started with loading the 4 chosen centroids from Approach A, saved in `centroids_chosen.csv`. *Droneport* coordinates were validated: Droneport 1 located at latitude -23,240856 and longitude -45,896882; Droneport 2 at latitude -23,206760 and longitude -45,859530; Droneport 3 at latitude -23,183598 and longitude -45,805135; and Droneport 4 at latitude -23,301550 and longitude -45,967545.

Statistical analysis of coordinates revealed a spatial distribution representative of the urban region of São José dos Campos, with standard deviation of 0,051283 degrees in latitude and 0,068196 degrees in longitude, indicating adequate dispersion for demand coverage.

Droneport distance analysis: To evaluate infrastructure, the geodesic distance matrix between all *droneport* pairs was calculated using the Haversine formula. Results indicated minimum distance of 5.38 km between Droneport 1 and Droneport 2, maximum distance of 21.15 km between Droneport 3 and Droneport 4, mean distance of 11.52 km, median of 10.61 km, and standard deviation of 5.42 km.

The minimum distance of 5.38 km ensures that *droneports* are not excessively close, avoiding overlap of influence zones, while the maximum distance of 21.15 km indicates adequate dispersion for broad regional coverage. Table 4.3 presents the complete distance matrix between *droneports*.

TABLE 4.3 – Matrix of geodesic distances between *droneports* (in kilometers).

	Droneport 1	Droneport 2	Droneport 3	Droneport 4
Droneport 1	0,00	5,38	11,33	9,88
Droneport 2	5,38	0,00	6,13	15,26
Droneport 3	11,33	6,13	0,00	21,15
Droneport 4	9,88	15,26	21,15	0,00

Restaurant association to droneports: A critical stage was determination of the influence zone of each *droneport*, associating each restaurant to the nearest *droneport*. Unlike

the simple geometric approach (Haversine distance), the **Google Routes API** was used to calculate real road distances between restaurants and *droneports*, providing a more accurate representation of time and terrestrial transport distance necessary to bring orders from restaurants to *droneports*.

The process involved extraction of 56 unique restaurants from the order dataset, identified by `Hashed Store ID`; calculation of road distance and travel time for each restaurant–*droneport* pair (totaling 224 API requests); identification of the nearest *droneport* for each restaurant based on shortest road distance; and parallel processing with 10 simultaneous *workers* to optimize execution time.

Association results: Distribution of restaurants by *droneport* revealed significant asymmetry: Droneport 1 was associated with 22 restaurants (39,3% of total), Droneport 2 with 13 restaurants (23,2% of total), Droneport 3 with 10 restaurants (17,9% of total), and Droneport 4 with 11 restaurants (19,6% of total).

Road distance statistics between restaurants and their associated *droneports* showed minimum distance of 0.42 km, maximum of 11.87 km, mean of 4.23 km, and standard deviation of 2.65 km.

The mean distance of 4.23 km indicates that, on average, restaurants are positioned at a moderate road distance from *droneports*, facilitating initial terrestrial transport of orders to takeoff points. Figure 4.13 visually illustrates influence zones, showing restaurants colored according to their associated *droneport* and evidencing spatial distribution of coverage zones. The visualization confirms that each *droneport* has a well-defined influence area, with concentration of restaurants nearby and distribution that reflects demand density in the São José dos Campos and Jacareí region.

Droneport–customer distance calculation: After restaurant association to *droneports*, geodesic distance (Haversine) between each *droneport* and the corresponding customer delivery address was calculated. This metric is fundamental to determine viability of direct drone deliveries, verifying if distance is within the 5 km operational radius.

For the 23 335 processed orders, statistics of *droneport*–customer distances were: minimum distance of 0.029 km, maximum of 14.332 km, mean of 3.515 km, median of 3.093 km, and standard deviation of 2.162 km.

Analysis by *droneport* revealed distribution variations: Droneport 1 concentrated 12 630 orders (54,1%) with mean distance of 3.563 km; Droneport 2 totaled 6708 orders (28,7%) with mean distance of 3.883 km; Droneport 3 served 1647 orders (7,1%) with mean distance of 2.517 km; and Droneport 4 served 2350 orders (10,1%) with mean distance of 2.912 km.

Demand concentration at Droneport 1 (more than half of orders) reflects presence

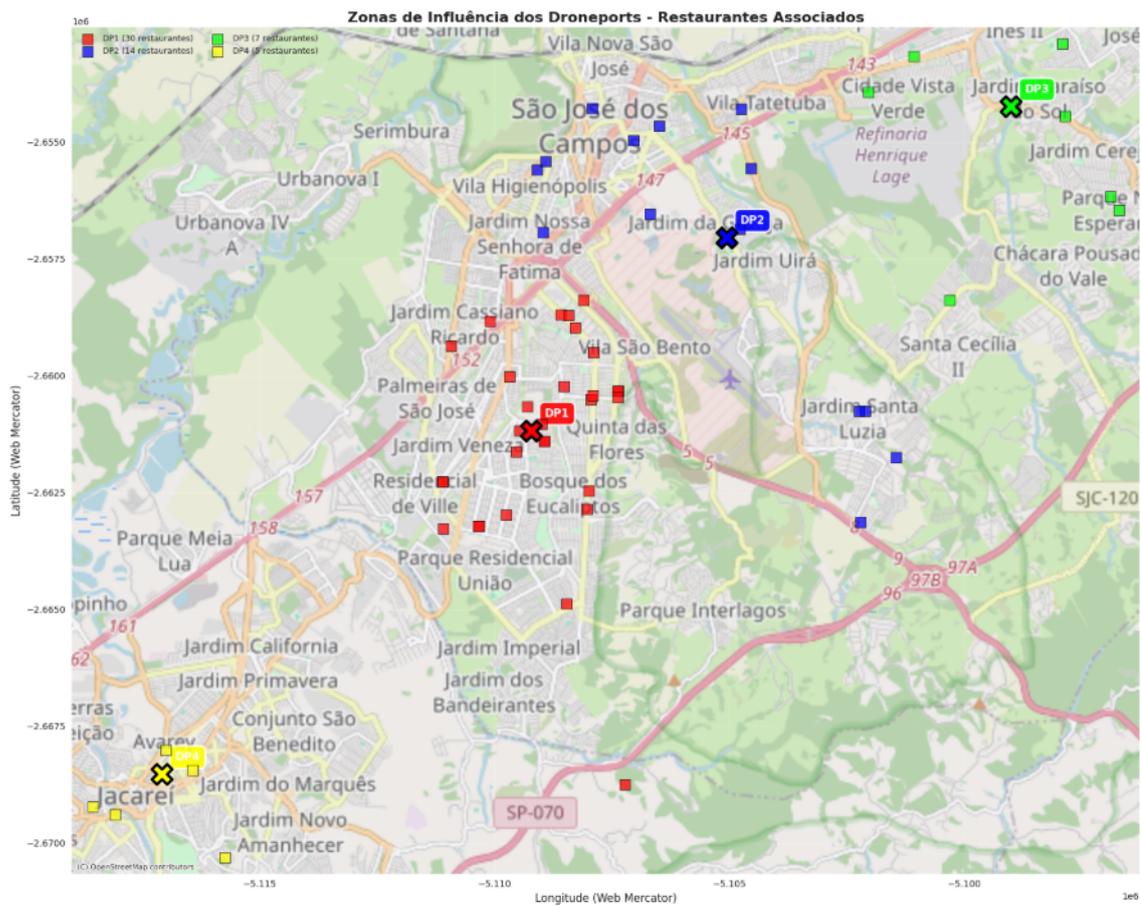


FIGURE 4.13 – Influence zones of *droneports* in the São José dos Campos region. Markers with 'X' represent the 4 *droneports* (DP1 in red, DP2 in blue, DP3 in green, DP4 in yellow), while colored squares represent restaurants associated with each *droneport* according to shortest road distance calculated by Google Routes API. Spatial distribution evidences natural partition of the region into coverage zones based on real road proximity, allowing logistical optimization of terrestrial transport to takeoff points.

of high-volume restaurants in this region, while Droneport 3 shows the shortest mean distance to customers (2.517 km), indicating potential operational efficiency for fast deliveries.

Integration with flyable matrix: The notebook concluded with integration of the meteorological viability matrix (`flyable_matrix_conservador.csv`) to orders, associating each order to flight status (`flyable`) corresponding to its request date and time. The join was performed through a composite key of date and hour, adjusting year from 2024 (meteorological data period) to 2025 (order period) to maintain temporal coherence.

Results indicated high operational availability: 23 290 out of 23 335 orders (99,8%) had flight status found; among these, 22 765 (97,7%) were classified as viable (`flyable = true`), while 525 (2,3%) were classified as restricted (`flyable = false`).

The 97,7% rate of orders in viable meteorological conditions validates the conservative scenario choice and demonstrates that most orders could be served by drone from a climatic perspective, disregarding other operational restrictions (autonomy, time, cost).

Generated files: Processing generated three main files: `restaurants_droneport_assignment.csv`, containing restaurant–*droneport* association table with road distances and transport times; `orders_with_complete_routes.csv`, containing complete order dataset with associated *droneport* information, geodesic distances, and meteorological status; and `distances_between_droneports.csv`, containing distance matrix between *droneports* for infrastructure analyses.

These files provide the consolidated database necessary for subsequent stages of the operational pipeline, including modal order classification (stage 7) and fleet sizing via M/M/c model (stage 8), as described in following sections of this chapter.

Stage 10: Modal Order Classification — Comparison between Delivery Modes

The processing executed in *notebook 07 — 07_delivery_mode_comparison.ipynb* — implemented the seventh stage of the operational pipeline (Section 3.3), performing modal classification of each order through systematic comparison between two distinct delivery modes. This stage is fundamental to determine the optimal logistical strategy for each order, balancing delivery time, operational viability, and technical restrictions.

Context and objective: Given strategically positioned *droneport* infrastructure and validated meteorological conditions, each order can be served by one of two operational strategies. The first strategy consists of 100% motorcycle delivery, involving direct transport from restaurant to customer using traditional motorcycle delivery, utilizing the optimized road route calculated by Google Routes API. The second strategy consists of mixed delivery, involving terrestrial transport from restaurant to the nearest *droneport*

(via motorcycle), followed by takeoff, drone flight to customer, and final delivery.

The choice between these strategies must minimize total delivery time, respecting operational and technical restrictions.

Decision model and methodological simplifications: This stage represents a critical and sensitive moment of the simulation, where the decision on delivery mode is made for each order. The problem can be visualized through a triangular diagram formed by three vertices: the *droneport*, the restaurant, and the customer. From this geometric configuration, the optimal logistical strategy must be determined. Figure 4.14 conceptually illustrates the two possible delivery modes: direct mode (restaurant \rightarrow customer via motoboy) and mixed mode (restaurant \rightarrow *droneport* via motoboy, followed by *droneport* \rightarrow customer via drone).

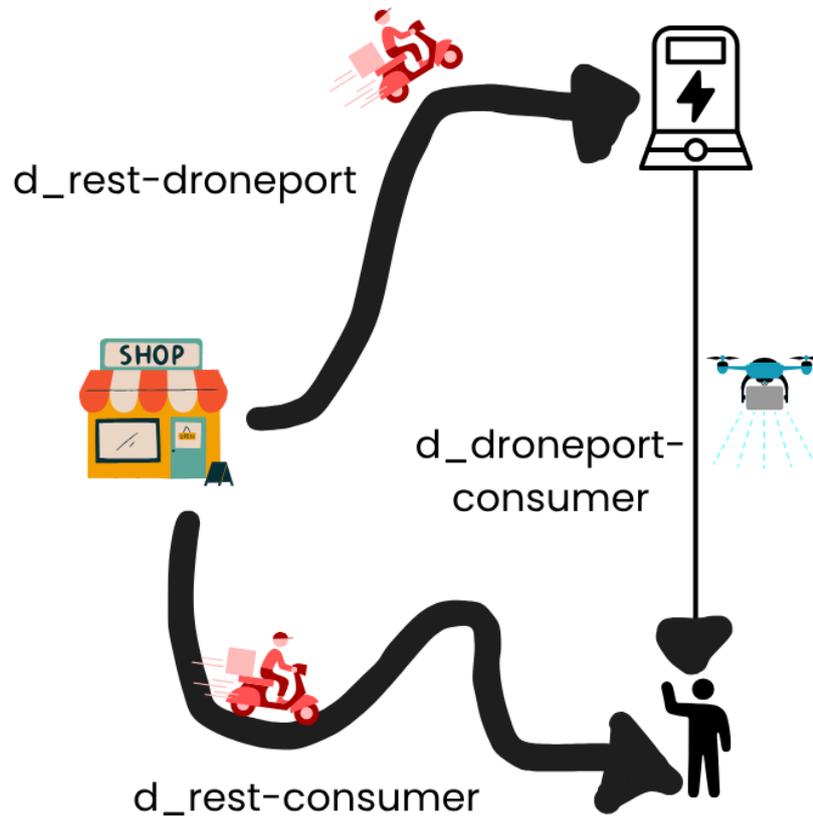


FIGURE 4.14 – Conceptual triangular diagram of delivery modes. The diagram illustrates the two possible operational strategies for an order: (1) direct mode, where the motoboy transports the order directly from restaurant to customer, covering distance $d_{\text{rest-consumer}}$; (2) mixed mode, where the motoboy transports the order from restaurant to *droneport* (distance $d_{\text{rest-droneport}}$), and then the drone completes delivery from *droneport* to customer (distance $d_{\text{droneport-consumer}}$). The decision between modes depends on comparison of total times of each strategy, considering meteorological and operational range restrictions.

Obvious decision cases are those where technical restrictions prevent use of mixed

mode. When meteorological conditions are not viable (`flyable = false`), flight is impossible and choice automatically falls on motorcycle. When distance between *droneport* and customer exceeds the 5 km operational limit, the drone does not have sufficient range and again motorcycle is the only viable option.

The complex case occurs when both strategies are technically viable: at this moment, decision requires time comparison, considering multiple operational variables that could impact choice in a real scenario.

Assumed simplifications: To enable analysis and make the model computationally tractable, important methodological simplifications were assumed that need to be explicit.

The first simplification refers to independence between orders: each order is treated as an independent entity, without temporal or spatial relation with subsequent or previous orders. In a real operation, there would be a sequence of orders arriving over time, creating operational dependencies (for example, a drone busy with a previous order would not be available for an immediate new order). This simplification allows treating each order in isolation, but does not capture the real temporal dynamics of the operation.

The second simplification refers to instant resource availability: it is assumed that both drone and motoboy are always available at the moment of order request. In practice, both resources may be busy serving other orders, and availability would be a stochastic variable dependent on the current system state. This simplification eliminates capacity restrictions that could alter modal decision in high-demand scenarios.

The third simplification refers to unit capacity of the motoboy: it is assumed that each motoboy can transport only one order at a time, both for direct deliveries to customer and for transports to *droneport*. In reality, especially in the case of transport to *droneport*, a motoboy could optimize its route collecting multiple orders from nearby restaurants and delivering them simultaneously to the *droneport*, which would then distribute to final customers via drones. This consolidation strategy could significantly reduce terrestrial transport times and make mixed mode more competitive. The adopted simplification represents a pessimistic scenario, where there is no order consolidation, simplifying analysis but potentially underestimating mixed mode efficiency.

These simplifications are intentional and conservative: by assuming less favorable conditions for mixed mode (no consolidation, no complex temporal dependencies), obtained results can be interpreted as a lower estimate of the hybrid system's operational potential. In a practical implementation, with route optimization, order consolidation, and dynamic resource management, mixed mode could present superior performance to that observed in this analysis.

Operational parameters: Time calculations for mixed mode were based on technical parameters established according to commercial drone specifications and scientific liter-

ature. Drone cruise speed was defined at 36 km/h (10 m/s), typical operation speed in urban environment for small delivery drones. *Droneport* operation time was established at 120 s (2 min), including landing time, cargo transfer, safety verification, and takeoff. Maximum *droneport*–customer distance was defined at 5 km, operational range limit of the drone considered in spatial viability analysis.

Comparison methodology: Modal classification was performed through a sequential decision algorithm applied to each order, according to the following rules:

Rule 1 — Meteorological verification: If `field_flyable` equals 0 (non-viable meteorological conditions), the order is automatically classified as *motorcycle*, regardless of other considerations. This restriction is imposed by operational safety and regulatory compliance.

Rule 2 — Range verification: If geodesic distance between associated *droneport* and customer is greater than 5 km, the order is classified as *motorcycle*, as it is beyond the drone’s operational range.

Rule 3 — Time comparison: For orders that passed the first two verifications, the decision criterion is simple and direct: calculate both delivery times and choose the mode that presents shorter total time.

Motorcycle time corresponds to **Route Duration (s)** already calculated by Google Routes API, representing direct road travel time from restaurant to customer.

Mixed mode time is calculated through the formula:

$$t_{\text{mixed}} = t_{\text{moto_droneport}} + t_{\text{operation}} + \frac{d_{\text{droneport_customer}}}{v_{\text{cruise}}} \quad (4.4)$$

where $t_{\text{moto_droneport}}$ is terrestrial transport time from restaurant to *droneport* (obtained from `time_to_droneport_s`), $t_{\text{operation}} = 120$ s is operation time at *droneport*, $d_{\text{droneport_customer}}$ is geodesic distance between *droneport* and customer, and $v_{\text{cruise}} = 36$ km/h is drone cruise speed.

The chosen mode is the one that presents shorter total delivery time. This direct time comparison is possible thanks to assumed methodological simplifications (instant resource availability, independence between orders), which allow treating each order as a static optimization problem.

General classification results: Application of the classification algorithm to the 23 335 processed orders resulted in the following distribution:

Analysis of decision reasons reveals limiting factors for mixed mode adoption:

Predominance of motorcycle mode (88,78% of orders) indicates that, for most cases in the analyzed region, direct delivery via motorcycle presents shorter time than mixed

TABLE 4.4 – Distribution of chosen delivery modes and respective decision reasons.

Chosen Mode	Number of Orders	Percentage
Motorcycle	20 717	88,78%
Mixed	2618	11,22%
Total	23 335	100,00%

TABLE 4.5 – Distribution of orders by decision reason in modal classification.

Decision Reason	Number of Orders	Percentage
Motorcycle time shorter	17 191	73,67%
Maximum distance exceeded	2956	12,67%
Mixed time shorter	2618	11,22%
Non-viable meteorological conditions	570	2,44%
Total	23 335	100,00%

strategy, even when this is viable from meteorological and range perspectives. Figure 4.15 visually illustrates distribution of chosen modes, decision reasons, comparison of delivery times, and distribution of time savings obtained.

Delivery time analysis: Statistical comparison of delivery times by mode reveals important characteristics of operational efficiency.

For orders classified as *motorcycle*, mean time is 525.72 s (8.76 min), median time is 480.00 s (8.00 min), and range varies from a few minutes for nearby deliveries to approximately 20 min for more distant routes.

For orders classified as *mixed*, mean time is 693.95 s (11.57 min) and median time is 711.89 s (11.86 min).

Observation that mean time of mixed mode (when chosen) is superior to mean time of motorcycle mode may seem contradictory; however, this global metric does not capture the individual comparison performed by the algorithm. Mixed mode was chosen only when it presented shorter time than motorcycle for that specific order, but its absolute times tend to be greater due to overhead of transport to *droneport* and operation.

Detailed analysis of mixed orders: To adequately understand the operational gain of mixed mode, it is necessary to analyze specifically the 2618 orders where this strategy was chosen. In these cases, comparing time that would be spent by pure motorcycle versus effective time of mixed mode, pure motorcycle time (hypothetical) shows mean of 777.95 s (12.97 min) and median of 775.50 s (12.93 min); mixed time (real) shows mean of 690.56 s (11.51 min) and median of 708.66 s (11.81 min); time savings shows mean of 87.39 s (1.46 min) and median of 63.13 s (1.05 min); and percentage savings shows mean of 11,41% and median of 8,24%.

