



Master's Degree Program in **Information Management**

Specialization in
Business Intelligence

Dashboards for Urban Incident Reporting: Improving Decision Support with Power BI Using Na Minha Rua Lx Data

João Pedro Brito Ramos

NOVA Information Management School
Instituto Superior de Estatística e Gestão de Informação
Universidade Nova de Lisboa

**Dashboards for Urban Incident Reporting: Improving Decision Support with
Power BI Using Na Minha Rua Lx Data**

by

João Pedro Brito Ramos

Master Thesis presented as partial requirement for obtaining the Master's degree in
Information Management, with a specialization in Business Intelligence

Supervised by

Fátima Trindade Neves, Ph.D., NOVA Information Management School

February, 2026

STATEMENT OF INTEGRITY

I hereby declare having conducted this academic work with integrity. I confirm that I have not used plagiarism, any form of undue use of information or falsification of results along the process leading to its elaboration. I further declare that I have fully acknowledged the Rules of Conduct and Code of Honor from the NOVA Information Management School.

26/02/2026

João Pedro Brito Ramos

DEDICATION

To my parents, for their unconditional love, sacrifices, and belief in my education. To my girlfriend, for her patience, encouragement, and for sharing every up and down of this journey. To my friends, for their constant support, humour, and timely distractions. And to my supervisor, for her guidance, trust, and invaluable feedback.

This dissertation is dedicated to you.

ACKNOWLEDGEMENTS

I would first like to express my deepest gratitude to my supervisor, Fátima Trindade Neves, for her guidance, availability, and constructive feedback throughout this dissertation. Her critical questions and encouragement were essential to refining both the research design and the final artefact.

I am also grateful to NOVA IMS for providing the academic environment and resources that made this work possible. I would like to thank the team at Câmara Municipal de Lisboa and the Na Minha Rua LX initiative for making incident reporting data openly available and for the broader commitment to transparency and open government data that underpins this study.

Special thanks are due to my classmates and colleagues who shared ideas, challenged assumptions, and provided practical suggestions during the development of the dashboards and the empirical study. Their feedback helped to align the artefact with real decision support needs.

I am profoundly thankful to my parents for their continuous support, values, and encouragement throughout all stages of my education. To my girlfriend, thank you for your patience, understanding, and motivation during the most demanding periods of this work. I am equally grateful to my friends, whose humour and friendship provided balance and perspective whenever it was most needed.

To all of you, my sincere thanks.

ABSTRACT

This dissertation addresses the translation gap between raw citizen-reported urban incident data and structured decision support in municipal governance. Focusing on Lisbon's Na Minha Rua LX platform, it investigates how Business Intelligence dashboards, grounded in disciplined dimensional modelling, can enhance decision support for managing urban incident reporting. Guided by the Design Science Research paradigm and informed by the Kimball methodology, the study designs and implements a dimensional data warehouse and a governed Power BI semantic layer. A declared analytical grain and a clear bus matrix with conformed dimensions (Date, Location, Incident Type, and Channel) ensure semantic consistency across dashboards. Data pipelines implemented in Microsoft Fabric follow a Medallion (Bronze–Silver–Gold) architecture to guarantee lineage, data quality, and traceability. At the core, a star schema supports interpretable measures for daily incident volumes, seasonal baselines, anomaly detection, territorial pressure, and weather sensitivity. Empirical analysis reveals structured regularities in Lisbon's incident reporting dynamics, including stable seasonal cycles, persistent territorial concentration, category-level asymmetries, and directional associations with rainfall in selected domains. These patterns are interpreted descriptively rather than causally and illustrate how dimensional discipline enables consistent detection and contextualisation of analytical signals. Evaluation combines scenario-based walkthroughs and artefact assessment criteria to examine business coverage, interpretability, traceability, semantic consistency, and usability. Findings indicate that improvements in decision support arise primarily from structural clarity rather than analytical complexity. The study demonstrates that a governed semantic layer and a conformant dimensional architecture can transform open citizen-reporting data into reproducible, interpretable, and action-oriented dashboards without relying on opaque predictive models. The dissertation contributes a reusable dimensional representation of urban incident reporting, documented implementation patterns that connect governance principles to analytical design, and empirical evidence that semantic discipline is a foundational enabler of reliable decision support in smart-city contexts.