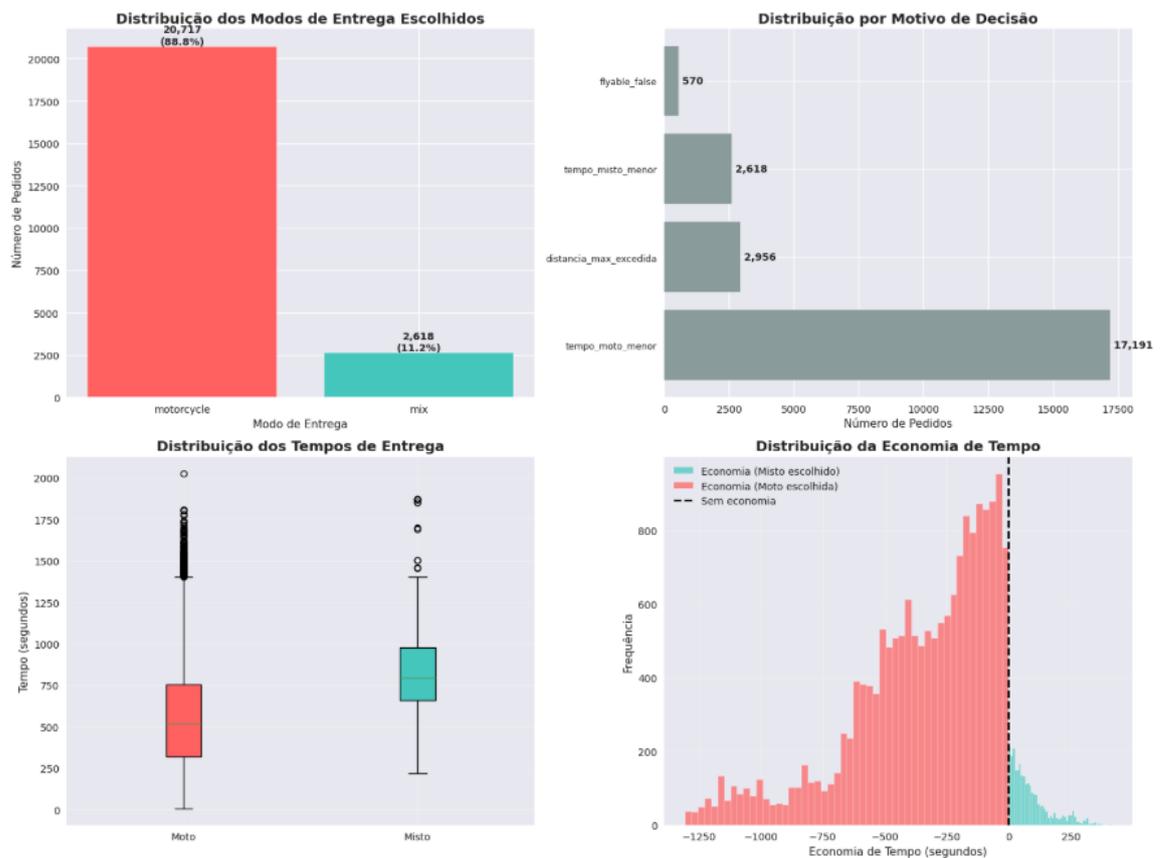


FIGURE 4.15 – Comparative analysis of delivery modes: (a) distribution of chosen modes, showing predominance of motorcycle mode (88,78%); (b) distribution by decision reason, evidencing that shorter motorcycle time is the main factor; (c) comparison of delivery time distributions through boxplots, indicating that motorcycle mode presents shorter median times; (d) time savings distribution, showing that when mixed mode is chosen, there is positive savings relative to motorcycle mode.

Percentile analysis of time savings reveals an asymmetric distribution: 10th percentile of 10.25 s (0.17 min) with minimum savings of 1,49%, 25th percentile of 26.97 s (0.45 min) with savings of 3,77%, 50th percentile (median) of 63.13 s (1.05 min) with savings of 8,24%, 75th percentile of 118.42 s (1.97 min) with savings of 14,73%, 90th percentile of 219.42 s (3.66 min) with savings of 26,34%, 95th percentile of 260.90 s (4.35 min) with savings of 39,22%, and 99th percentile of 329.20 s (5.49 min) with savings of 46,84%.

These results demonstrate that, when mixed mode is selected, it offers consistent time gains, with half of orders showing savings superior to 1 min and 25% of cases showing savings superior to approximately 2 min. Figure 4.16 presents detailed analysis of orders classified as mixed, comparing time that would be spent by pure motorcycle versus effective time of chosen mixed mode, visually evidencing time savings obtained.

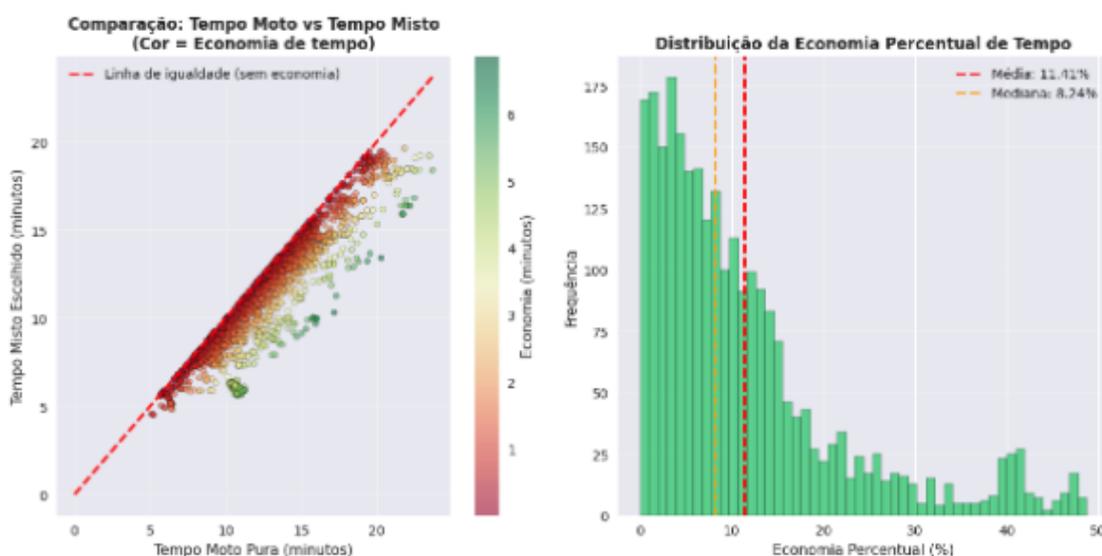


FIGURE 4.16 – Detailed analysis of orders with mixed delivery: (a) time comparison through scatter plot, where each point represents an order and color indicates time savings obtained; points below equality line indicate positive savings of mixed mode; (b) percentage time savings distribution, showing that most orders present savings between 5% and 15%, with mean of 11,41% and median of 8,24%. Distribution is right-skewed, indicating that some orders present significantly larger percentage savings.

Distance analysis: Distribution of chosen modes as a function of distance between *droneport* and customer reveals important patterns about operational viability:

TABLE 4.6 – Distribution of delivery modes by *droneport*–customer distance range.

Distance Range	Mixed	Motorcycle	Total	% Mixed
0–2 km	612	6023	6635	9,22%
2–4 km	1332	6660	7992	16,67%
4–6 km	674	5010	5684	11,86%
6–10 km	0	2792	2792	0,00%
>10 km	0	232	232	0,00%
Total	2618	20 717	23 335	11,22%

It is observed that the 2–4 km distance range presents the highest proportion of orders served by mixed mode (16,67%), indicating that this is the zone of highest relative efficiency for the aerial mode. For very short distances (0–2 km), overhead of transport to *droneport* and operation makes mixed mode less competitive. For distances between 4–6 km, still within operational range, proportion of mixed mode decreases, possibly due to increased flight time that compensates less the gain relative to road routes. Figure 4.17 visually illustrates this percentage distribution by distance range, clearly evidencing that mixed mode is used exclusively for distances within the 5 km operational radius.

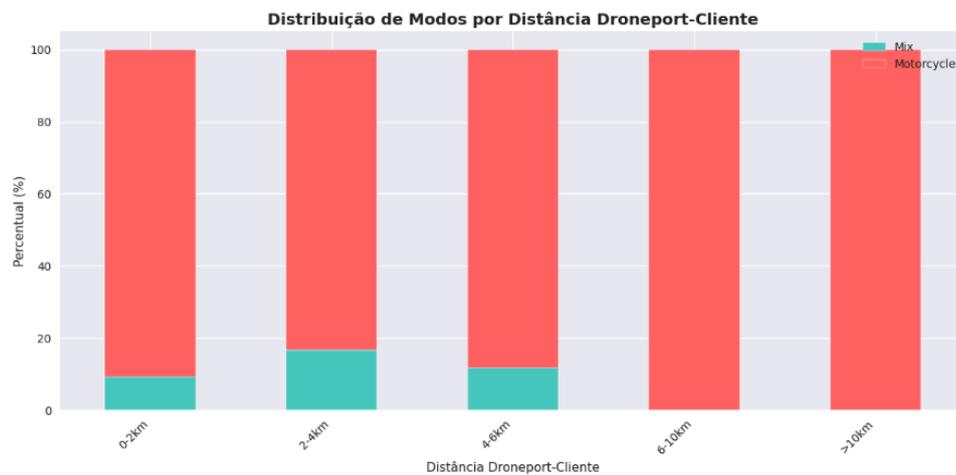


FIGURE 4.17 – Percentage distribution of delivery modes by distance range between *droneport* and customer. The stacked bar chart shows that mixed mode (in turquoise blue) presents highest proportion in the 2–4 km range (16,67%), decreasing for smaller distances (9,22% in 0–2 km) and larger distances (11,86% in 4–6 km). For distances superior to 6 km, mixed mode is not used due to operational range restriction, with 100% of orders served by motorcycle.

Analysis of road distances (**Route Distance**) by chosen mode reveals that mixed orders show mean Route Distance of 5.12 km and median of 4.87 km, while motorcycle orders show mean Route Distance of 4.74 km and median of 3.92 km.

Interestingly, orders classified as mixed show slightly larger road distances on average, which suggests that mixed mode is especially advantageous when the direct road route is longer or presents unfavorable traffic conditions.

Geographic distribution of orders by delivery mode, presented in Figure 4.18, reveals important spatial patterns about hybrid system operation. It is observed that orders served by mixed mode (in cyan) concentrate in well-defined clusters near *droneports*, while orders served by motorcycle (in red) show much broader and denser distribution throughout the region. This geographic visualization confirms that mixed mode operates essentially within *droneport* coverage radius, while motorcycle mode serves both covered and uncovered areas by aerial infrastructure.

Interpretation and operational implications: Modal classification analysis reveals that only 11,22% of orders are optimized through mixed mode, while 88,78% are served via

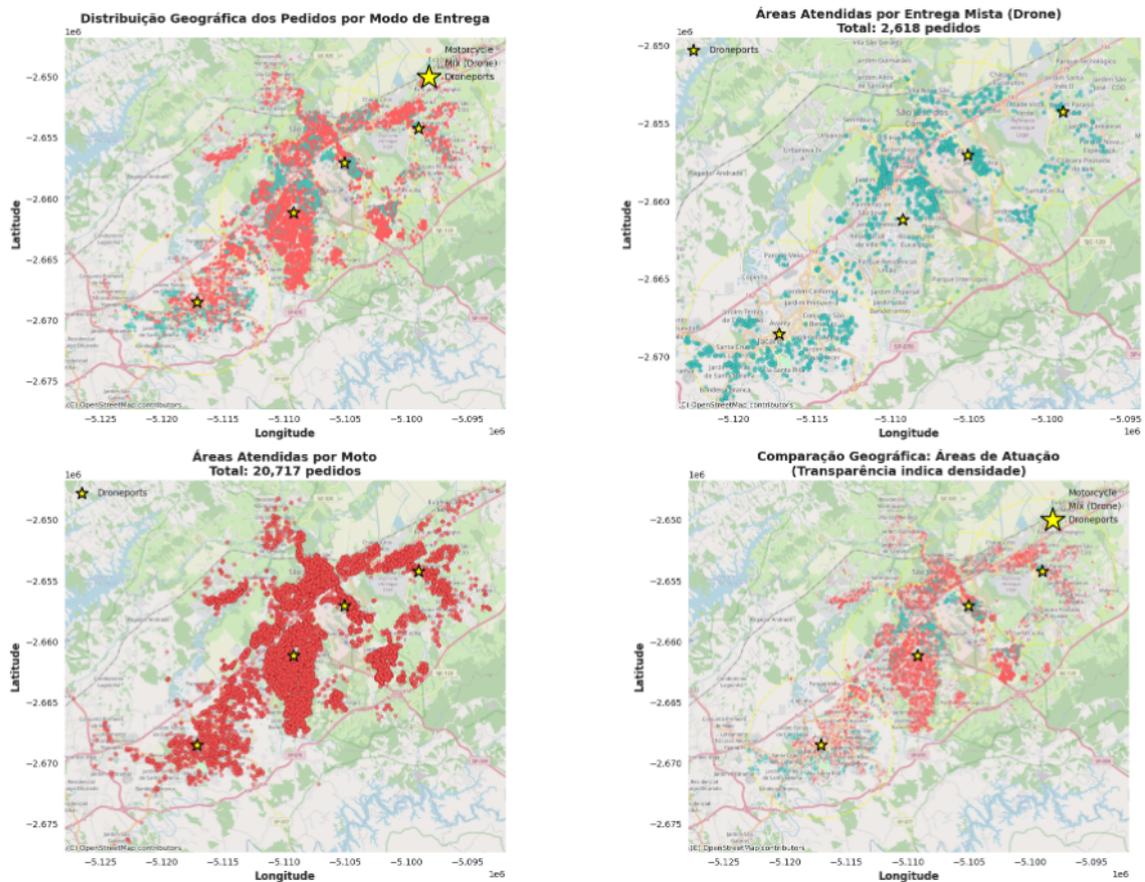


FIGURE 4.18 – Geographic distribution of orders by delivery mode in the São José dos Campos region: (a) general view showing all orders (red = motorcycle, cyan = mixed, yellow stars = *droneports*); (b) areas served exclusively by mixed delivery (2,618 orders), evidencing concentrated clusters near *droneports*; (c) areas served by motorcycle (20,717 orders), showing much broader coverage; (d) overlaid comparison with transparency indicating density, evidencing that motorcycle has dominant spatial distribution while mixed mode operates in specific zones.

traditional motorcycle. This result reflects various operational realities.

Regarding meteorological restrictions, approximately 2,4% of orders are excluded from mixed mode due to non-viable climatic conditions, validating the importance of meteorological analysis performed in Stage 8.

Regarding range restrictions, about 12,7% of orders are beyond the 5 km operational radius of *droneports*, being automatically directed to motorcycle. Orders with distance superior to the 5 km operational radius were automatically classified as motorcycle deliveries (motorcycle mode). This limitation could be mitigated through additional infrastructure positioning or expansion of drone operational range.

Regarding relative efficiency of terrestrial mode, for 73,7% of orders viable from meteorological and range perspectives, direct delivery via motorcycle still presents shorter time than mixed mode. This occurs because overhead of transport to *droneport* and operation time (2 min) frequently exceeds the gain provided by direct flight, especially for already optimized road routes.

Regarding optimization potential, orders where mixed mode was chosen concentrate at intermediate distances (2–6 km from *droneport*), where geometric gain of direct flight compensates operational overhead. Mean savings of 1.46 min in these cases demonstrates operational value, especially in contexts of rigid SLA.

Implications for sizing: The proportion of 11,22% of orders eligible for mixed mode is a critical input for drone fleet sizing, as described in the M/M/c queueing analysis of pipeline stage 8. This rate indicates that demand for drone services is substantial but not dominant, requiring a hybrid operation model where most orders continue to be served by traditional motorcycle delivery.

Generated files: Processing generated file `orders_delivery_mode_comparison.csv`, containing for each order the chosen mode (`modo_escolhido`: `motorcycle` or `mix`), decision reason (`motivo_decisao`), calculated times for both modes (`tempo_moto_s`, `tempo_misto_s`), and all information necessary for subsequent fleet sizing and economic evaluation analyses.

This dataset consolidates modal classification and serves as the basis for infrastructure sizing and economic-financial evaluation analyses described in following sections of this chapter.

Exploratory Analysis: Alternative Droneport Positioning Based on Demand

As part of exploratory analysis of notebook 07, an additional investigation was performed on alternative *droneport* positioning strategies aiming to increase penetration of deliveries via mixed mode. Analysis revealed that order distribution among restaurants is

highly asymmetric: while mean orders per restaurant is approximately 416,7, some establishments concentrate substantially larger volumes. Figure 4.19 illustrates this asymmetry through order volume of the 20 restaurants with highest demand, evidencing that the top 6 restaurants concentrate a disproportionately high proportion of total orders.

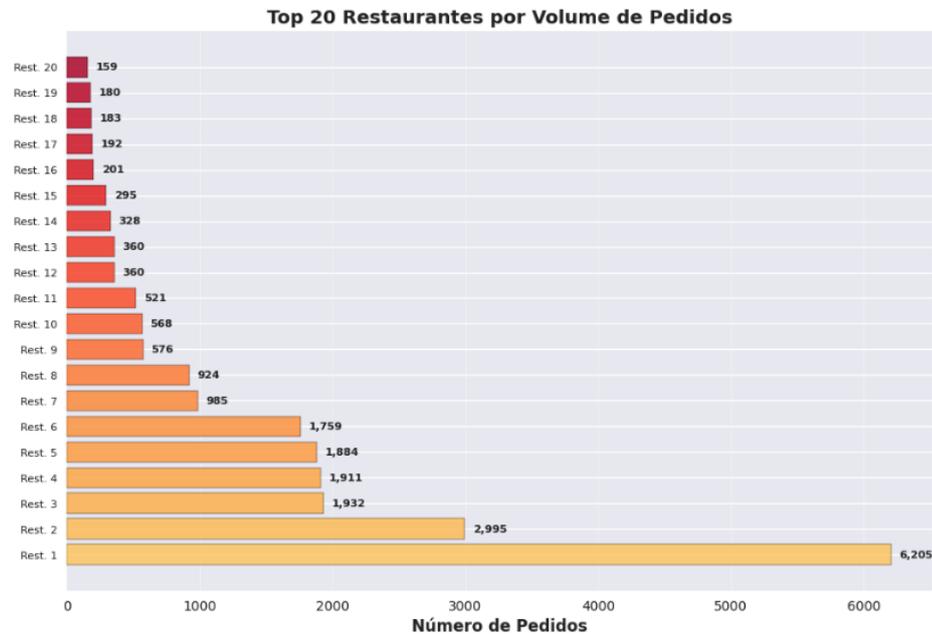


FIGURE 4.19 – Order volume distribution among the 20 restaurants with highest demand. The chart evidences high concentration of orders in top restaurants: the leading restaurant concentrates 6,205 orders, while following ones show significantly smaller volumes. The top 6 restaurants concentrate a substantial proportion of total demand, justifying *droneport* positioning strategies based on demand instead of uniform spatial distribution.

High-volume restaurant identification: Analysis of order distribution by restaurant identified that approximately 6 restaurants concentrate a significant proportion of total demand. This concentration suggests that positioning *droneports* directly at these high-volume restaurants could eliminate overhead of terrestrial transport to *droneport*, making mixed mode more competitive for these specific establishments.

Figure 4.20 presents geographic distribution of restaurants coded by order volume, revealing spatial location of demand centers and visually illustrating observed concentration. This visualization demonstrates that the mean does not necessarily represent the best practical solution: while centroids calculated via K-means position *droneports* at locations that optimize mean spatial distribution, high-volume restaurants may be displaced from these centroids, resulting in significant terrestrial transport times that reduce mixed mode competitiveness. Positioning *droneports* directly at highest-volume restaurants better captures the operational reality of concentrated demand.

Alternative analysis methodology: A demand-based positioning strategy was implemented, where *droneports* are positioned directly at restaurants with highest order volume, instead of using centroids calculated by K-means algorithm. The main advantage

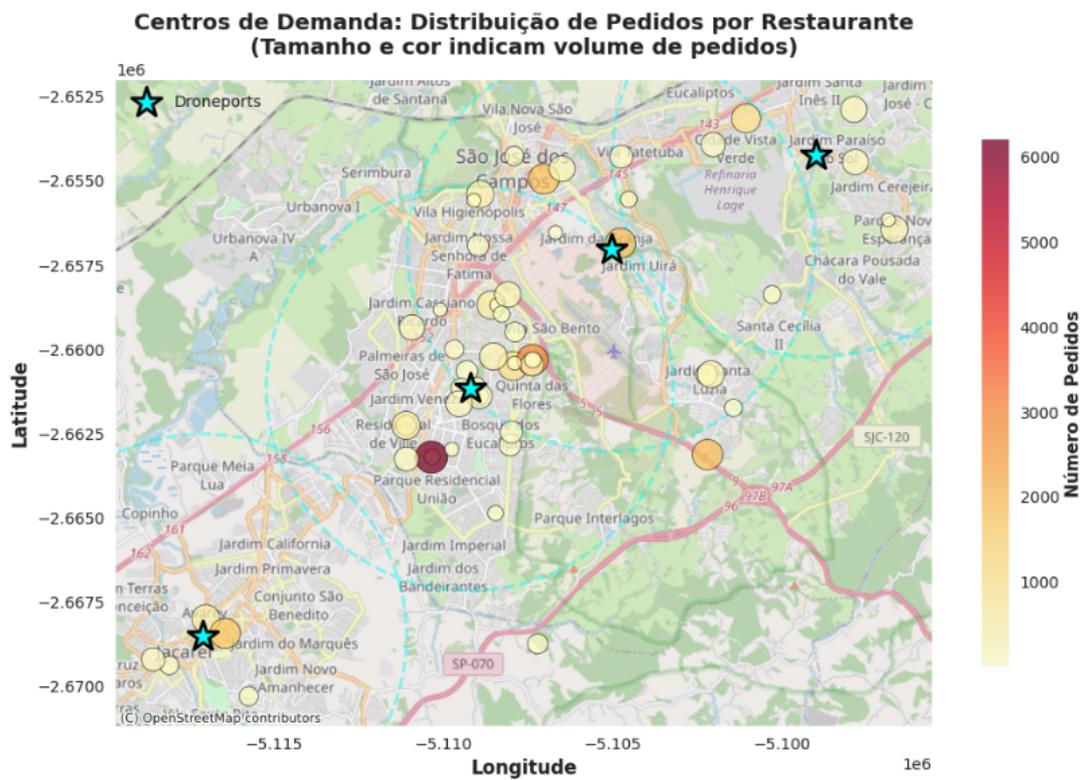


FIGURE 4.20 – Geographic distribution of restaurants coded by order volume. Size and color of circles indicate order volume of each restaurant, varying from light yellow (low volume) to dark red (high volume). *Droneports* positioned via K-means are marked with cyan stars. Visualization evidences that highest-volume restaurants (large red circles) do not necessarily coincide with calculated centroids, demonstrating that strategies based on spatial mean may not adequately capture real demand concentration.

of this approach is that, when the *droneport* is located at the restaurant itself, terrestrial transport time from restaurant to *droneport* ($t_{\text{moto_droneport}}$) becomes zero, eliminating this component from mixed time calculation.

For a restaurant with its own *droneport*, the mixed time formula simplifies to:

$$t_{\text{mixed}} = t_{\text{operation}} + \frac{d_{\text{droneport_customer}}}{v_{\text{cruise}}} \quad (4.5)$$

where $t_{\text{operation}} = 120$ s (2 minutes) and terrestrial transport time to *droneport* is eliminated.

Analyzed scenarios: Three alternative scenarios were evaluated, positioning *droneports* at top 4, top 5, and top 6 restaurants with highest order volume in the analyzed period. Each configuration was compared with the base configuration (4 *droneports* positioned via K-means) using the same modal classification algorithm described previously.

Exploratory analysis results: Alternative analysis results revealed significant gains in mixed mode penetration:

TABLE 4.7 – Comparison between *droneport* positioning configurations.

Configuration	Mixed Orders	Motorcycle Orders	% Mixed
Current (4 K-Means)	2618	20 717	11,22%
Alternative (4 Demand)	9351	13 984	40,07%
Alternative (5 Demand)	10 790	12 545	46,24%
Alternative (6 Demand)	12 501	10 834	53,57%
Total	23 335	23 335	—

Results demonstrate that demand-based strategy offers substantial gains. Configuration with 4 demand-based *droneports* showed increase of 28,85 percentage points in mixed mode penetration (from 11,22% to 40,07%), representing 6733 additional orders served via mixed mode. Configuration with 5 demand-based *droneports* showed increase of 35,02 percentage points (to 46,24%), with 8172 additional orders. Configuration with 6 demand-based *droneports* showed increase of 42,35 percentage points (to 53,57%), with 9883 additional orders.

Figure 4.21 presents geographic distribution of orders by delivery mode for configuration with 4 *droneports* positioned at top 4 restaurants with highest order volume. This configuration, which will be adopted as the basis for subsequent analyses of this work, demonstrates substantial increase in mixed mode penetration (40,07%) relative to original K-means configuration (11,22%), while maintaining a manageable number of *droneports* for operational and economic analyses.

Result interpretation: Significant gain observed in demand-based configurations can

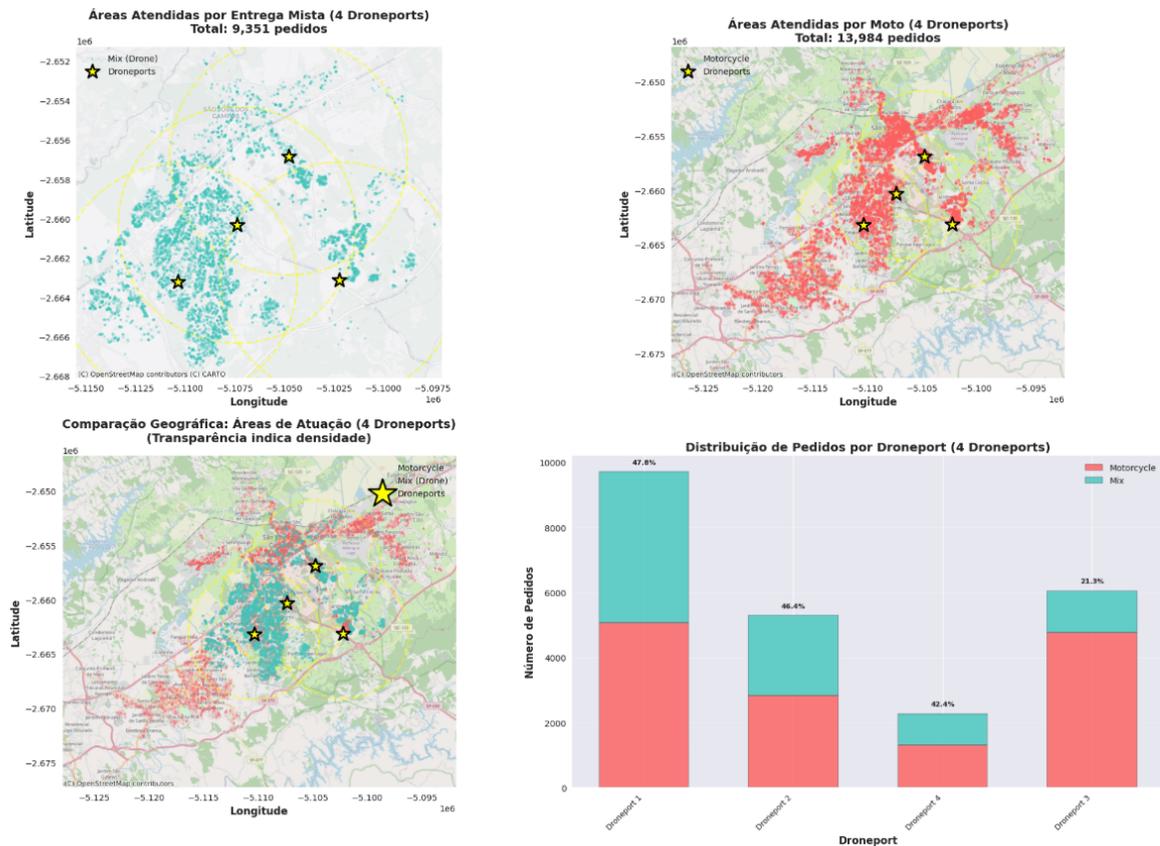


FIGURE 4.21 – Geographic distribution of orders by delivery mode for configuration with 4 *droneports* positioned at top 4 restaurants with highest order volume: (a) areas served by mixed delivery (9,351 orders), showing concentration near *droneports*; (b) areas served by motorcycle (13,984 orders), with broader distribution; (c) overlaid geographic comparison, evidencing that mixed mode operates mainly within *droneport* coverage radius; (d) order distribution by *droneport*, showing that each *droneport* serves both orders from the restaurant where it is located and from nearby restaurants in the region. This configuration presents 40,07% mixed mode penetration and will be used as the basis for subsequent analyses.

be attributed mainly to elimination of terrestrial transport overhead to *droneport* for high-volume restaurants. When the *droneport* is located at the restaurant itself, mixed time reduces essentially to operation time (2 min) plus flight time to customer, making this strategy competitive even for smaller distances that would not be viable in K-means configuration.

Additionally, *droneports* positioned at high-volume restaurants continue to serve orders from nearby restaurants in the region, maintaining operational flexibility and potentially increasing overall system efficiency.

Stage 11: Drone Fleet Sizing via M/M/c Model

The processing executed in *notebook* 08 — `08_analysis_top4_restaurants.ipynb` — implemented the eighth stage of the operational pipeline (Section 3.3), performing sizing of drone fleet necessary to meet the 20-minute SLA through the M/M/c queueing model. This stage is fundamental to determine operational capacity of each *droneport*, ensuring that the system can meet expected demand without compromising agreed delivery times.

Context and data preparation: Sizing was performed using the configuration chosen in previous exploratory analysis: 4 *droneports* positioned at top 4 restaurants with highest order volume. Only orders classified as *mixed* (9,351 orders, 40,07% of total) were considered for sizing, as these are the ones that arrive at the *droneport* queue to be served by drones. Orders classified as *motorcycle* are served directly via terrestrial transport and do not require drone resources.

Operational parameters of sizing: Calculations were based on established technical and operational parameters. SLA was defined at 1200 s (20 min), representing maximum total time in the system (queue waiting time + service time). Customer discharge time was established at 60 s (1 min), representing time necessary to deliver the order to the customer. Drone cruise speed was defined at 36 km/h, both for outbound and return flight. *Droneport* operation time was established at 120 s (2 min), including battery swap, verification, and preparation for next mission.

Complete cycle time calculation: For each mixed order, complete drone cycle time was calculated, which comprises four sequential components. The first component is outbound flight, corresponding to flight time from *droneport* to customer, calculated as $\frac{d_{\text{droneport_customer}}}{v_{\text{cruise}}}$. The second component is customer discharge, with fixed time of 60 s for order delivery. The third component is return flight, corresponding to flight time from customer back to *droneport*, equal to outbound flight. The fourth component is *droneport* operation, with fixed time of 120 s for next mission preparation.

Complete cycle time is given by:

$$t_{\text{cycle}} = 2 \cdot \frac{d_{\text{droneport_customer}}}{v_{\text{cruise}}} + t_{\text{discharge}} + t_{\text{operation}} \quad (4.6)$$

For the 9,351 mixed orders analyzed, complete cycle time showed mean of 782.06 s (13.03 min), median of 789.19 s (13.15 min), ranging from 195.44 s (3.26 min) to 1379.86 s (23.00 min).

Temporal demand distribution analysis: A critical stage of sizing was analysis of temporal distribution of mixed orders, aiming to identify demand patterns that directly impact capacity sizing. Analysis revealed that demand is not uniform over time, showing significant concentrations in specific periods.

Figure 4.22 presents distribution of mixed orders by hour of day and by day of week, revealing important patterns. Regarding distribution by hour of day, demand is very low during early morning and morning (0h–16h), starting significant growth from 17h. Peaks occur between 19h and 21h, with approximately 1,900 orders at 20h, followed by gradual decline until midnight. Regarding distribution by day of week, demand increases progressively throughout the week, with Monday showing lowest volume (approximately 550 orders) and Saturday showing highest volume (approximately 1,900 orders). Weekends (Friday, Saturday, and Sunday) concentrate substantially larger volumes than weekdays.



FIGURE 4.22 – Temporal distribution of mixed orders: (a) distribution by hour of day, evidencing demand concentration in evening period (17h–23h) with peak at 20h; (b) distribution by day of week, showing progressive increase throughout the week with larger volumes on weekends. Temporal analysis is fundamental to identify peak periods that require greater capacity and system idleness periods.

Combined analysis day of week \times hour of day is essential to adequately capture demand patterns, as different days of week present distinct hourly profiles. Figure 4.23 presents a detailed heatmap showing order distribution by all possible combinations of day of week and hour of day.

It is observed that periods of highest demand concentrate in hours 18h–22h, especially on weekends, with intensities superior to 350 orders per day+hour combination. On the other hand, early morning periods (0h–7h) and morning (8h–15h) show very low demand,

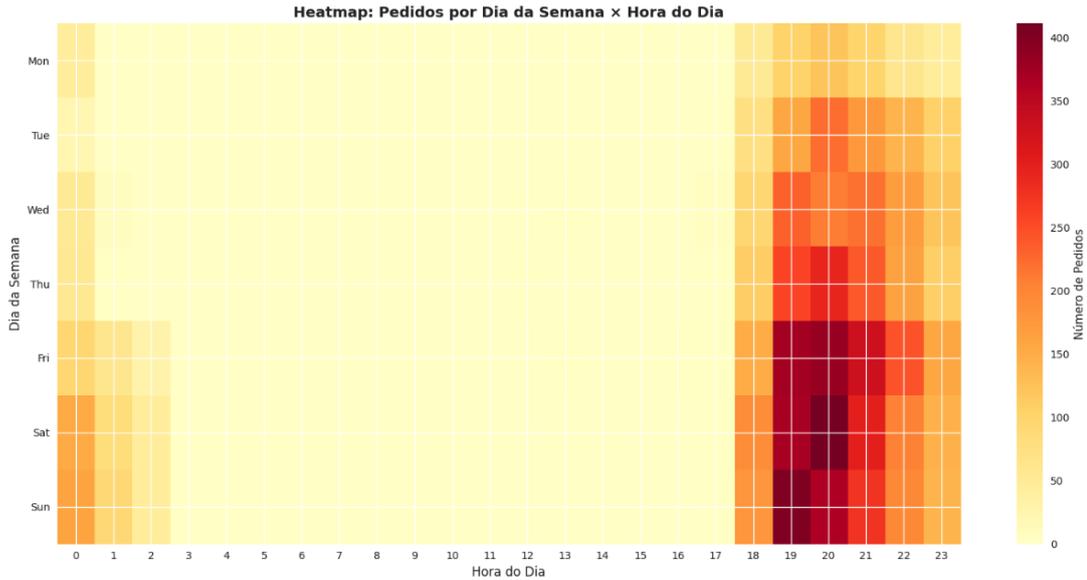


FIGURE 4.23 – Heatmap of mixed order distribution by day of week and hour of day. Color intensity indicates order volume, varying from light yellow (low volume) to dark red (high volume). The heatmap clearly evidences that periods of highest demand concentrate in hours 18h–22h, especially on weekends (Friday, Saturday, and Sunday). Early morning and morning periods show very low demand, indicating potential significant system idleness during these hours.

frequently inferior to 50 orders per combination, indicating that the system sized to meet peaks will show significant idleness periods during these hours.

Arrival rate (λ) calculation: Choice of method to calculate arrival rate λ is critical for sizing, as it will determine necessary system capacity. Three approaches were evaluated. The first consists of calculating λ as simple mean of all orders divided by total time (uniform mean), a simple method but that does not capture demand peaks. The second consists of using the largest value observed in any day+hour combination (absolute peak), a very conservative method that may oversize the system. The third consists of using the 95th percentile of mean rate by combination (day of week + hour), a balanced method that ensures robustness without excessive conservatism.

The 95th percentile method considering day of week + hour combination was chosen, as this method adequately captures specific patterns of each day of week, ensures that the system is sized to serve 95% of expected demand scenarios, avoids excessive oversizing that would result in even larger idleness periods, and offers operational robustness against demand variations.

For each *droneport*, arrival rate λ was calculated using this method, resulting in specific values that reflect spatial and temporal demand distribution of each *droneport*.

Service rate (μ) calculation: Service rate μ (orders served per hour per drone) was calculated for each *droneport* based on mean complete cycle time of orders assigned to that *droneport*. As cycle time varies according to distance to customer, each *droneport*

presents a distinct service rate μ , reflecting spatial characteristics of its coverage area.

Service rate is calculated as:

$$\mu = \frac{3600}{t_{\text{cycle_mean}}} \quad (4.7)$$

where $t_{\text{cycle_mean}}$ is mean complete cycle time of orders assigned to the *droneport*, in seconds.

Sizing via M/M/c model: For each *droneport*, the M/M/c queueing model was applied to determine the minimum number of servers (drones) c necessary to guarantee that mean total time in system W is less than or equal to the 20 min SLA.

The M/M/c model assumes arrivals following Poisson process with rate λ , exponential service times with rate μ , c identical servers operating in parallel, and FIFO (First In, First Out) queue discipline.

For each *droneport*, the minimum number of drones c that satisfies condition $W \leq 1200$ s was calculated iteratively. Metrics calculated by the model include $\rho = \frac{\lambda}{c\mu}$ (system utilization rate), P_0 (empty system probability), P_w (probability of finding queue upon arrival), W_q (mean queue waiting time), and W (mean total time in system, waiting + service).

Sizing results: Application of the M/M/c model to the 4 *droneports* resulted in the following sizing:

TABLE 4.8 – M/M/c sizing results by *droneport*.

Droneport	λ (ord/h)	μ (ord/h)	c	ρ (%)	W (min)	P(exceed SLA) (%)
Droneport 1	12,45	4,57	4	68,12	17,25	5,73
Droneport 2	10,14	4,11	4	61,63	17,52	3,76
Droneport 3	3,99	5,01	2	39,88	14,25	3,06
Droneport 4	4,57	5,99	2	38,21	11,74	1,80
Total	—	—	12	—	—	—

Figure 4.24 presents detailed visualizations of sizing results, including number of drones per *droneport*, utilization rate, mean time in system compared to SLA, and relationship between arrival rate and total capacity.

SLA guarantee analysis: In addition to mean time in system, analysis of SLA exceedance probability was performed through calculation of $P(W > 1200$ s). For M/M/c queues in steady state, probability of time in system exceeding a value t is given by:

$$P(W > t) = P_w \cdot \exp(-(c \cdot \mu - \lambda) \cdot t) \quad (4.8)$$

where P_w is probability of finding queue and $(c \cdot \mu - \lambda)$ is queue reduction rate when all servers are busy.

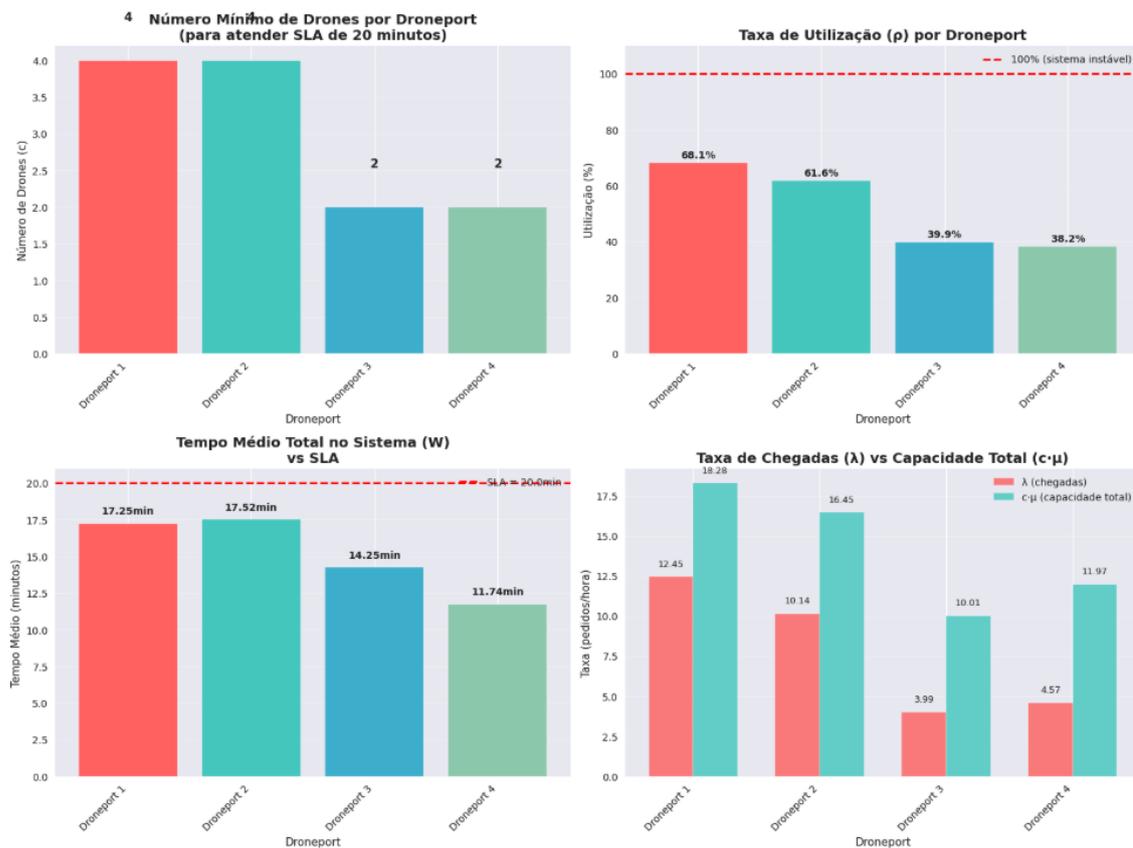


FIGURE 4.24 – M/M/c sizing analysis by *droneport*: (a) minimum number of drones necessary to meet 20-minute SLA, showing that Droneports 1 and 2 require 4 drones each, while Droneports 3 and 4 require 2 drones each; (b) utilization rate (ρ) by *droneport*, evidencing that all operate below 100% (instability limit), with Droneports 1 and 2 showing higher utilization; (c) mean total time in system (W) compared to SLA, confirming that all *droneports* meet SLA; (d) relationship between arrival rate (λ) and total capacity ($c \cdot \mu$), showing that capacity exceeds demand in all cases.

Figure 4.25 presents detailed SLA exceedance analysis by *droneport*, including exceedance probability, expected number of orders that would exceed SLA, and relationship between mean time and exceedance probability.

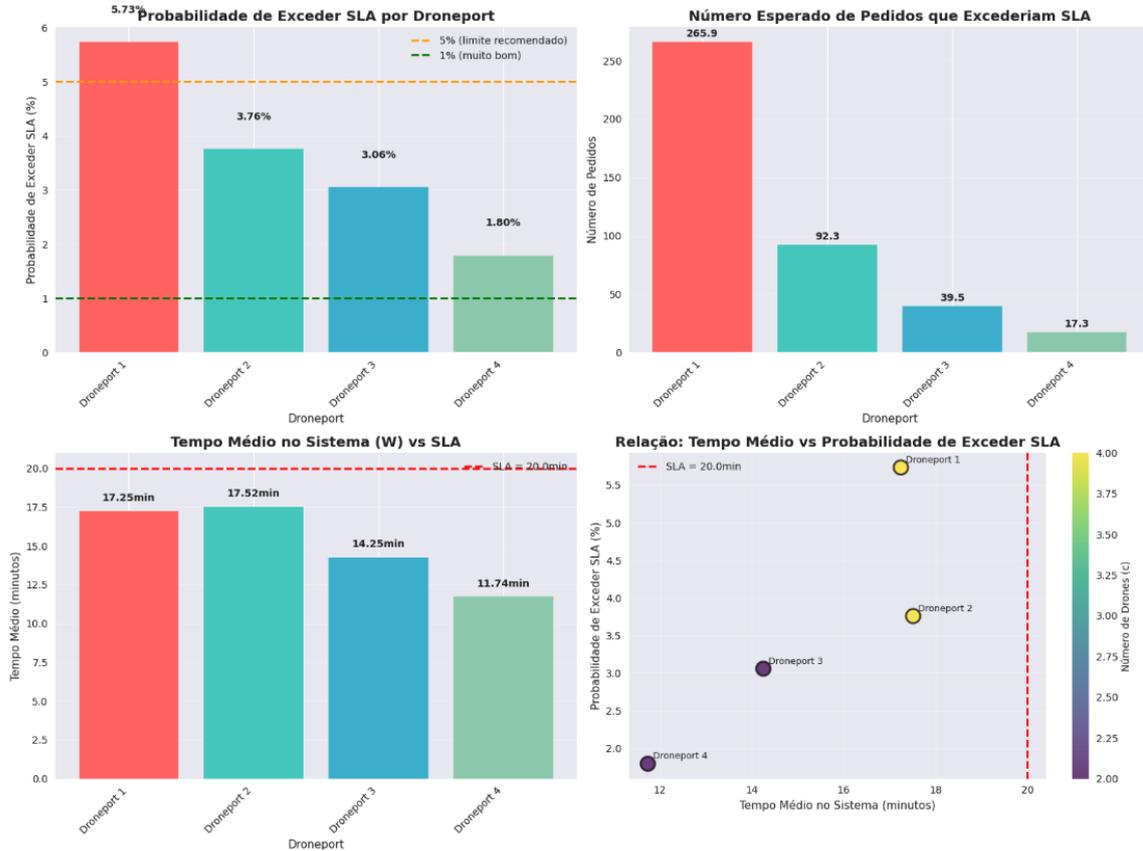


FIGURE 4.25 – SLA exceedance analysis by *droneport*: (a) probability of exceeding 20-minute SLA, showing that Droneport 1 presents highest probability (5,73%), while Droneport 4 presents lowest probability (1,80%); (b) expected number of orders that would exceed SLA, evidencing that approximately 415 orders (4,4% of total) could exceed SLA; (c) mean time in system vs SLA, confirming that all meet SLA on average; (d) relationship between mean time and exceedance probability, showing positive correlation between these metrics.