KEYWORDS

Business Intelligence (BI); Dimensional Data Warehousing; Urban Incident Reporting; Decision Support Systems (DSS); Smart Cities; Microsoft Power BI

Sustainable Development Goals (SDG):



TABLE OF CONTENTS

- Statement of Integrity ii
- Dedication..... iii
- Acknowledgements iv
- Abstract v
- List of Figures viii
- List of Tables ix
- List of Abbreviations and Acronyms x
- 1. Introduction 1
- 2. Literature review 3
 - 2.1 Purpose of the Approach 3
 - 2.2 Citizen Reporting Platforms and Open Government Data (OGD) 3
 - 2.3 Dashboards for Decision Support in the Public Sector 3
 - 2.4 Baseline Monitoring and Anomaly Detection for Municipal Workloads 4
 - 2.5 Kimball as the Methodological Lens for Data Architecture 4
 - 2.6 Operationalizing Kimball with the Medallion Architecture in Microsoft Fabric 5
- 3. Methodology 7
 - 3.1 Research design and methodological stance 7
 - 3.2 Business needs and questions 8
 - 3.3 Data Description 11
 - 3.4 Dimensional design with Kimball and Medallion 12
 - 3.5 Measure design, analysis, validity, and ethics 13
- 4. Empirical Study 15
 - 4.1 Overview of the dashboard artefact and shared layout 15
 - 4.1.1 Shared filter panel and semantic consistency 15
 - 4.2 Homepage - Overview and navigation 16
 - 4.2.1 Objectives of the homepage 17
 - 4.2.2 Visual composition 17
 - 4.3 Page “Workload & Spatial Distribution” (BN1) 18
 - 4.3.1 Objective and link to BN1 18
 - 4.3.2 Visual composition and interactions 19
 - 4.3.3 Analytical use 20
 - 4.4 Page “Seasonality & Baselines” (BN2) 21
 - 4.4.1 Objective and link to BN2 21
 - 4.4.2 Visual composition and interactions 21

4.4.3 Analytical use	22
4.5 Page “Anomalies & Hotspots” (BN3)	23
4.5.1 Objective and link to BN3	23
4.5.2 Visual composition and interactions	23
4.5.3 Analytical use	24
4.6 Page “Territory Prioritization Categories & Hotspots” (BN4).....	25
4.6.1 Objective and link to BN4	25
4.6.2 Visual composition and interactions	25
4.6.3 Analytical use	27
4.7 Page “Weather Context” (BN5)	27
4.7.1 Objective and link to BN5	27
4.7.2 Visual composition and interactions	28
4.7.3 Analytical use	29
5. Results and Discussion	31
5.1 Scenario-Based Evaluation	31
5.1.1 Scenario 1 – City-wide workload in 2024 (BN1)	32
5.1.2 Scenario 2 – Is 2024 unusually busy? (BN2)	34
5.1.3 Scenario 3 – Second-semester anomaly wave (BN3).....	35
5.1.4 Scenario 4 – Prioritising Saneamento and Adding Weather Context (BN4 & BN5) 38	
5.2 Principal Patterns in Na Minha Rua LX Incident Reporting	41
5.2.1 Temporal Patterns.....	41
5.2.2 Territorial Distribution	41
5.2.3 Typological Concentration.....	42
5.2.4 Weather Sensitivity	42
5.2.5 Synthesis of Empirical Patterns.....	42
5.3 Discussion	43
6. Conclusions and Future Research	46
References	48
Appendix A	52

LIST OF FIGURES

- Figure 5-1 - BN1 - Workload & Spatial Distribution..... 32
- Figure 5-2 - BN1 - Workload & Spatial Distribution Tooltip 33
- Figure 5-3 - BN2 - Seasonality & Baselines..... 34
- Figure 5-4 - BN2 - Seasonality & Baselines selected 35
- Figure 5-5 - BN3 - Anomalies & Hotspots..... 36
- Figure 5-6 - BN3 - Anomalies & Hotspots Selected..... 37
- Figure 5-7 - BN4 - Territory Prioritisation Categories & Hotspots 39
- Figure 5-8 - BN5 - Weather Context..... 40
- Figure A-1 - Homepage: Overview and navigation 57
- Figure A-2 - BN1 page: Workload and spatial distribution 58
- Figure A-3 - BN2 page: Seasonality and baseline monitoring 58
- Figure A-4 - BN3 page: Anomalies and hotspots..... 59
- Figure A-5 - BN4 page: Territory prioritisation (pressure and ranking) 59
- Figure A-6 - BN5 page: Weather context (rain indicators and comparative rates).... 60

LIST OF TABLES

Table 5-1 - Scenario-based walkthroughs: mapping business questions, evaluation criteria, and observable evidence (BN1–BN5)..... 31