Exceedance analysis results indicate that mean exceedance probability is 3,59% of total mixed orders. Sizing guarantees that, on average, 96,41% of orders will be served within the 20-minute SLA. Droneport 1 presents highest exceedance probability (5,73%), while Droneport 4 presents lowest probability (1,80%). Approximately 415 orders (4,4% of total of 9,351 mixed orders) could exceed SLA.

Final sizing decision: Based on M/M/c analysis and SLA guarantee evaluation, the following sizing configuration was defined: Droneport 1 with 4 drones (utilization rate 68,12%, mean time 17,25 min), Droneport 2 with 4 drones (utilization rate 61,63%, mean time 17,52 min), Droneport 3 with 2 drones (utilization rate 39,88%, mean time 14,25 min), and Droneport 4 with 2 drones (utilization rate 38,21%, mean time 11,74 min).

Total drones: 12 drones for the complete system (mean of 3 drones per *droneport*).

Decision justification: This configuration was chosen to guarantee that all *droneports* present mean time in system inferior to 20 minutes (mean SLA), that mean exceedance probability of 3,59% implies that approximately 96,41% of orders will be served within SLA (SLA in 95% of orders), and that *droneports* with highest demand (1 and 2) receive greater capacity, while those with lower demand (3 and 4) receive proportionally smaller capacity (operational balance).

Important observation on individual guarantee: It is important to highlight that sizing guarantees SLA on average and for approximately 95% of orders, but does not guarantee that each *droneport* individually will always serve 100% of orders within SLA. During exceptional peak periods or extreme events, some orders may exceed SLA even with adequate sizing. This is an inherent limitation of stochastic queueing models, which work with probabilities and means instead of absolute deterministic guarantees.

Operational implications: Sizing analysis reveals several important operation characteristics. Regarding idleness periods, as sizing was performed to meet demand peaks (95th percentile), the system will show significant idleness periods during low-demand hours (early morning, morning, and early afternoon). Mean utilization rate of 51,96% indicates that, on average, approximately half of system capacity will be idle.

Regarding demand concentration, Droneports 1 and 2, which serve highest-volume restaurants, show higher arrival rates and consequently require greater capacity. This reflects asymmetry in demand distribution identified in previous exploratory analysis.

Regarding capacity vs cost trade-off, sizing was optimized to guarantee SLA without excessive oversizing. Alternatives with greater number of drones would further reduce exceedance probability, but would increase investment and operation costs, in addition to resulting in higher mean idleness.

This sizing provides the basis for subsequent economic-financial analyses, where operational and investment costs necessary to maintain this service capacity will be evaluated.

Stage 12: Energy Consumption Analysis and Battery Sizing

The processing executed in *notebook* 09 — `09_energy_consumption_analysis.ipynb` — implemented the ninth stage of the operational pipeline (Section 3.3), performing calculation of energy consumption of each aerial mission using the *segment-by-segment* model presented in Section 2.2 of Chapter 2. This stage is fundamental to determine not only energy viability of missions, but also to adequately size the number of batteries necessary per *droneport*, ensuring continuous operation without interruptions due to lack of charged batteries.

Applied energy model: The model used decomposes consumed power into two main

components, as presented in Section 2.2. The first component is hover power (P_{hover}), responsible for stationary flight and transition phases, calculated by Froude theory according to Equation (2.8):

$$P_{\text{hover}} = \frac{(mg)^{3/2}}{\eta\sqrt{2\rho A}} \quad (4.9)$$

where m is total mass (drone + cargo), g is gravitational acceleration, η is global electrical-mechanical efficiency, ρ is air density, and A is rotor disk area.

The second component is parasite power (P_{parasite}), resulting from body drag during cruise, calculated according to Equation (2.9):

$$P_{\text{parasite}} = \frac{1}{2}\rho SC_D(V + v_w)^3 \quad (4.10)$$

where S is frontal body area, C_D is drag coefficient, V is drone forward speed, and v_w is headwind component projected in flight direction.

Physical and operational parameters: Calculations were performed using technical parameters based on commercial class 3 cargo drone specifications: empty mass of 7.2 kg, mean cargo of 1.0 kg, and maximum cargo capacity of 2.4 kg; rotor radius of 0.5 m, frontal body area of 0.15 m, and drag coefficient of 0,5; global electrical-mechanical efficiency (η) of 0,72; cruise speed of 36 km/h (10 m/s), climb and descent speed of 18 km/h (5 m/s); cruise altitude of 50 m; and total battery capacity of 400 Wh, with usable capacity of 360 Wh (90% of total capacity, reserving 10% as safety margin).

Flight segmentation: Each delivery mission was segmented into six distinct phases, each with its own mass, speed, and aerodynamic configuration characteristics: climb (loaded), corresponding to flight from *droneport* to cruise altitude with loaded drone; outbound cruise (loaded), corresponding to horizontal flight from *droneport* to customer maintaining cruise altitude with loaded drone; descent (loaded), corresponding to descent from cruise altitude to ground at customer location with loaded drone; return cruise (empty), corresponding to horizontal flight from customer back to *droneport* maintaining cruise altitude with empty drone; return climb (empty), corresponding to climb back to cruise altitude after return with empty drone; and landing hover (empty), corresponding to hover phase before landing at *droneport* with fixed duration of 30 s with empty drone.

For each segment, hover and parasite powers were calculated, considering segment mass (loaded or empty) and segment duration time, obtaining consumed energy through temporal integration: $E_k = (P_{\text{hover},k} + P_{\text{parasite},k})\Delta t_k$.

Integration with wind data: An important characteristic of the model is integration of ERA5 wind data to consider impact of meteorological conditions on energy consumption. For each order, headwind component (v_w) projected in flight direction was calculated, using flight azimuth calculation (geographic flight direction droneport \rightarrow customer, calcu-

lated from geographic coordinates), wind projection (wind component in flight direction, considering that headwind $v_w > 0$ increases consumption and tailwind $v_w < 0$ reduces consumption), and segmented application (wind component was calculated separately for outbound and return, as flight direction is opposite in these segments).

Wind data used are from ERA5 database for the period June to August 2024, mapped to 2025 orders through correspondence by month, day, and hour (ignoring year). This approach allows capturing seasonal and hourly wind patterns without need for specific 2025 data.

Figure 4.26 presents distribution of headwind components for cruise segments in outbound and return. In outbound, distribution shows slight tendency for tailwind (negative values), with median of -0.086 m/s. In return, distribution shows slight tendency for headwind (positive values), with median of 0.086 m/s. Both distributions are approximately symmetric and centered near zero, indicating that extreme wind conditions are relatively rare.

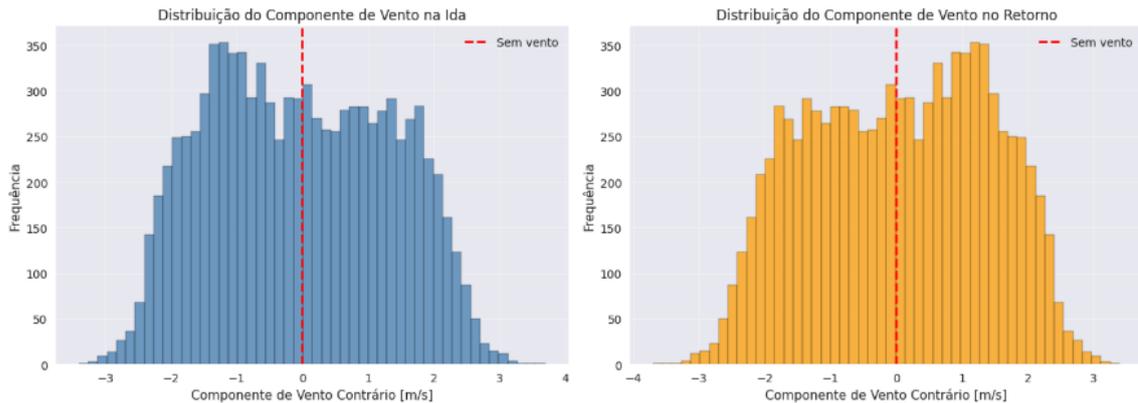


FIGURE 4.26 – Distribution of headwind components in cruise segments: (a) outbound distribution, showing slight tendency for tailwind (negative values); (b) return distribution, showing slight tendency for headwind (positive values). The dashed red line indicates no-wind condition ($v_w = 0$).

Figure 4.27 presents relationship between headwind component and energy consumption in cruise segments, evidencing significant impact of wind on consumption. Positive correlation between headwind and consumed energy is observed: in outbound cruise, correlation is $0,346$, indicating that headwind substantially increases consumption; in return cruise, correlation is $0,361$, similar to outbound. Impact is non-linear: the cubic term in parasite power equation results in more pronounced impact for strong headwinds — a 20% increase in v_w elevates P_{parasite} by approximately 73%.

Energy consumption results: The model was applied to the 9,351 orders classified as mixed delivery, resulting in mean energy per order of 157.21 Wh (39,30% of battery capacity), median energy of 158.74 Wh (39,69% of capacity), minimum energy of 13.14 Wh (3,28% of capacity), maximum energy of 332.54 Wh (83,13% of capacity), and standard deviation of 66.13 Wh.

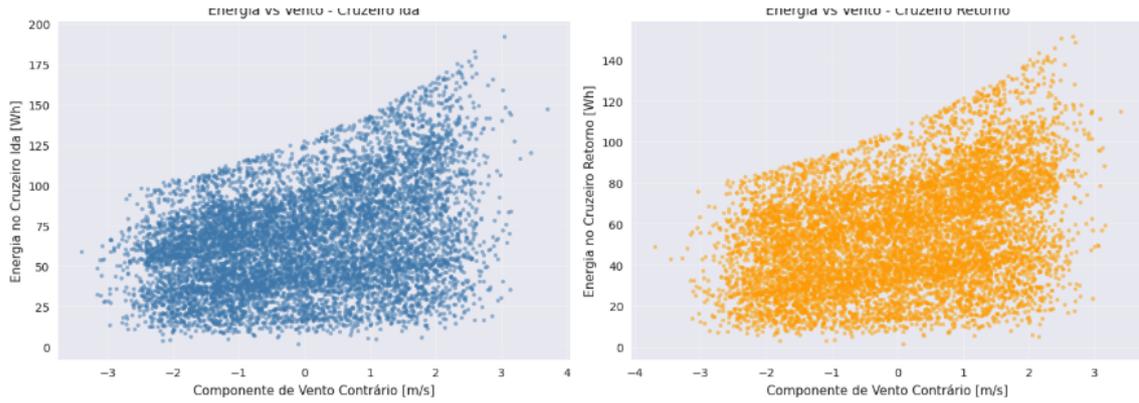


FIGURE 4.27 – Relationship between headwind component and energy consumption in cruise segments: (a) outbound cruise, showing positive correlation (0,346) between headwind and energy; (b) return cruise, showing similar correlation (0,361). Consumption in return cruise is slightly smaller than outbound due to lower mass (empty drone).

Energy viability: All 9,351 mixed orders (100%) are viable from an energy perspective, that is, consume energy inferior to the usable limit of 360 Wh. This confirms that battery capacity of 400 Wh is adequate to serve all considered missions, maintaining 10% safety margin.

Figure 4.28 presents mean contribution of each segment to total energy consumption. Outbound cruise is the largest consumer, representing 41,76% of total (65.65 Wh on average). Return cruise is the second largest consumer, representing 35,18% of total (55.30 Wh on average). Together, the two cruise segments represent 76,94% of total consumption. Climb and return climb contribute moderately (7,88% and 6,50%, respectively), while descent and landing hover are the smallest consumers (5,54% and 3,15%, respectively).

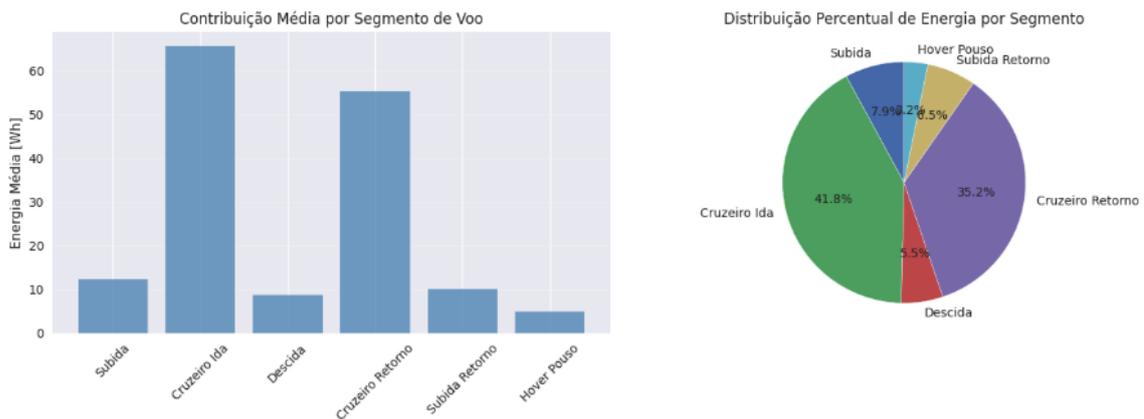


FIGURE 4.28 – Energy distribution by flight segment: (a) bar chart with mean contribution in Wh; (b) pie chart with percentage distribution. Cruise segments (outbound and return) are responsible for approximately 77% of total consumption.

Battery usage profile: Figure 4.29 presents distribution of energy consumed per order. Distribution shows mode near 100 Wh to 120 Wh, indicating concentration of short-distance orders. Distribution is slightly right-skewed, with tail extending to approximately

330 Wh. Mean (157.2 Wh) is close to median (158.74 Wh), indicating relatively balanced distribution. All orders are well below the usable limit of 360 Wh.

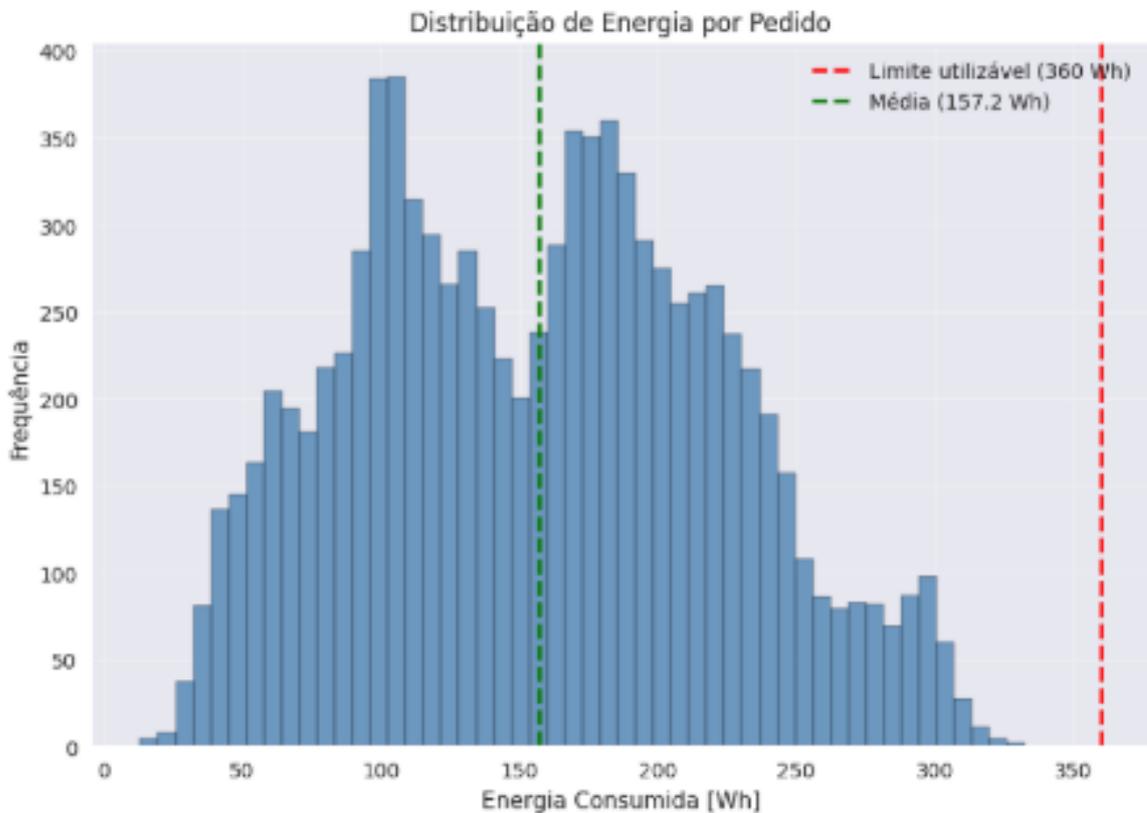


FIGURE 4.29 – Distribution of energy consumed per order. The dashed green line indicates mean (157.2 Wh) and the dashed red line indicates usable battery limit (360 Wh). All orders are within safety limit.

Figure 4.30 presents relationship between energy consumed and flight distance, confirming expected relationship. Strong positive correlation between distance and energy consumed is observed, with relatively linear distribution and proportional energy increase as distance increases. Orders of shorter distance (up to 2 km) typically consume between 50 Wh and 150 Wh, while orders of longer distance (above 5 km) typically consume between 250 Wh and 330 Wh. All points are below the usable limit, confirming energy viability for all considered distances.

Battery sizing by droneport: Once number of drones per *droneport* was determined (Stage 11), it was necessary to size the number of batteries necessary to guarantee continuous operation. Sizing considered the following operational parameters: charging time of 45 min per battery, *droneport* operation time of 120 s (2 min) including battery swap and preparation, 24 hours of operation per day (continuous operations), 365 days of operation per year, and usable battery capacity of 360 Wh (90% of 400 Wh).

Calculation methodology: Number of batteries per drone was determined considering that a battery can perform approximately 2,3 orders before needing to be recharged (based

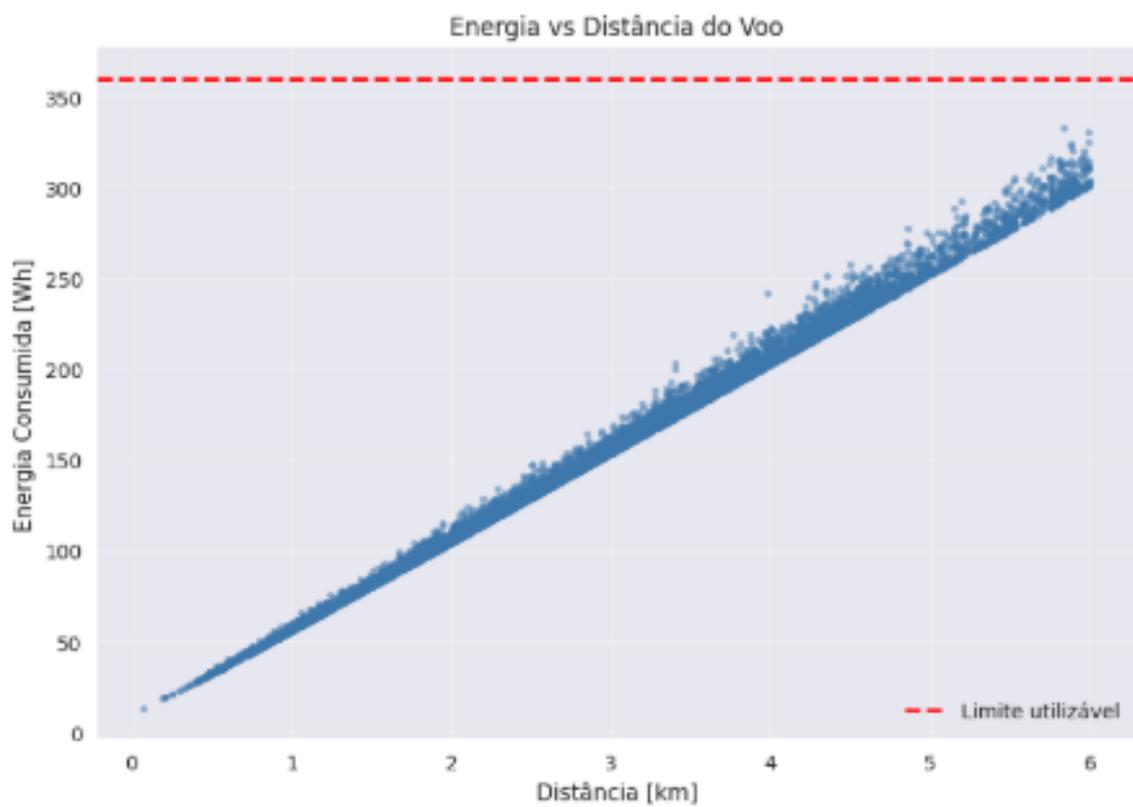


FIGURE 4.30 – Relationship between energy consumed and flight distance. The dashed red line indicates usable battery limit (360 Wh). Strong positive correlation between distance and energy is observed, with all orders within safety limit.

on mean energy of 157.21 Wh per order), flight time until discharge is approximately 0.60 h (36 minutes) until reaching usable limit, to maintain continuous operation while a battery is charging (0.75 h) the drone needs another battery to continue flying, and a 20% safety margin was added to consider demand variations and operational conditions.

Calculation resulted in need of **3 batteries per drone**, rounded up from theoretical minimum ratio of 2,69 batteries/drone.

Sizing results: Table 4.9 presents battery sizing by *droneport*, considering number of drones determined in Stage 11.

TABLE 4.9 – Battery sizing by droneport

Droneport	Drones	Batteries/drone	Total batteries	Orders/year
Droneport 1	4	3	12	18,552
Droneport 2	4	3	12	9,824
Droneport 3	2	3	6	5,172
Droneport 4	2	3	6	3,856
Total	12	3	36	37,404

Charging cycles throughout the year: Considering that analysis was performed for 3 months (June, July, and August), values were multiplied by 4 to obtain annual estimates. Total energy consumed throughout an operational year is 5880.43 kWh, distributed among the 36 batteries of the system.

Number of charging cycles per battery throughout the year varies according to demand of each *droneport*: Droneport 1 presents 153,41 cycles/battery/year (highest demand), Droneport 2 presents 92,22 cycles/battery/year, Droneport 3 presents 76,33 cycles/battery/year, Droneport 4 presents 44,95 cycles/battery/year (lowest demand), and general mean is 408,36 cycles/battery/year.

These charging cycle values are important to evaluate battery life and estimate replacement costs over time. Considering that typical lithium batteries present life of approximately 400 to 500 complete charging cycles, proposed sizing results in operation within expected life limits, with adequate margin for operational variations.

Battery utilization distribution: Analysis reveals that mean utilization is 39,30% of battery capacity per order, median utilization is 39,69% of capacity, maximum observed utilization is 83,13% of capacity, and percentiles are P10 = 17,25%, P50 = 39,69%, P90 = 60,90%, and P99 = 75,37%.

This distribution indicates that battery capacity is adequately sized, with most orders using less than half of available capacity. This allows sufficient safety margin to deal with adverse conditions (strong winds, cargo variations), operational flexibility to perform multiple consecutive orders without immediate need for recharge, and extended life, as

lower depth of discharge generally results in greater battery life.

This energy analysis provides the basis for subsequent economic-financial calculations, where electricity costs (OpEx) and battery depreciation (CapEx) will be evaluated throughout the operational life cycle.

4.2 Discussion of Economic-Financial Results

Completed operational and energy analysis of the DaaS system, economic-financial viability evaluation constitutes the final stage to determine if the proposed model is sustainable from a financial perspective. This section presents detailed analysis of capital costs (CapEx) and operational costs (OpEx), establishing direct comparisons with the motorcycle delivery model to identify operational viability points.

The adopted methodological approach starts from a fundamental premise: before evaluating sophisticated investment attractiveness metrics — such as net present value (NPV), internal rate of return (IRR), and risk analyses — it is essential to establish if the DaaS model can operate with a **competitive OpEx per order** relative to the terrestrial mode. This is a necessary, although not sufficient, condition: if OpEx per order via drone is significantly greater than OpEx per order via motorcycle, there is no point in advancing to investor return analyses, as the model would not be operationally sustainable even considering a possible premium on drone order price.

4.2.1 CapEx Analysis (Capital Expenditure)

Initial investment (CapEx) represents all capital expenditures made in **year zero** of the project, constituting the amount necessary to start DaaS service operation. This analysis was performed in *notebook 10 — 10_financial_analysis.ipynb* — and considers seven main categories of initial investment, as described in Section 3.4 of Chapter 3.

Premises and Limitations

An important limitation of this analysis refers to acquisition costs of drones used. Calculations are based on data from **Keeta G-4** (generation 4) drones, which are aircraft of own manufacture by Keeta company, commercialized through direct sale and not available in the open market. As a result, acquisition cost values used in this analysis are **estimates** based on benchmarking with similar commercial class 3 BVLOS drones and do not correspond to official public sale price values.

Need for estimates results from the scope of this work, which seeks to evaluate economic

viability of the DaaS model using real technical specifications of drones in commercial operation. Although this limitation should be recognized, estimated values were validated through comparison with market prices of commercial drones with similar characteristics (class 3, cargo capacity of 2,27 kg, range up to 24 km) and are within reasonable ranges for equipment of this category.

In addition to drone acquisition costs, other CapEx categories also incorporate estimates when it was not possible to obtain official values or direct quotes, especially for *droneport* infrastructure items and certified safety equipment. All estimates were documented and are subject to refinement through direct quotes with suppliers in subsequent analyses.

Another important premise is the choice of **acquisition** modality instead of *leasing* for CapEx calculation. This decision simplifies initial analysis and allows evaluating total necessary investment, and can be refined in future analyses that consider different financing structures.

CapEx Categories

CapEx was calculated considering seven main categories, described below:

1. Drone Acquisition: Corresponds to purchase cost of drones necessary for fleet operation. Number of drones was determined through sizing performed in Stage 11 (Section ??), using M/M/c queueing model to guarantee 20-minute SLA service. Calculation considers 12 drones in total, distributed among the 4 *droneports* as presented in Table 4.9. Estimated unit cost is USD 20,000,00 (BRL 100,000,00) per drone, resulting in total investment of USD 240,000,00 (BRL 1,200,000,00).

2. Extra Batteries: Refers to additional batteries necessary to guarantee continuous operation, allowing drones to remain in operation while other batteries are in charging process. Battery sizing was performed in Stage 12 (Section 4.1.1), resulting in need of 3 batteries per drone, totaling 36 batteries for the entire fleet. Estimated unit cost is USD 1,000,00 (BRL 5,000,00) per battery, resulting in total investment of USD 36,000,00 (BRL 180,000,00).

3. Safety Equipment: Includes mandatory equipment according to RBAC-E 94 regulation (CIVIL, 2023b): certified parachute and remote *kill-switch* system. Each drone requires a complete set of this equipment. Estimated cost is USD 1,500,00 (BRL 7,500,00) per certified parachute and USD 350,00 (BRL 1,750,00) per *kill-switch*, totaling USD 1,850,00 (BRL 9,250,00) per drone. For the 12 drones, total investment is USD 22,200,00 (BRL 111,000,00).

4. Droneport Infrastructure: Comprises three main components per *droneport*: phys-

ical structure (landing platform, maintenance area), fast charging stations for batteries, and meteorological protection (coverage, security system). Estimated cost per *droneport* is USD 15,500,00 (BRL 77,500,00), distributed between physical structure (USD 10,000,00), charging stations (USD 3,500,00), and meteorological protection (USD 2,000,00). For the 4 *droneports*, total investment is USD 62,000,00 (BRL 310,000,00).

5. Software and Communication: Includes fleet management software (operational control, monitoring, routing) and C2 (Command and Control) communication infrastructure for each *droneport*. Fleet management software is considered a unique investment for the entire operation, with estimated cost of USD 30,000,00 (BRL 150,000,00). C2 communication infrastructure (antennas, radios, communication systems) has estimated cost of USD 3,500,00 (BRL 17,500,00) per *droneport*. Total investment is USD 44,000,00 (BRL 220,000,00).

6. Regulatory Licensing: Corresponds to initial certification and homologation costs necessary for commercial operation: RBAC-E 94 certification (ANAC) (CIVIL, 2023b) and ANATEL homologation for communication equipment (OES, 2019). Estimated cost is USD 10,000,00 (BRL 50,000,00) for RBAC-E 94 certification and USD 2,000,00 (BRL 10,000,00) for ANATEL homologation, totaling USD 12,000,00 (BRL 60,000,00).

7. Training: Refers to initial training of remote operators and technical support personnel. Calculation considers training of 3 remote operators and 2 maintenance technicians, with estimated cost of USD 1,000,00 (BRL 5,000,00) per person. Total investment is USD 5,000,00 (BRL 25,000,00).

Consolidated Results

Table 4.10 presents detailed CapEx breakdown by component, consolidating all initial investments necessary to start DaaS service operation.

TABLE 4.10 – Detailed CapEx breakdown by component

Component	Total Cost (USD)	Total Cost (BRL)	Percentage
Drone Acquisition	240,000,00	1,200,000,00	56,98%
Extra Batteries	36,000,00	180,000,00	8,55%
Safety Equipment	22,200,00	111,000,00	5,27%
Droneport Infrastructure	62,000,00	310,000,00	14,72%
Software and Communication	44,000,00	220,000,00	10,45%
Regulatory Licensing	12,000,00	60,000,00	2,85%
Training	5,000,00	25,000,00	1,19%
Total CapEx	421,200,00	2,106,000,00	100,00%

Total CapEx necessary to start operation is **USD 421,200,00** (BRL 2,106,000,00). Fig-

ure 4.31 presents visualization of CapEx distribution by component, evidencing that drone acquisition represents the largest portion of investment (56,98%), followed by *droneport* infrastructure (14,72%) and software and communication (10,45%).

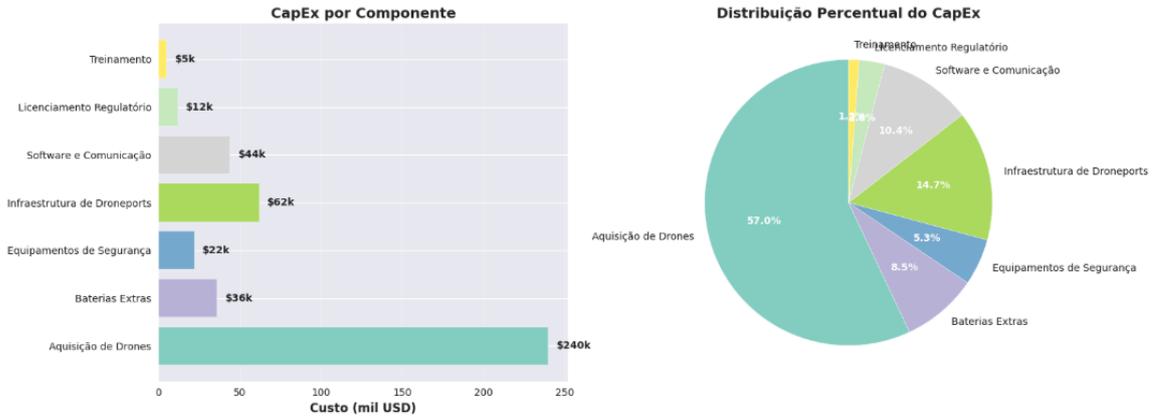


FIGURE 4.31 – CapEx distribution by component: (a) horizontal bar chart with absolute values in thousand USD; (b) pie chart with percentage distribution. Drone acquisition represents approximately 57% of total investment.

CapEx by Droneport

Figure 4.32 presents CapEx breakdown distributed by *droneport*, evidencing significant differences between larger *droneports* (Droneport 1 and 2) and smaller ones (Droneport 3 and 4).

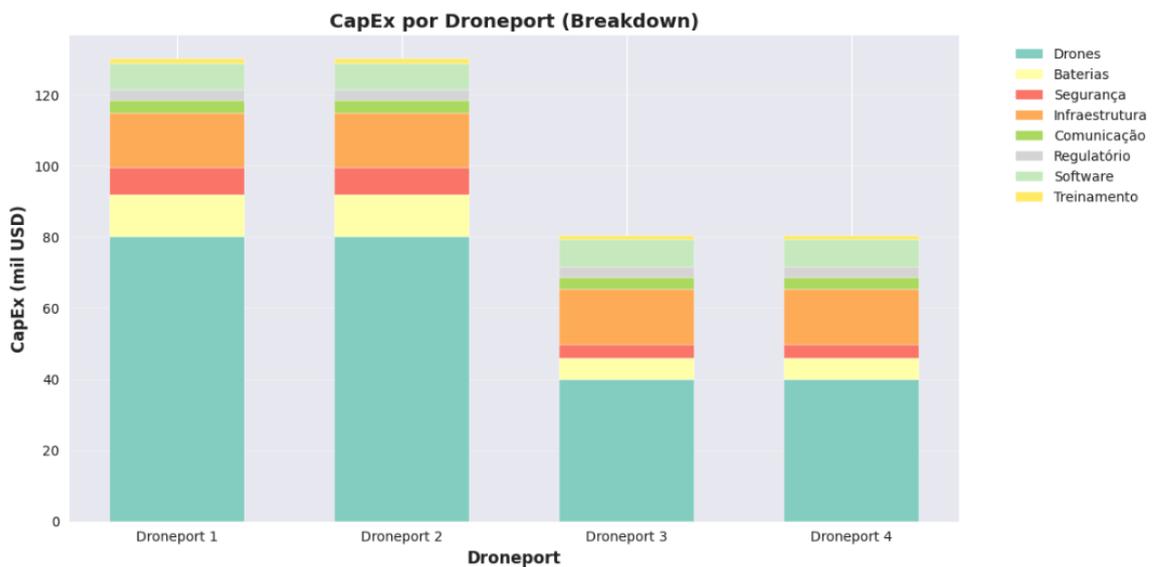


FIGURE 4.32 – CapEx by droneport with breakdown by component. Droneports 1 and 2 show higher investment due to greater number of drones (4 drones each), while Droneports 3 and 4 operate with 2 drones each. Shared costs (regulatory, software) were allocated proportionally.

Droneports 1 and 2 show total CapEx of approximately USD 130,000,00 each, while Droneports 3 and 4 show total CapEx of approximately USD 80,000,00 each. This differ-

ence results mainly from greater number of operational drones in the first *droneports* (4 drones) compared to the latter (2 drones), as determined by fleet sizing (Stage 11).

Shared costs — regulatory licensing, fleet management software, and training — were allocated proportionally among *droneports* based on number of drones of each unit. This allocation methodology reflects that these investments benefit the entire operation, but their utilization is distributed according to scale of each *droneport*.

Detailed results of CapEx calculation were exported to file `capex_breakdown.csv`, available in `data/results/` of *notebook* 10, allowing transparency and reproducibility of analysis.

4.2.2 OpEx Analysis (Operational Expenditure)

Operational costs (OpEx) represent all recurring expenditures necessary to maintain DaaS service operation throughout a year, different from CapEx which corresponds to non-recurring initial investments. This analysis complements CapEx evaluation and is fundamental to determine operational viability of the model, especially through the **OpEx per order** metric, which allows direct comparison with the motorcycle delivery model.

OpEx calculation was performed in *notebook* 10 — `10_financial_analysis.ipynb` — considering eight main categories of recurring operational costs, as described in Section 3.4 of Chapter 3.

OpEx Categories

OpEx was calculated considering eight main categories, described below:

1. Electricity Consumption: Corresponds to cost of electrical energy necessary to charge drone batteries throughout the year. Calculation uses mean electricity tariff of USD 0,63 per kWh (approximately BRL 3,15 per kWh, considering exchange rate of USD 1 = BRL 5), applied to total energy consumption calculated in Stage 12 (Section 4.1.1). Total energy consumed throughout an operational year is 5880.43 kWh, considering the 37,404 orders served and charging cycles necessary to maintain continuous operation. Annual electricity cost is USD 882,07 (BRL 4,292,72).

2. Maintenance: Includes preventive and corrective maintenance costs of drones throughout the year. Preventive maintenance is performed according to operational intervals recommended by manufacturer, while corrective maintenance covers replacement of components such as rotors, sensors, and other elements subject to wear. Calculation considers mean annual costs per drone, totaling USD 6,600,00 (BRL 33,000,00) for the entire fleet of 12 drones.

3. Battery Replacement: Refers to cost of replacing batteries that reach end of their useful life throughout the operational year. Calculation is based on number of charging cycles per battery and expected battery life (approximately 400 to 500 complete cycles). Considering that mean cycles per battery throughout the year is 408,36 cycles (Stage 12) and assuming useful life of 450 cycles, annual replacement rate is approximately 90,75% of total number of batteries. Replacement cost uses the same unit value as CapEx (USD 1,000,00 per battery), resulting in annual cost of USD 32,669,08 (BRL 163,345,40).

4. Insurance: Comprises two types of mandatory insurance: aeronautical civil liability insurance (RC-RETA), with minimum coverage of R\$ 500,000 according to regulation, and equipment insurance, which covers drones and equipment against damage and losses. RC-RETA insurance has annual cost of USD 2,400,00 (BRL 12,000,00) per drone, while equipment insurance corresponds to approximately 2% of insured value per year. Total annual insurance cost is USD 34,764,00 (BRL 173,820,00).

5. Salaries and Operation: Represents the largest portion of OpEx and includes costs with personnel necessary for service operation. Calculation considers two main groups: remote operators and technical support personnel. For remote operators, the necessary for complete 24-hour operation coverage was considered, which demands 3 work shifts (8 hours each shift). Each shift requires 1 remote operator, totaling 3 operators for continuous coverage. Additionally, 2 technical support people available during the week were considered for maintenance, operational support, and emergency response. Hourly values consider Brazilian market reality for these specialized functions. Total annual salary and operation cost is USD 113,700,00 (BRL 568,500,00).

6. Amortization: Corresponds to linear amortization of investments in software and infrastructure made in CapEx, recognizing that these assets have limited useful life and must have their value depreciated over time. Amortization considers two main components: fleet management software (useful life of 5 years) and communication and *droneport* infrastructure (useful life of 7 years). Although amortization does not represent an effective cash disbursement, it reflects economic degradation of investments and is important for adequate financial evaluation. Annual amortization cost is USD 16,857,14 (BRL 84,285,71).

7. Operational Licenses: Includes fees and annual renewals necessary to maintain regulatory operation: RBAC-E 94 certification renewal (ANAC) (CIVIL, 2023b), SARPAS fees (Sistema de Autorização de Rastreamento de Produtos Aeronáuticos) (AÉREO, 2022), and ANATEL homologation renewal (OES, 2019). Although these licenses are part of operational cost, they also reflect regulatory cost of maintaining legalized operation. Total annual operational license cost is USD 3,300,00 (BRL 16,500,00).

8. Facility Rental: Refers to rental cost of physical spaces where *droneports* are

located. Calculation considers a monthly value per *droneport*, which was estimated at a high value relative to mean commercial rental in Brazil, reflecting need for strategic locations, adequate for drone operation (access, security, proximity to high-demand areas). For the 4 *droneports*, annual rental cost is USD 48,000,00 (BRL 240,000,00).

Consolidated Results

Table 4.11 presents detailed OpEx breakdown by component, consolidating all annual operational costs necessary to maintain DaaS service operation.

TABLE 4.11 – Detailed OpEx breakdown by component (annual)

Component	Annual Cost (USD)	Annual Cost (BRL)	Percentage
Electricity Consumption	882,07	4,292,72	0,34%
Maintenance	6,600,00	33,000,00	2,57%
Battery Replacement	32,669,08	163,345,40	12,72%
Insurance	34,764,00	173,820,00	13,54%
Salaries and Operation	113,700,00	568,500,00	44,28%
Amortization	16,857,14	84,285,71	6,57%
Operational Licenses	3,300,00	16,500,00	1,29%
Facility Rental	48,000,00	240,000,00	18,69%
Total Annual OpEx	256,772,29	1,283,743,83	100,00%

Total annual OpEx necessary to maintain operation is **USD 256,772,29** (BRL 1,283,743,83). Figure 4.33 presents visualization of OpEx distribution by component, evidencing that salaries and operation represent the largest portion of operational cost (44,28%), followed by facility rental (18,69%) and insurance (13,54%).

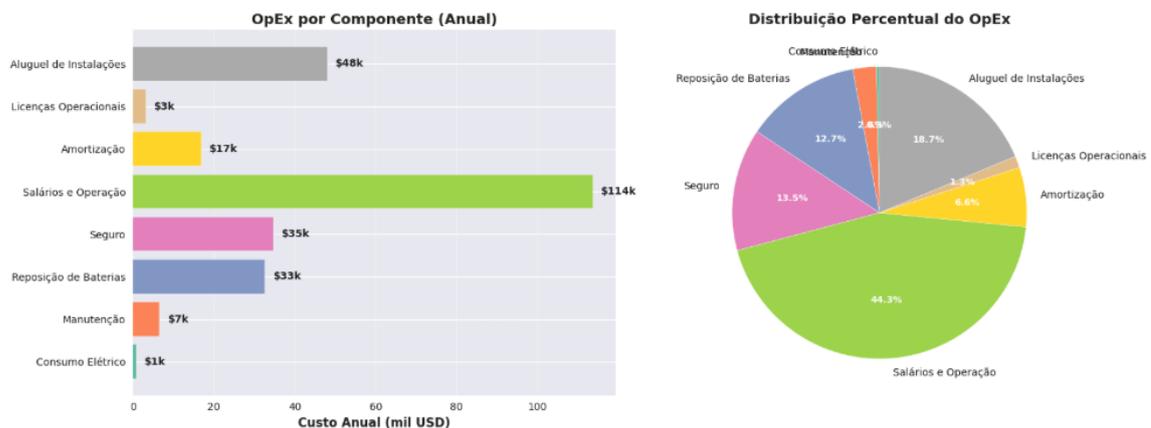


FIGURE 4.33 – OpEx distribution by component: (a) horizontal bar chart with absolute values in thousand USD; (b) pie chart with percentage distribution. Salaries and operation represent approximately 44% of total operational cost.

OpEx per Order

A fundamental metric for viability evaluation is **OpEx per order**, which allows directly comparing operational cost of the DaaS model with operational cost of the motorcycle delivery model. This indicator is calculated by dividing total annual OpEx by number of orders served throughout the year.

Considering that the system serves 37,404 orders throughout an operational year (based on 3-month analysis multiplied by 4 for annual projection), OpEx per order is calculated as:

$$\text{OpEx per order} = \frac{\text{Total annual OpEx}}{\text{Number of annual orders}} = \frac{\text{USD } 256,772,29}{37,404} = \text{USD } 6,86 \quad (4.11)$$

OpEx per order via drone is USD 6,86 (approximately BRL 34,30). This metric will be fundamental for comparative analysis with the motorcycle delivery model, which will be presented in the next subsection, allowing evaluation if the DaaS model can operate with competitive costs relative to traditional terrestrial mode.

It is important to highlight that salaries and operation represent the largest portion of OpEx (44,28%), followed by facility rental (18,69%). These two components together represent more than 60% of total operational cost, indicating that optimizations in operational efficiency and facility cost negotiation can have significant impact on OpEx per order reduction.

Electricity consumption, in turn, represents only 0,34% of total OpEx, confirming that energy is not a limiting factor from an economic perspective, despite being fundamental for technical viability of operations.

Detailed results of OpEx calculation were exported to file `opex_breakdown.csv`, available in `data/results/` of *notebook* 10, allowing transparency and reproducibility of analysis.

Exploratory Analysis: Vertical Scale and OpEx per Order Variation

Initial sizing presented previously considers an operation at relatively small scale, serving 37,404 orders throughout an operational year. In this context, it was identified that salaries and operation represent 44,28% of total OpEx, while facility rental represents 18,69% — totaling more than 60% of operational cost in components that are essentially fixed, regardless of volume of orders served.

This result may seem counterintuitive at first glance, given that we are replacing an

intrinsically human operation (deliveries by humans and vehicles) with an autonomous operation. However, this cost distribution precisely reflects the scale problem: in small operations, fixed costs — especially operational personnel and infrastructure — are diluted over a limited volume of orders, resulting in high OpEx per order.

To understand how these costs vary with operation expansion in the same city, an exploratory analysis was performed in *notebook 10 — 10_financial_analysis.ipynb* — exploring **vertical scale** strategies, that is, increasing served demand without increasing existing infrastructure (number of drones, *droneports*, operational personnel, or facility rental). This analysis aims to identify potential for OpEx per order reduction through better utilization of installed capacity.

Figure 4.34 presents demand distribution by hour of day observed in historical data, evidencing enormous concentration of orders in evening hours (18h–23h) and practically non-existent demand during daytime period (5h–17h). This asymmetric distribution is a reflection of the profile of restaurants already served by Brendi platform, which concentrate in establishments with predominantly evening operation.



FIGURE 4.34 – Demand distribution by hour of day in the analyzed period. Concentration of orders in evening hours (18h–23h) evidences significant infrastructure idleness during daytime period, especially between 5h and 17h. Evening period (21h–4h) is highlighted in light purple.

First Approach: Theoretical Maximum Load Calculation The first exploratory approach aims to determine the **theoretical minimum OpEx** that would be possible to achieve operating existing infrastructure at its maximum capacity, minimizing idleness while guaranteeing 20-minute SLA service for 95% of orders.

Methodology: Calculation was performed assuming current infrastructure — 4 *droneports* and 12 drones — and using arrival rate λ of the 95th percentile of the $M/M/c$ model used in fleet sizing (Stage 11, Section ??). This rate represents the maximum load that the system can support maintaining probability of exceeding SLA at maximum 5% (that is, guaranteeing that 95% of orders are served within 20 minutes). Calculation maintains fixed all OpEx components that do not vary with number of orders: salaries and

operation, facility rental, amortization, and operational licenses. Only components proportional to order volume are recalculated: electricity consumption, maintenance, battery replacement, and equipment insurance.

Results: Maximum theoretical annual demand that the system can serve maintaining 95% SLA is **68,050 orders**, representing an increase of 81,9% relative to current demand (37,404 orders). Total annual OpEx increases due to variable components that are proportional to order volume (electricity consumption, maintenance, battery replacement, and equipment insurance). However, as fixed components — salaries and operation, facility rental, amortization, and operational licenses — represent more than 70% of total OpEx and do not vary with order volume, increase in total OpEx is proportionally smaller than demand increase. As a result, OpEx per order at theoretical maximum load is **USD 3,77**, representing a reduction of 45,0% relative to initial OpEx of USD 6,86 per order, achieved through dilution of fixed costs over a larger volume of orders.

Table 4.12 presents detailed distribution of OpEx per order at maximum load, evidencing how each cost category contributes to theoretical minimum OpEx.

TABLE 4.12 – OpEx per order distribution at theoretical maximum load (68,050 orders/year)

Category	OpEx/Order (USD)	% of Total
Electricity Consumption	0,0130	0,34%
Maintenance	0,0970	2,57%
Battery Replacement	0,4801	12,72%
Insurance	0,5109	13,54%
Salaries and Operation	1,6708	44,28%
Amortization	0,2477	6,57%
Operational Licenses	0,0485	1,29%
Facility Rental	0,7054	18,69%
TOTAL	3,7733	100,00%

It is important to highlight that this result represents a **theoretical minimum limit** of OpEx per order, assuming that it would be possible to achieve this demand distributed ideally throughout the year, maintaining the system operating near maximum capacity during the entire period. In practice, this scenario is difficult to achieve due to seasonal, intraday, and between-days-of-week variations observed in historical data, in addition to need to maintain operational safety margins.

Second Approach: Demand Mirroring A more realistic and plausible approach to achieve consists of exploring temporal demand redistribution strategies, taking advantage of infrastructure idleness periods. As evidenced by Figure 4.34, the system shows significant idleness during daytime period, especially between 10h and 15h, traditional lunch time.

Although current restaurant portfolio of Brendi platform concentrates in establishments with evening operation, this does not mean that there is no potential delivery demand at lunch time in São José dos Campos. A *go-to-market* (GTM) strategy directed at restaurants that operate predominantly at lunch time could increase demand in this period, improving utilization of installed infrastructure.

Methodology: A **demand mirroring** strategy was implemented, where order distribution observed in evening period (17h–22h) is replicated to lunch period (10h–15h), applying a volume correction factor to reflect that potential lunch demand may not reach the same volumes observed in evening period. Specifically, a mirroring factor was applied that maps each hour of evening period to a corresponding hour in lunch period:

For each pair of mapped hours, a volume correction factor of 60% was applied, that is, 60% of demand observed in evening hour is added to already existing demand in corresponding hour of lunch period. This factor reflects a conservative estimate of potential demand that could be captured through directed GTM strategy.

Results: Total annual demand after mirroring is **57,645 orders**, representing an increase of 54,1% relative to current demand. Total annual OpEx increases due to variable components that are proportional to order volume (electricity consumption, maintenance, battery replacement, and equipment insurance), which grow with greater number of flights performed. However, as fixed components — salaries and operation, facility rental, amortization, and operational licenses — do not vary with order volume and represent the largest portion of total OpEx, increase in total OpEx is proportionally smaller than demand increase. As a result, OpEx per order after mirroring is **USD 4,45**, representing a reduction of USD 2,39 (35,0%) relative to initial OpEx, achieved through dilution of fixed costs over a larger volume of orders.

Figure 4.35 presents hourly demand distribution before and after mirroring, visually evidencing redistribution of orders from evening period to lunch period.

Table 4.13 presents detailed distribution of OpEx per order after mirroring, including absolute and percentage reduction of each category relative to initial scenario.

Figure 4.36 presents a visual comparison of OpEx per order before and after mirroring, highlighting reduction of USD 2,39 per order achieved through this strategy.

It is important to highlight that demand mirroring represents a **plausible and achievable** scenario through business strategies, different from theoretical maximum load scenario that assumes ideal operational conditions difficult to replicate in practice. Mirroring strategy does not require changes in physical infrastructure or personnel structure, being dependent only on capacity to attract new restaurants with operational profile complementary to current profile.

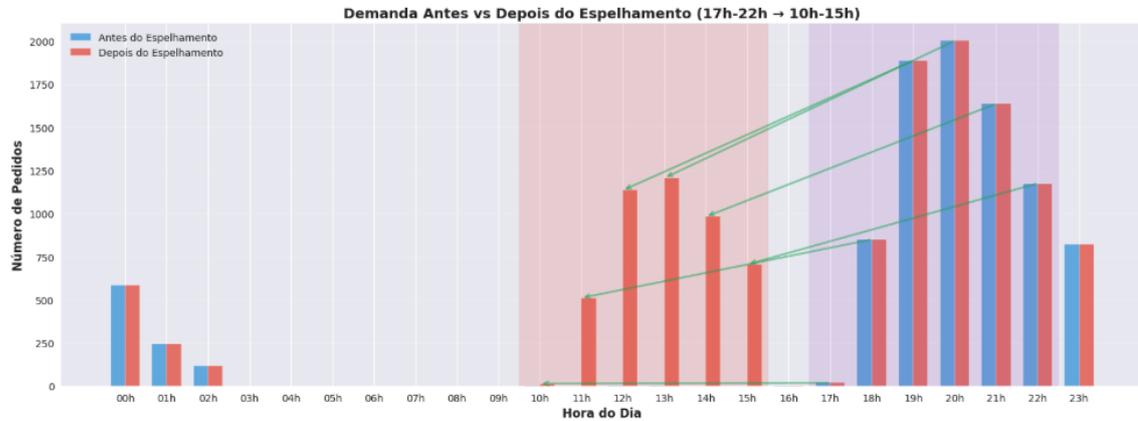


FIGURE 4.35 – Demand distribution before and after mirroring (17h–22h → 10h–15h). Blue bars represent original demand, while red bars represent demand after mirroring. Green arrows visually indicate mapping between evening hours and lunch hours. Shaded areas highlight periods involved in mirroring: light purple for source evening period (17h–22h) and light red for target lunch period (10h–15h).

TABLE 4.13 – OpEx per order distribution after demand mirroring (57,645 orders/year)

Category	OpEx/Order (USD)	% of Total	Reduction (USD)
Salaries and Operation	1,9724	44,28%	1,0591
Facility Rental	0,8327	18,69%	0,4471
Insurance	0,6031	13,54%	0,3238
Battery Replacement	0,5667	12,72%	0,3043
Amortization	0,2924	6,56%	0,1571
Maintenance	0,1145	2,57%	0,0615
Operational Licenses	0,0572	1,28%	0,0308
Electricity Consumption	0,0153	0,34%	0,0082
TOTAL	4,4543	100,00%	2,3919



FIGURE 4.36 – OpEx per order comparison before and after demand mirroring. Reduction of USD 2,39 per order (from USD 6,86 to USD 4,45) represents significant improvement in operational efficiency through better utilization of existing infrastructure.

Comparative Discussion of Scenarios Table 4.14 presents a consolidated comparison of the three analyzed scenarios, allowing visualization of potential for OpEx per order reduction through vertical scale strategies.

TABLE 4.14 – OpEx per order comparison in different vertical scale scenarios

Scenario	Annual Demand	OpEx/Order (USD)	Reduction
Current Demand	37,404	6,86	—
Demand Mirroring	57,645	4,45	35,0%
Theoretical Maximum Load	68,050	3,77	45,0%

Results demonstrate that vertical scale strategies — increasing served demand without increasing infrastructure — can result in significant OpEx per order reductions, especially through better utilization of fixed components such as salaries and facility rental.

The **demand mirroring** scenario presents itself as the most promising from a practical perspective, offering a 35% reduction in OpEx per order (from USD 6,86 to USD 4,45) through a viable business strategy. This reduction is achieved mainly through dilution of fixed costs over a larger volume of orders, without need for additional infrastructure investments.

The **theoretical maximum load** scenario, in turn, represents the theoretical minimum limit of OpEx per order (USD 3,77), but assumes ideal operational conditions that may be difficult to achieve in practice due to seasonal variations, need for safety margins, and limitations in capacity to attract demand distributed uniformly over time.

It is important to highlight that both scenarios maintain fixed OpEx components related to salaries, rental, amortization, and operational licenses, reflecting vertical scale premise. In a long-term analysis, it would be possible to explore also horizontal expansion (increasing number of *dronereports* and drones) combined with vertical scale strategies, potentially achieving even larger demand multipliers and additional OpEx per order reductions through scale economies in components such as software, regulatory licensing, and centralized management.

Detailed results of this exploratory analysis were exported to CSV files in *notebook* 10, allowing transparency and reproducibility of calculations.

4.3 Model Validation and Limitations

Validation of the developed operational cost model is fundamental to guarantee that obtained results are representative and can be compared with real market operations. This section presents a comparative validation using the iFood Delivery Plan business model as reference, in addition to discussing main assumed limitations and their implications for presented results.

4.3.1 Validation with Reference Model (iFood Delivery Plan)

To validate results of the OpEx per order model developed for DaaS service, a comparative analysis was performed using as reference the business model of **iFood Delivery Plan**, which offers food delivery service through partner deliverers (motoboy and cyclists). This choice is appropriate as it represents a real and consolidated case in the Brazilian market, functioning as a demand centralizer for deliverers to make deliveries of orders from restaurants registered on the platform.

The iFood Delivery Plan model presents characteristics that facilitate OpEx per order estimation: commission rates and cost structure are public and well documented, the model operates as a marketplace platform connecting restaurants and deliverers, cost structure does not have strong scale/volume dependence once deliverers are autonomous partners who register on the platform, and operational costs are essentially variable per order (payment to deliverer and processing fees).

Estimation Methodology

Analysis was performed in *notebook* 12 — `12_motofrete_opex_comparison.ipynb` — using the same database of delivered orders analyzed in the drone model, totaling **23,335**

delivered orders. Methodology is based on official iFood rates, as documented in public sources¹.

iFood fee structure comprises the **Basic Plan (Marketplace)**, with 12% commission on order value when restaurant makes its own delivery; the **Delivery Plan (Fullservice)**, with 23% commission when iFood manages deliveries; the difference between plans of 11% (23% - 12%) represents revenue destined to subsidize delivery service; an **online payment fee** of 3,2% on orders paid via iFood; and a **minimum fee for deliverer** of R\$ 7,50 for motorcycles and cars.

OpEx per order calculation model:

For each order, OpEx was calculated as the sum of two main components. The first is **deliverer cost**, represented by a variable delivery fee based on actual distance traveled (Route Distance), calculated as:

$$\text{Delivery fee} = \max(\text{R\$ } 7,50, \text{R\$ } 7,50 + 0,80 \times \text{Distance (km)}) \quad (4.12)$$

where base fee of R\$ 7,50 corresponds to minimum guaranteed to deliverer and additional R\$ 0,80 per km reflects cost variation with distance traveled. This model was adjusted so that mean fee results in approximately R\$ 9,00, mean delivery value observed in the market.

The second component is **online payment fee** of 3,2% of order value, assuming that all orders are paid via iFood (conservative scenario that considers worst case for OpEx).

Therefore, total OpEx per order is given by:

$$\text{OpEx per order} = \text{Deliverer cost} + \text{Online payment fee} \quad (4.13)$$

It is important to highlight that this model considers only **direct operational costs** (payment to deliverer and payment processing fee), not including indirect costs such as marketing, technology development, customer support, IT infrastructure, and overhead, which would be part of total OpEx of a complete operation. This choice is intentional and allows a more direct comparison with the drone model, which also considers only direct operational costs related to delivery execution.

Analysis Results

Applying the described methodology to the base of 23,335 delivered orders, the following results were obtained: mean OpEx per order of R\$ 13,48, median of R\$ 12,91, minimum of R\$ 7,59, and maximum of R\$ 32,55.

¹iFood rates available at: <https://blog-parceiros.ifood.com.br/taxas-ifood/>

Mean composition of OpEx per order shows that **mean delivery cost** corresponds to R\$ 11,33 (84,0% of total OpEx), while **mean payment fee** corresponds to R\$ 2,15 (16,0% of total OpEx).

Converting to dollars using the same exchange rate adopted in the drone model (USD 1 = BRL 5), **mean OpEx per order of the iFood Delivery Plan model is USD 2,70**.

Comparison with Drone Model

Table 4.15 presents a consolidated comparison of OpEx per order between the drone model and the reference model (iFood Delivery Plan).

TABLE 4.15 – OpEx per order comparison: drone model vs iFood Delivery Plan

Model/Scenario	OpEx/Order (USD)
iFood Delivery Plan (reference)	2,70
Drones — Current Demand	6,86
Drones — Demand Mirroring	4,45
Drones — Theoretical Maximum Load	3,77

Results demonstrate that the drone model presents OpEx per order significantly superior to the reference model (iFood) in the current demand scenario (USD 6,86 vs USD 2,70), representing a differential of USD 4,16 per order (154% more expensive).

However, through vertical scale strategies, the drone model approaches the reference model. In the **demand mirroring** scenario, OpEx reduces to USD 4,45, still 65% superior to the iFood model, but with absolute difference reduced to USD 1,75 per order. In the **theoretical maximum load** scenario, OpEx reduces to USD 3,77, representing a difference of only USD 1,07 per order (40% superior to the iFood model).

Comparison Limitations and Considerations

It is fundamental to recognize inherent limitations of this comparison, which result from structural differences between business models.

Regarding **business model differences**, the iFood model operates as a marketplace platform, connecting restaurants and autonomous deliverers, without need for investment in own delivery infrastructure. In contrast, the drone model requires significant investment in infrastructure (drones, *droneports*, communication systems) and own operation with specialized personnel. iFood model costs are essentially variable (payment per delivery performed), while the drone model has substantial fixed costs (salaries, rental, amortization) that are diluted over order volume.

Scalability presents significant differences between models. The iFood model has practically infinite scalability through addition of new partner deliverers, without need for additional infrastructure investment. The drone model, in turn, has capacity limits defined by installed infrastructure (number of drones and *dronereports*), requiring additional investment (CapEx) for expansion beyond theoretical maximum capacity.

Cost coverage considered in the analysis also presents particularities. iFood model analysis considers only direct delivery and payment processing costs, not including indirect costs such as marketing, technology development, customer support, and IT infrastructure. The drone model also considers only direct operational costs related to delivery execution, maintaining comparison on similar bases.

Regarding **volume dependence**, the iFood model does not have strong scale dependence, once deliverers are autonomous partners who register on the platform according to demand. The drone model presents strong scale dependence due to substantial fixed costs, requiring reaching high volumes to adequately dilute fixed costs and make OpEx per order competitive.

Regarding **service quality**, the iFood model offers terrestrial deliveries with delivery times typically superior to 30 minutes, subject to traffic conditions and deliverer availability. The drone model offers aerial deliveries with 20-minute SLA, providing greater speed and reliability, especially in routes where aerial route offers significant distance gains (high DSR).

Finally, regarding **economic viability**, despite drone model OpEx per order being superior to the iFood model, it is important to consider that the drone model offers a value differential through delivery speed (20-minute SLA vs 30+ minutes) and operational reliability independent of traffic conditions. This differential can justify a price premium that compensates higher OpEx, especially for market segments that value speed and reliability.

Validation Conclusion Comparison with the iFood Delivery Plan model validates that the developed cost model produces coherent and comparable results with real market operations. Drone model OpEx per order (USD 6,86 in current scenario) is superior to the reference model (USD 2,70), but through vertical scale strategies can be reduced to close values (USD 3,77 in theoretical maximum load scenario). Observed difference is explained by structural characteristics of models: iFood has essentially variable costs and infinite scalability through autonomous partners, while the drone model has substantial fixed costs that require high volumes for adequate dilution, but offers value differential through superior speed and reliability. This validation confirms that the cost model is well-founded and produces realistic results for economic viability analyses.

4.3.2 Horizontal Scale Analysis

While the vertical scale analysis presented previously explores strategies to increase served demand without expanding installed infrastructure, **horizontal scale** analysis investigates the scenario of geographic and market expansion to serve a significant fraction of the total *food delivery* market of São José dos Campos, not only orders currently intermediated by the Brendi platform.

This analysis is fundamental to understand market potential and establish bases for future evaluations of economic viability of a full-scale operation, where the DaaS service would serve not only current partner restaurants, but all potential food delivery demand in the city.

It is important to highlight that, at this stage, the focus is to establish context and estimate market potential through calculation of current and maximum plausible market share. Infrastructure sizing necessary to serve expanded demand and OpEx and ROI analysis in the horizontal scale scenario will be presented in subsequent analyses.

Current Market Share and Potential Demand

Current simulation considers only orders intermediated by the Brendi platform in the analyzed period. As presented in *notebook 14 — 14_market_share_analysis.ipynb* —, analysis of historical data of 3 months (90 days, from June to August 2025) resulted in **23,335 delivered orders**, which when projected annually equate to approximately **94,636 orders per year**.

To estimate the total *food delivery* market in São José dos Campos, a methodology based on official national data from iFood (IFOOD, 2025; IFOOD; IPEC, 2024) was used. iFood intermediated 120 million orders per month in 2025, totaling 1,44 billion orders per year (IFOOD, 2025). Ipsos-Ipec Research indicates that iFood has 25,7% market share in the Brazilian delivery market in 2024 (IFOOD; IPEC, 2024). From these data, it is calculated that the total Brazilian market is approximately 5,6 billion orders per year. Considering Brazilian population of 203,06 million inhabitants (IBGE 2024), a metric of **27,59 orders per inhabitant per year** is obtained.

Applying this metric to São José dos Campos population — 697,428 inhabitants according to IBGE 2024 estimates (ESTATÍSTICA, 2024) —, it is estimated that the total *food delivery* market in the city is approximately **19,244,162 orders per year**.

Comparing orders observed in current analysis (94,636 orders/year) with estimated total market (19,244,162 orders/year), it is concluded that current analysis comprises only **0,49%** of the total delivery market of São José dos Campos. This result reflects that the Brendi platform is only one of several existing intermediation platforms in the market,

competing with iFood, Rappi, Uber Eats, and other established platforms.

Maximum Plausible Market Share Estimate

Considering the DaaS service business model, which allows intermediating orders from any restaurant without need for contractual exclusivity, it is possible to estimate a maximum plausible market share superior to current. Different from models that depend on exclusive partnerships, the DaaS service can provide delivery support for all restaurants in the city, regardless of where the order was originally made (iFood, Rappi, restaurant's own app, etc.), as long as delivery is requested to the DaaS service.

This flexible characteristic of the model allows potentially high market penetration, limited mainly by operational factors such as geographic coverage of *droneports*, operational capacity (number of drones), and market acceptance.

To estimate maximum plausible market share, it is considered as reference that iFood, Brazilian market leader, has 25,7% market share (IFOOD; IPEC, 2024). For an established and successful DaaS operation in São José dos Campos, it is considered that a market share of **20% of total market** represents an ambitious but plausible scenario, reflecting significant service penetration in the city.

Therefore, **maximum plausible market share is estimated at 20% of total market**, resulting in potential demand of approximately **3,848,832 orders per year** (20% of 19,244,162 orders/year).

This market estimate will be fundamental for subsequent analyses that will evaluate necessary infrastructure sizing (number of *droneports* and drones) to serve this expanded demand, total OpEx and OpEx per order in the horizontal scale scenario, and economic viability through ROI (return on investment) analysis, considering necessary CapEx and projected cash flows.

Expansion Modeling: Cost Dilution Analysis up to 20% Market Share

To evaluate how market share expansion from 0,49% to 20% affects necessary infrastructure and cost dilution per order, an expansion simulation model was developed presented in *notebook* 15 — `15_expansion_market_share.ipynb`.

The model considers gradual market share expansion through increments of 1% at a time, from current scenario (0,49%) to target market share (20%). For each market share scenario, the model calculates:

- **Infrastructure sizing:** Minimum number of drones and *droneports* necessary to serve expanded demand, using the M/M/c queueing model with 20-minute SLA

and respecting physical limit of 15 drones per *droneport*

- **Human resources:** Operators and technicians necessary according to operation scale (minimum of 4 operators and 2 technicians up to 50 drones, scaling with 0,08 operators/drone and 0,04 technicians/drone above this limit)
- **Batteries:** Total batteries necessary considering 3 batteries per drone to guarantee continuous operational availability
- **CapEx:** Investments in fixed infrastructure (*droneport* structure, charging stations, communication) and variable (drones, batteries, safety equipment, training)
- **OpEx:** Annual operational costs fixed (rental, regulatory licenses) and variable (energy, maintenance, battery replacement, salaries, insurance)

Demand distribution among *droneports* maintains uniform proportions, adjusting dynamically as demand grows and new *droneports* are added when capacity limit is reached.

Figure 4.37 presents a consolidated view of expansion results, showing evolution of infrastructure (drones and *droneports*), total costs (CapEx and OpEx), and cost dilution per order (CapEx and OpEx per order) as a function of market share.

Results demonstrate that market share expansion generates significant dilution of costs per order due to greater operational scale. Table 4.16 presents a detailed comparison of four representative scenarios: current market share (0,49%), intermediate (5,49% and 10,49%), and target (20,49%).

Results evidence significant impacts of expansion on cost structure. Expansion from 0,49% to 20,49% market share — corresponding to a demand increase of approximately **41,8 times** — requires proportionally smaller infrastructure increase: **1,8 times** more *droneports* (from 4 to 7) and **8,8 times** more drones (from 12 to 105). This mismatch between demand growth and infrastructure reflects dilution of fixed costs and operational efficiency gains at scale.

The most significant effect of expansion is **cost dilution per order**: CapEx per order reduces from USD 11,29 to USD 1,79 (**reduction of 84,2%**), while OpEx per order reduces from USD 6,26 to USD 0,67 (**reduction of 89,3%**). This dramatic dilution demonstrates the strong scale effect of the operational model, where substantial fixed costs (*droneport* infrastructure, communication systems, minimum personnel) are distributed over a much larger volume of orders.

Mean system utilization increases significantly with expansion, evolving from 17,1% in current scenario to 77,1% at 5,49% market share, reaching 90,6% at 10,49% and stabilizing at 87,7% in the 20,49% scenario. This utilization increase reflects better utilization of installed capacity and contributes to cost per order reduction through fixed cost dilution.

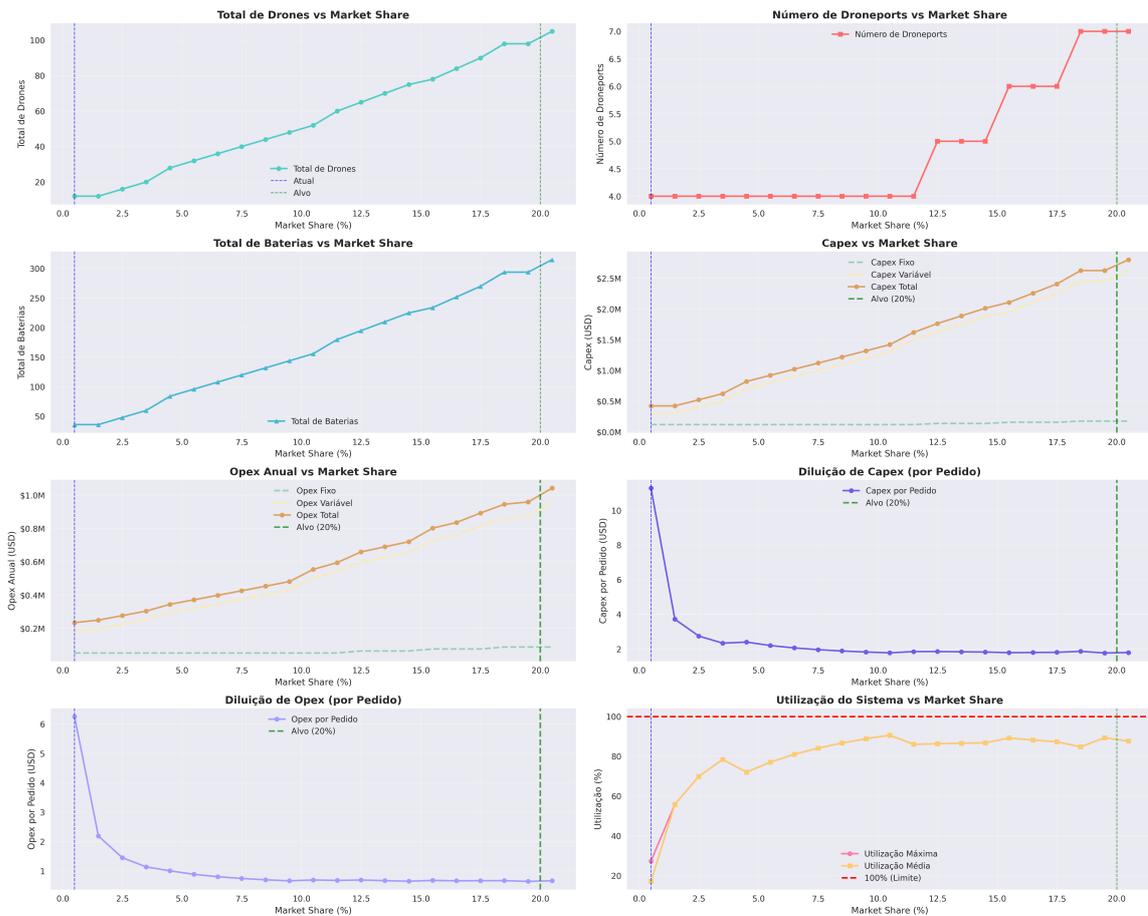


FIGURE 4.37 – Market share expansion modeling: infrastructure and cost dilution analysis. Charts show evolution of (a) total drones and *droneports*, (b) necessary batteries, (c) fixed, variable, and total CapEx, (d) fixed, variable, and total OpEx, (e) CapEx dilution per order, (f) OpEx dilution per order, and (g) mean and maximum system utilization. Dashed vertical lines indicate current market share (0,49%, blue) and target market share (20%, green).

TABLE 4.16 – Market share expansion scenario comparison: infrastructure, total costs, and dilution per order

Metric	0,49%	5,49%	10,49%	20,49%
Demand (orders/year)	37,404	419,077	800,751	1,564,098
Multiplier	1,0x	11,2x	21,4x	41,8x
Infrastructure				
Droneports	4	4	4	7
Drones	12	32	52	105
Batteries	36	96	156	315
Operators	4	4	5	9
Technicians	2	2	3	5
Total Costs				
Total CapEx (USD)	422,200	919,200	1,418,200	2,798,250
Total Annual OpEx (USD)	233,982	370,982	553,883	1,042,034
Dilution per Order				
CapEx per Order (USD)	11,29	2,19	1,77	1,79
OpEx per Order (USD)	6,26	0,89	0,69	0,67
Mean Utilization (%)	17,1	77,1	90,6	87,7

This analysis validates that horizontal expansion to maximum plausible market share of 20% is viable from an operational perspective and presents significant cost dilution benefits. Results indicate that reaching high volumes is fundamental to make the economic model competitive, substantially reducing both CapEx and OpEx per order through scale effect.

4.3.3 Temporal Growth Modeling: Financial Evolution over 5 Years

While the horizontal expansion analysis presented previously explores static scenarios of different market share levels, it is fundamental to understand how temporal business evolution impacts operation financial viability. In this way, an exponential growth model was developed that projects market share evolution over 5 years, starting from current scenario (0,49%) until reaching maximum plausible market share of 20%.

This analysis, performed in *notebook* 16 — `16_crescimento_exponencial_5_anos.ipynb` —, allows evaluating how operational costs (OpEx), incremental investments (incremental CapEx), and revenues evolve over time, providing insights on operation financial health year by year and identifying indicators such as break-even point, return on investment (ROI), and operating margin.

Exponential Growth Model (CAGR)

To model market share growth from 0,49% to 20% in 5 years, an exponential growth model based on compound annual growth rate (CAGR) was used. CAGR represents the constant annual growth rate that would be necessary to achieve final value from initial value over the considered period.

CAGR calculation formula is:

$$\text{CAGR} = \left(\frac{\text{Final Value}}{\text{Initial Value}} \right)^{\frac{1}{\text{Periods}}} - 1 \quad (4.14)$$

Applying this formula for market share growth from 0,49% to 20% in 5 years, a CAGR of **109,97% per year** is obtained, representing a very high but feasible growth rate in the context of an accelerated expansion operation in a growing market.

Table 4.17 presents evolution of market share, annual demand, and multiplier relative to initial scenario over the 5 years.

TABLE 4.17 – Temporal evolution of market share and demand over 5 years

Year	Market Share (%)	Annual Demand	Multiplier
0	0,49	94,296	1,00x
1	1,03	197,998	2,10x
2	2,16	415,747	4,41x
3	4,54	872,962	9,26x
4	9,52	1,832,999	19,44x
5	20,00	3,848,832	40,82x

Evolution shows that, in year 5, annual demand increases to **3,848,832 orders**, representing growth of approximately **40,8 times** relative to initial scenario. This accelerated growth reflects need for significant expansion of operational infrastructure, as discussed below.

Infrastructure and Cost Evolution

For each year of the projection, the model calculates necessary sizing (number of drones, *dronports*, batteries, and human resources) using the same methodology presented in Section 4.3.2, based on the M/M/c queueing model with 20-minute SLA.

The model considers that **incremental CapEx** corresponds only to investments in new resources necessary in each year, not to total accumulated investment. This approach is more realistic for cash flow analysis, as it represents effective disbursements of each year.

Annual OpEx is calculated based on scale of each year, including both fixed costs

(which increase as new *droneports* are installed) and variable costs (which scale with volume of orders processed).

Revenue Model

Revenue model considers a mean delivery price based on iFood reference model, as presented in Section 4.3.1, with addition of a premium that reflects value differential offered by drone delivery service (greater speed and reliability).

According to *notebook* 12 analysis, mean iFood delivery cost is R\$ 11,33 per order. The model adds a **premium of R\$ 4,00 per order**, resulting in a mean delivery price of **R\$ 15,33** (USD 3,07, considering exchange rate of USD 1 = BRL 5).

This premium is justified by value differential offered: deliveries with 20-minute SLA (vs. 30+ minutes of traditional model), lower dependence on traffic conditions, and greater operational reliability, characteristics that can be especially valued by market segments that prioritize speed and punctuality.

Annual revenue for each year is calculated multiplying mean delivery price by projected annual demand.

Financial Evolution and Indicators

Table 4.18 presents detailed financial evolution over the 5 years, including revenue, OpEx, incremental CapEx, free cash flow (FCF), operating result (profit/loss), and operating margin.

TABLE 4.18 – Detailed financial evolution over 5 years

Year	MS (%)	Annual Revenue		OpEx	CapEx	FCF	Op. Margin
		(USD)	(BRL)	(USD)	(USD)	(USD)	(%)
0	0,49	289,161	1,445,807	244,324	422,200	-377,362	15,51
1	1,03	607,166	3,035,828	277,169	99,400	230,596	54,35
2	2,16	1,274,894	6,374,467	370,904	397,600	506,390	70,91
3	4,54	2,676,953	13,384,763	593,386	697,800	1,385,767	77,83
4	9,52	5,620,921	28,104,604	1,175,525	1,616,000	2,829,396	79,09
5	20,00	11,802,507	59,012,534	2,488,211	3,584,750	5,729,545	78,92

Results evidence important patterns in operation financial evolution. In **Year 0** (current scenario), operation presents positive operating result (profit of USD 44,838), but free cash flow is negative (USD -377,362) due to high initial CapEx investment (USD 422,200). Initial operating margin is 15,51%, indicating that the model is viable from an operational perspective even at small scale.

In **Year 1**, with market share of 1,03%, FCF becomes positive (USD 230,596), but

accumulated FCF is still negative (USD -146,766), reflecting that initial investment has not yet been fully recovered.

Break-even point is reached in **Year 2**, when accumulated FCF reaches USD 359,624, indicating that all investments made until then were recovered through positive cash flows from previous years. From this point, operation generates positive return on invested capital.

Subsequent evolution shows accelerated growth: in **Year 5**, with market share of 20% and demand of 3,848,832 orders/year, annual revenue reaches USD 11,802,507 (BRL 59,012,534), while OpEx is USD 2,488,211, resulting in operating result of USD 9,314,295 and operating margin of 78,92%.

Total accumulated CapEx over the 5 years is USD 6,817,750, while **total accumulated FCF** is USD 10,304,332, resulting in an **accumulated ROI of 151,14%** after 5 years. This high ROI demonstrates financial viability of the model when target market share is reached, reflecting the strong scale effect that dilutes fixed costs over a much larger volume of orders.

Figure 4.38 presents a consolidated view of financial evolution, showing main indicators over the 5 years in eight complementary charts: market share evolution, demand growth, revenue evolution, OpEx and incremental CapEx, annual FCF, accumulated FCF (with break-even line), operating margin, and infrastructure evolution (drones and *droneports*).

Considerations on Pricing and Economic Viability

Adopted pricing model considers a mean delivery price of R\$ 15,33 per order, resulting from combination of iFood base cost (R\$ 11,33) with a premium of R\$ 4,00. This pricing strategy was developed to balance two fundamental objectives: maintain competitiveness in the short term and guarantee economic viability in the long term.

In the **short term** (Year 0 and beginning of Year 1), the model may present negative cash flows due to high initial infrastructure investments. However, even in this initial period, operation presents positive operating result (revenue exceeds OpEx), demonstrating that the model is viable from an operational perspective and that adopted pricing is compatible with operational costs even at small scale.

In the **long term** (from Year 2), the model becomes increasingly profitable, with operating margin evolving from 15,51% in Year 0 to approximately 79% from Year 4, demonstrating economic viability of the model when adequate scale is reached. Break-even point is reached in Year 2, when accumulated FCF becomes positive, indicating complete recovery of initial investments.

It is important to highlight that, in practice, delivery costs are **partitioned between**

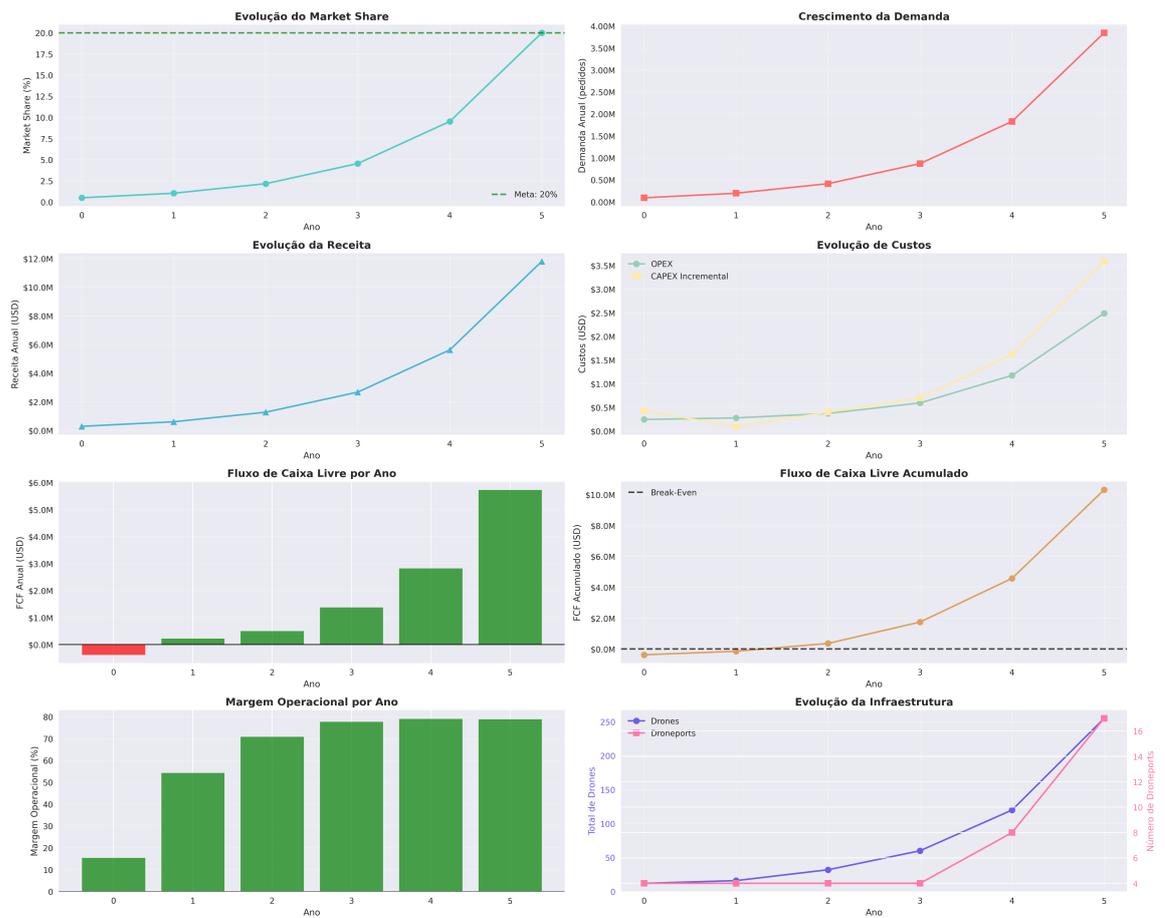


FIGURE 4.38 – Financial and operational evolution over 5 years: (a) market share evolution, (b) demand growth, (c) annual revenue evolution, (d) OpEx and incremental CapEx, (e) annual free cash flow, (f) accumulated free cash flow (with break-even line), (g) operating margin, and (h) infrastructure evolution (drones and *droneports*). Dashed line in chart (f) indicates break-even point (Year 2).

restaurant and consumer, according to common model in the *food delivery* market. A portion of delivery fee is paid directly by consumer (as visible delivery fee in the app), while another portion may be incorporated into product prices by restaurant or charged through commission fees on order. This partitioning allows flexibility in pricing strategy, enabling adjustments for different market segments and different business models (e.g., monthly unlimited deliveries plan vs. per-delivery charge).

The R\$ 4,00 premium over the iFood model can be distributed between restaurant and consumer proportionally, maintaining service competitiveness while guaranteeing economic viability through sufficient revenue to cover operational costs and generate adequate return on invested capital.

This temporal analysis validates that the DaaS business model presents economic viability when considered an adequate planning horizon (5 years), as long as it is possible to achieve projected market share growth through marketing strategies, restaurant partnerships, and market acceptance of drone delivery service (see assumptions spreadsheet in Appendix A).

5 Conclusion

This chapter presents the conclusion of the work, synthesizing the main results obtained, evaluating the research hypothesis formulated in the introduction, highlighting the contributions of the study and pointing out limitations and directions for future work.

5.1 Synthesis of Main Results

This work developed an integrated analysis of logistical, energetic and economic-financial viability of a Drone-as-a-Service (DaaS) service for food deliveries in the city of São José dos Campos, using real data from the Brendi platform. The results obtained demonstrate both technical viability and economic challenges for implementation of the proposed model.

Logistical Results

Processing of the original dataset of 1,048,575 orders resulted in a final set of 23 335 valid orders after geographic filtering and outlier removal, representing successfully delivered orders in the São José dos Campos region. Analysis of the *Distance-Shortening Rate* (DSR) revealed an average value of 0.3524, indicating that the air mode offers potential reduction of approximately 35% in distance traveled relative to the ground mode, validating the hypothesis of significant geometric gain for urban deliveries.

Optimal location of *droneports* through the K-means algorithm resulted in 4 strategic points distributed in the region, capable of covering 85.48% of demand (19,947 orders) within the operational radius of 5 km. Fleet sizing through the M/M/c queue model determined the need for 12 drones distributed among the 4 *droneports*, guaranteeing Service Level Agreement (SLA) of 20 minutes for 95% of orders. This result confirms the operational viability of the proposed model from a logistical perspective.

Economic-Financial Results

The economic-financial analysis revealed a total initial investment (CapEx) of USD 421,200 (BRL 2,106,000), with drone acquisition responsible for 56.98% of total investment. Annual operational costs (OpEx) totaled USD 256,772 (BRL 1,283,744), resulting in OpEx per order of USD 6.86 in the current demand scenario (37,404 orders/year).

The exploratory analysis of vertical scale demonstrated that strategies for better utilization of installed infrastructure can significantly reduce OpEx per order: through demand mirroring (temporal redistribution), OpEx reduces to USD 4.45 (35% reduction), while at theoretical maximum load OpEx reaches USD 3.77 (45% reduction). Comparison with the reference model (iFood Delivery Plan) revealed that the drone model's OpEx per order is higher in the current scenario (USD 6.86 vs USD 2.70), but approaches the ground model with scale strategies.

The horizontal expansion analysis up to 20% market share showed operational viability and significant cost dilution benefits. The temporal growth model projected for 5 years indicated break-even in Year 2, with accumulated ROI of 151.14% in Year 5 and operating margin evolving from 15.51% (Year 0) to approximately 79% (Years 4-5), demonstrating economic viability when adequate scale is reached.

5.2 Response to Hypothesis and Research Problem

The hypothesis formulated in Chapter 1 established that *“a DaaS service dedicated to meal delivery is logistically viable — capable of meeting a minimum Service Level Agreement (SLA) of 20 minutes and reducing total delivery time compared to the traditional motorcycle-based model — and economically viable — with operational costs comparable to or lower than motorcycle delivery, considering the necessary infrastructure, regulatory costs and the operation scale of the Brendi platform”*.

Logistical Viability Assessment

The logistical viability hypothesis was **confirmed**. The M/M/c queue model demonstrated that it is possible to meet the 20-minute SLA for 95% of orders with a fleet of 12 drones distributed in 4 strategically positioned *droneports*. DSR analysis revealed an average geometric gain of 35% in distance traveled, validating the potential for delivery time reduction relative to the ground mode. Coverage of 85.48% of demand within the 5 km operational radius confirms that the proposed infrastructure is adequate to serve a significant portion of the market.

Economic Viability Assessment

The economic viability hypothesis was **partially confirmed**. In the current demand scenario, the drone model's OpEx per order (USD 6.86) is significantly higher than the motorcycle delivery model (USD 2.70), representing a 154% differential. However, through vertical scale strategies — especially demand mirroring and operation at theoretical maximum load — OpEx per order reduces to USD 4.45 and USD 3.77, respectively, approaching the ground model.

It is important to highlight that the drone model offers a value differential through delivery speed (20-minute SLA vs 30+ minutes of the traditional model) and operational reliability independent of traffic conditions. This differential can justify a price premium that compensates for the higher OpEx, especially for market segments that value speed and punctuality. The temporal growth analysis demonstrated that, with adequate market share expansion, the model presents full economic viability, with break-even in Year 2 and accumulated ROI of 151.14% in Year 5.

Response to Research Problem

The research problem formulated: *“Would it be logistically and economically viable to implement, in a Brazilian city, a Drone-as-a-Service (DaaS) service dedicated to meal delivery, taking as a pilot phase the order portfolio already intermediated by the Brendi platform?”*

The answer is: **Yes, logistically viable; economically viable with reservations**. Logistical viability was confirmed through demonstration of operational capacity to meet the proposed SLA with the sized infrastructure. Economic viability requires adequate scale to dilute substantial fixed costs and justify the initial investment. Full economic viability depends on market share growth and operational optimization strategies, as demonstrated by the temporal growth analysis. The model presents potential for positive return when adequate scale is reached, but requires significant initial investment and demand capture strategies to achieve profitability.

5.3 Contributions of the Work

This work contributes to the literature and practice of drone deliveries in multiple ways, standing out for the integration of different analysis dimensions and adaptation to the Brazilian context.

Methodological Contributions

The work integrates logistical (VRP-D), energetic and economic-financial modeling in a unique framework for DaaS viability evaluation, allowing holistic analysis that captures interactions between demand, infrastructure, climate and costs. The adaptation of consolidated methodologies — DSR, K-means, M/M/c — to the Brazilian regulatory context (RBAC-E 94, ICA 100-40) demonstrates how established techniques can be applied considering specific restrictions of each country. The development of a complete pipeline from real data to viability metrics offers a reproducible methodology that can be adapted to other cities and contexts.

Empirical Contributions

The use of real data from the Brendi platform — 23,335 successfully delivered orders — captures specific patterns of the Brazilian market, including geographic distribution of restaurants and customers, temporal demand patterns and characteristics of the urban road network. The specific analysis for São José dos Campos incorporates vehicle flow data and local meteorological conditions (ERA5, INMET), providing contextualized results that reflect Brazilian operational reality. Validation with reference model (iFood Delivery Plan) and comparison with international commercial operators (Speedbird Aero, Wing, Meituan) ensures that results are realistic and comparable with real operations.

Regulatory Contributions

The complete incorporation of RBAC-E 94 and ICA 100-40 restrictions into the models — including maximum takeoff weight, operational ceiling, mandatory safety equipment, SARPAS-NG flight plans and RC-RETA insurance — provides a viability analysis that fully reflects the legal compliance cost. The analysis of regulatory cost impact on economic viability demonstrates how these restrictions affect model competitiveness, offering insights for regulators and operators on the impact of different regulatory policies.

Practical Contributions

The quantitative CapEx/OpEx comparative analysis with motorcycle delivery establishes objective operational viability criteria, allowing stakeholders to evaluate the model based on clear and comparable metrics. The identification of scale strategies (vertical and horizontal) for economic viability offers practical paths for operators to achieve profitability, highlighting the importance of demand growth and operational optimization.

The temporal growth projection with financial indicators (break-even, ROI, operating margin) provides a quantitative basis for investment decisions and strategic planning.

5.4 Limitations and Future Work

It is fundamental to recognize the limitations inherent to this study and identify directions for future research that can refine or extend the presented results.

Recognized Limitations

Cost estimates: The drone acquisition costs used in the analysis are estimates based on benchmarking with similar commercial equipment, not corresponding to official public values. The Keeta G-4 drones used as reference are commercialized through direct sale and not available on the open market, requiring estimates that may differ from real acquisition values.

Queue model: The M/M/c model is adopted as a long-term equilibrium approximation, recognizing that SARPAS-NG operational windows and hourly peaks may introduce non-Poissonian behavior. The need for flight plans with 24-hour advance notice may create arrival patterns that differ from the assumed Poisson process, potentially impacting waiting times and fleet sizing.

Geographic generalization: The analysis was focused exclusively on the city of São José dos Campos, a medium-sized city with specific characteristics of road network, population density and demand patterns. Generalization of results to other Brazilian cities — especially metropolises such as São Paulo and Rio de Janeiro — requires additional validation considering different urban, regulatory and market characteristics.

Meteorological scenario: A conservative scenario of meteorological conditions was adopted to ensure pessimistic analysis, which may underestimate real operational availability. The analysis of three scenarios (conservative, moderate and aggressive) offers variability, but the choice of the conservative scenario may not reflect average operational conditions throughout the year.

Pricing model: The revenue model assumes a fixed premium of R\$ 4.00 over iFood's base cost, not exploring dynamic pricing strategies, market segmentation or different business models (e.g., monthly unlimited delivery plan vs per-delivery charge). Consumer willingness to pay for drone delivery service was not empirically investigated.

Safety protocols and physical delivery operationalization: The model assumes that the drone arrives at the customer and delivery is complete, without considering actual

procedures of drone descent, physical package delivery to the customer, safety protocols on landing, customer interaction, procedures in case of customer absence, and other operational aspects of the final delivery stage. This simplification may impact both delivery times and operational costs, especially if there is need for support personnel on site or automated delivery systems (retractable cables, secure compartments, dedicated delivery stations). The analysis did not consider different physical delivery modalities — such as direct delivery to customer, delivery in secure compartment, or hybrid delivery with motorcycle completing the last meters — which may have significant implications for operational and economic viability.

Future Work

Validation in other cities: Extending the analysis to other Brazilian cities with different urban characteristics — especially metropolises such as São Paulo, Rio de Janeiro and Belo Horizonte — would allow validating generalization of results and identifying specific factors of each context that impact model viability.

Sensitivity analysis: Conducting more detailed sensitivity analysis of cost parameters — especially drone acquisition costs, electricity tariffs, insurance premiums and infrastructure costs — would allow identifying which components most impact economic viability and establishing critical value ranges.

Advanced queue models: Incorporating more sophisticated queue models that capture non-Poissonian behavior resulting from SARPAS-NG operational windows, seasonal demand patterns and intradaily variations would allow more precise fleet sizing and better waiting time estimation.

Risk analysis: Developing risk analysis and probabilistic scenarios for financial projections — incorporating uncertainties in demand, operational costs, meteorological conditions and market acceptance — would offer a more robust viability evaluation considering variability and risks.

Pricing strategies: Investigating dynamic pricing strategies and market segmentation — including consumer willingness to pay analysis, demand elasticity and different business models — would allow optimizing revenues and improving model economic viability.

Environmental impact: Conducting comparative environmental impact analysis — especially CO₂ emissions and other pollutants — between the drone model and the motorcycle delivery model, considering complete equipment life cycle and energy sources, would offer an additional dimension of model sustainability evaluation.

Route optimization: Integrating machine learning models for real-time route optimization — considering dynamic meteorological conditions, traffic patterns and drone

availability — would allow better utilization of installed capacity and additional operational cost reduction.

Consumer acceptance: Conducting empirical study of consumer acceptance and willingness to pay for drone delivery service — through surveys, focus groups or choice experiments — would provide fundamental data for pricing and marketing strategies.

Physical delivery protocols: Developing detailed analysis of different physical delivery modalities — including direct delivery to customer, automated delivery systems, and hybrid models — considering operational costs, delivery times, safety requirements and customer experience, would allow identifying the most viable modality for different contexts.

Network analysis: Extending the analysis to multiple cities simultaneously, considering operational synergies, infrastructure sharing and economies of scale in centralized management, would allow evaluating network operation viability and geographic expansion strategies.

In summary, this work demonstrated that a Drone-as-a-Service for food deliveries is logistically viable and presents potential for economic viability when adequate scale strategies and market share growth are considered. The results offer a solid quantitative basis for investment decisions and strategic planning, while identifying paths for future research that can refine and extend the presented analyses.

Bibliography

AÉREO, D. de Controle do E. **ICA 100-40 – Regras do Espaço Aéreo e SARPAS-NG**. [S.l.], 2022. Instrução do Comando da Aeronáutica sobre regras do espaço aéreo e Sistema de Autorização de Rastreamento de Produtos Aeronáuticos.

ARTHUR, D.; VASSILVITSKII, S. k-means++: The advantages of careful seeding. *In: Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms. Proceedings* [...]. SIAM, 2007. p. 1027–1035. Available at: <https://dl.acm.org/doi/10.5555/1283383.1283494>.

BINE, T.; MARQUES, J.; OCHIENG, W. Drone delivery: Why, where, and when? quantifying the distance–shortening rate in urban networks. *In: Proceedings of the 20th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM '23). Proceedings* [...]. Association for Computing Machinery, 2023. p. 221–229. Artigo que propõe o índice Distance-Shortening Rate (DSR) utilizado neste trabalho. Available at: <https://www.updwg.org/wp-content/uploads/2023/11/Drone-Delivery-Where-when.pdf>.

BONACICH, P. Power and centrality: A family of measures. **American Journal of Sociology**, v. 92, n. 5, p. 1170–1182, 1987.

BRASIL, D. **Sistema de segurança para drones BVLOS**. 2023. <https://www.dronecenter.com.br/paraquedas-bvlos/>.

CIVIL, A. N. de Aviação. **Primeira autorização BVLOS emitida no Brasil sob RBAC-E 94**. [S.l.], 2023. Primeira certificação de aeronavegabilidade especial (CAER) para operações BVLOS no Brasil.

CIVIL, A. N. de Aviação. **RBAC-E n^o 94, Emenda 00 – Operações com Aeronaves não Tripuladas**. [S.l.], 2023. Regulamento Brasileiro da Aviação Civil – Especial – Operações com Aeronaves não Tripuladas.

CONCEIÇÃO, F. R.; RODRIGUES, D. N.; SANTOS, F. L. A branch-and-price heuristic for the truck-and-drone routing problem with time windows. **Transportation Research Part C: Emerging Technologies**, v. 142, p. 103749, 2022. Available at: <https://doi.org/10.1016/j.trc.2022.103749>.

CURBFLOW. **Relatório de Impacto de Fila-Dupla no Centro de São Paulo**. 2023. <https://curbflow.com/reports/paulista-double-parking-2023.pdf>. Dados mostram que motofretes respondem por mais de 20% das paradas em fila dupla.

DORLING, K.; HEINRICHS, J.; MESSIER, G.; MAGIEROWSKI, S. Vehicle routing problems for drone delivery. **IEEE Transactions on Systems, Man, and Cybernetics: Systems**, v. 47, n. 1, p. 70–85, 2017.

ESTATÍSTICA, I. B. de Geografia e. **Estimativas da População dos Municípios Brasileiros**. 2024. <https://www.ibge.gov.br/estatisticas/sociais/populacao/9103-estimativas-de-populacao.html>. Estimativas populacionais dos municípios brasileiros com data de referência em 1^o de julho de 2024. Valores consultados: São Paulo (12.325.232 hab.), Rio de Janeiro (6.747.815 hab.), Curitiba (1.963.726 hab.), São José dos Campos (750.000 hab.), Montes Claros (418.000 hab.).

FIGLIOZZI, M. A.; WADE, M.; ROUMBOUTSOS, A. Case study of drone delivery reliability for time-sensitive medical products. **Transportation Research Record**, v. 2675, n. 1, p. 608–618, 2021.

FILIOPOULOU, M.; AL. et. Energy-aware routing of delivery drones under wind uncertainty. **Economics of Transportation**, v. 34, p. 100385, 2025.

FORUM, W. E. **The Future of the Last-Mile Ecosystem in Brazil**. Genebra: [s.n.], 04 2024. Relatório técnico que estima o faturamento do delivery brasileiro em US\$ 1,3 bi. Available at: <https://www.weforum.org/reports/future-of-last-mile-ecosystem-brazil>.

GROSS, D.; SHORTLE, J. F.; THOMPSON, J. M.; HARRIS, C. M. **Fundamentals of Queuing Theory**. 5. ed. Hoboken: Wiley, 2018. ISBN 978-1119450889.

GROUP, Y. **Meituan's Drone Service Takes Flight: Over 100 000 Orders Delivered in 2022**. 2022. <https://www.yolegroup.com/industry-news/meituan-drone-service-takes-flight-over-100000-orders-delivered-in-2022/>. Dados de escala da operação de drones em Shenzhen.

HAN, J.; LIU, Y.; LI, Y. Vehicle routing problem with drones considering time windows and dynamic demand. **Applied Sciences**, MDPI, v. 13, n. 24, p. 13086, 2023. Available at: <https://doi.org/10.3390/app132413086>.

IFOOD. **Portal de Transparência**. 2025.

<https://institucional.ifood.com.br/transparencia/>. Dados oficiais do iFood indicando 120 milhões de pedidos intermediados por mês em 2025 (1,44 bilhão por ano).

IFOOD; IPEC, I. **Pesquisa Ipsos-Ipec: Market Share do iFood no Brasil**. 2024. <https://institucional.ifood.com.br/estudos-e-pesquisas/pesquisa-ipsos-ipecc/>. Pesquisa indica que o iFood possui 25,7% de market share no mercado brasileiro de delivery em 2024.

KENDALL, D. G. Stochastic processes occurring in the theory of queues and their analysis by the method of the imbedded markov chain. **The Annals of Mathematical Statistics**, v. 24, n. 3, p. 338–354, 1953.

LEISHMAN, J. G. **Principles of Helicopter Aerodynamics**. 2. ed. [S.l.]: Cambridge University Press, 2006.

LI, Y.; LIU, Y.; HAN, J. Managing price and fleet size for courier service with shared drones. **Omega**, v. 101, p. 102301, propõe a formulação TW-VRP-D adotada neste trabalho, 2021. Available at: <https://doi.org/10.1016/j.omega.2020.102301>.

LLOYD, S. P. Least squares quantization in pcm. **IEEE Transactions on Information Theory**, v. 28, n. 2, p. 129–137, 1982.

OES, A. N. de T. **Resolução nº 715/2019**. [S.l.], 2019. Regulamentação sobre licenciamento de estações de rádio para aeronaves não tripuladas.

PRIVADOS, S. de S. **Circular nº 637/2021 – Seguro RC-RETA**. [S.l.], 2021. Regulamentação sobre seguro de responsabilidade civil para aeronaves remotamente pilotadas.

SEGUROS, M. **Tabela de prêmios para seguro de aeronaves não tripuladas (drones)**. 2024. <https://www.mapfre.com.br/drones-seguro/>.

SILVA, C. C.; OLIVEIRA, A. P.; CORRÊA, S. M. Comparative emission rates of motorcycles and light-duty vehicles in rio de janeiro. **Journal of the Air & Waste Management Association**, v. 63, n. 5, p. 550–556, 2013.

SINNOTT, R. W. Virtues of the haversine. **Sky and Telescope**, v. 68, n. 2, p. 159, apresenta a derivação moderna da fórmula de Haversine para distâncias geodésicas, 1984.

SIVRIOGLU, S.; TEMIZ, V. Energy consumption modeling of rotary-wing uavs for forward flight. **Aerospace**, v. 7, n. 10, p. 140, 2020.

SP, I. **Boletim Estatístico de Acidentes de Trânsito – Setembro 2024**. 10 2024. <https://www.infosiga.sp.gov.br/boletins/2024/09>. Registra 1 925 mortes de motociclistas em SP entre jan–set 2024.

VINCENT, J. Walmart and wing expand drone delivery to 100 stores. **Wired**, reportagem descreve mil entregas diárias e plano de expansão, 06 2025. Available at: <https://www.wired.com/story/walmart-wing-expand-drone-delivery/>.

WANG, X.; POIKONEN, S.; GOLDEN, B. The vehicle routing problem with drones: Several worst-case results. **Networks**, v. 73, n. 4, p. 409–426, apresenta o algoritmo Savings + Split usado como heurística, 2019. Available at: <https://doi.org/10.1002/net.21911>.

FOLHA DE REGISTRO DO DOCUMENTO

1. CLASSIFICAÇÃO/TIPO TC	2. DATA 07 de novembro de 2025	3. DOCUMENTO Nº DCTA/ITA/TC-072/2025	4. Nº DE PÁGINAS 133
5. TÍTULO E SUBTÍTULO: Operational and Economic Viability of a Drone-as-a-Service for Food Delivery: Case Study with Real Order Data from Brazil			
6. AUTOR(ES): João Paulo Penna			
7. INSTITUIÇÃO(ÕES)/ÓRGÃO(S) INTERNO(S)/DIVISÃO(ÕES): Instituto Tecnológico de Aeronáutica – ITA			
8. PALAVRAS-CHAVE SUGERIDAS PELO AUTOR: Drone Delivery; Last-Mile Logistics; Economic viability; Entrega por drones; Logística de última milha; Viabilidade econômica			
9. PALAVRAS-CHAVE RESULTANTES DE INDEXAÇÃO: Drone Delivery; Last-Mile Logistics; Economic viability; Entrega por drones; Logística de última milha; Viabilidade econômica; Engenharia Aeroespacial			
10. APRESENTAÇÃO: <input checked="" type="checkbox"/> Nacional <input type="checkbox"/> Internacional ITA, São José dos Campos. Curso de Graduação em Engenharia Aeroespacial. Orientador: Maj. Lucas Oliveira Barbacovi. Publicado em 2025.			
11. RESUMO: <p>Brazil’s rapidly expanding food-delivery market, with over \$12,8 billion in 2024 and expected to grow to \$27,9 billion in 2030, with a annual growth rate of 22%(FORUM, 2024), exposes the drawbacks of a motorcycle-centric logistics model that intensifies congestion, emissions and accidents. Recent regulatory advances — notably amendments to RBAC-E 94 and Brazil’s first BVLOS airworthiness certificate — now permit commercial drone operations, motivating a comprehensive assessment of a Drone-as-a-Service (DaaS) solution for urban last-mile delivery. This study draws on 1.05 million real orders from the Brendi platform and deploys an end-to-end pipeline: data geocoding and cleansing, computation of air (Haversine) and street distances, derivation of a Distance-Shortening Rate, placement of droneports through K-means and eigenvector centrality, fleet sizing with an M/M/c queue, and segment-level energy modelling that factors in wind and rainfall from ERA5 and INMET. These operational outputs feed a five-year discounted-cash-flow framework whose CapEx and OpEx explicitly incorporate Brazilian compliance costs.</p>			
12. GRAU DE SIGILO: <input checked="" type="checkbox"/> OSTENSIVO <input type="checkbox"/> RESERVADO <input type="checkbox"/> SECRETO			