Table 5-2 - Summary of main empirical findings..... 43

Table A-1 - Operational criteria for evaluating “improving decision support.” 52

Table A-2 - DimDate attribute dictionary..... 53

Table A-3 - DimLocation attribute dictionary..... 53

Table A-4 - DimIncidentType attribute dictionary..... 54

Table A-5 - FactIncidentDaily schema 54

Table A-6 - FactIncidentDaily Measures..... 55

Table A-7 - FactWeatherDaily schema 56

Table A-8 - FactWeatherDaily Measures..... 56

Table A-9 - Bus matrix (processes × dimensions) 57

LIST OF ABBREVIATIONS AND ACRONYMS

API – Application Programming Interface

BI – Business Intelligence

BK – Business Key

BN1–BN7 – Business Needs 1 to 7

BQx.y – Business Question x.y (e.g., BQ1.1, BQ2.3)

DAX – Data Analysis Expressions

DSR – Design Science Research

ELT – Extract, Load, Transform

ETL – Extract, Transform, Load

KPI – Key Performance Indicator

LY – Last Year

NMRLX – Na Minha Rua LX

OGD – Open Government Data

PK – Primary Key

SCD – Slowly Changing Dimension

SLA – Service Level Agreement

SK – Surrogate Key

YTD – Year to Date

YoY – Year-over-Year

1. INTRODUCTION

Urban incident reporting is a crucial component of data-informed municipal governance, as it facilitates efficient management and resolution of urban challenges by leveraging technology and citizen engagement (Volinz et al., 2025). In Lisbon, Na Minha Rua LX, the local counterpart to FixMyStreet channels citizen-generated reports on potholes, street lighting, waste, and related issues, thereby strengthening transparency, responsiveness, and civic engagement in city management (Meijer & Bolívar, 2016). While these platforms produce substantial volumes of open data, their systematic use for routine decision support in public administration remains comparatively underexplored, motivating a design-oriented approach that bridges raw reporting and operational insight (Gonçalves et al., 2024).

Against this backdrop, the present dissertation proposes a Business Intelligence artefact, Microsoft Fabric dataflow, and a Power BI semantic model serving interactive dashboards to transform Na Minha Rua LX reports into actionable knowledge for municipal managers in Lisbon. The general problem addressed is the translation gap between raw incident logs and timely, standardized planning signals. Although dashboards are widely adopted in organizations, their application to the specific context of smart city incident management remains limited and insufficiently studied, as highlighted by prior analyses of citizen reporting ecosystems (Meijer & Bolívar, 2016). The artefact explicitly targets spatial-temporal patterns (trend, seasonality, hotspots), prioritization, and operational follow-up through interpretable measures and traceable drilldowns.

The central research question guiding the study is: *How can Power BI dashboards built from Na Minha Rua LX data improve decision support for managing urban incident reporting in Lisbon?* From this, three objectives follow: (i) to identify the most relevant indicators and dimensions for incident analysis in the Lisbon context; (ii) to design and implement interactive dashboards in Power BI that provide clear visualization and drill-down capabilities for municipal stakeholders; and (iii) to evaluate the extent to which such dashboards enhance decision support for understanding, monitoring, and prioritising urban incidents in practice.

To operationalise the claim of “improving decision support”, the dissertation evaluates the artefact against five explicit criteria aligned with public-sector dashboard requirements and decision support evaluation practice. These criteria coverage of business questions (C1), interpretability (C2), traceability (C3), semantic consistency (C4), and actionability (C5)—are defined in Table A-1 (see Appendix A) and guide the artefact-centric evaluation through scenario-based walkthroughs in the empirical chapters (Matheus et al., 2018; OECD, 2019).

Methodologically, the dissertation follows a Design Science Research (DSR) methodology (Peffer et al., 2007): the study builds a dashboard as an artefact for a defined organisational problem and evaluates it against explicit criteria in its target

environment. In this research design, Design Science Research provides the overarching research strategy, guiding the build–evaluate–reflect cycle of the artefact. Kimball’s dimensional modelling approach, in turn, serves as the methodological design framework for structuring the data warehouse and semantic model. The Medallion architecture operationalises this design within Microsoft Fabric (Kimball & Ross, 2013).

The study builds upon a dimensional data warehousing architecture grounded in the Kimball tradition of Business Intelligence solution design. Practically, the work is organized around (a) scoping and standardization of the business process via a bus matrix that enumerates processes, facts, and conformed dimensions; (b) modelling of star-schemas (fact tables for incident events and service workflows; shared dimensions for time, geography, incident type, priority, and channel); and (c) ETL (Extract Transform and Load) pipelines in Microsoft Fabric to populate and govern these models before exposure in a Power BI semantic layer. This approach prioritizes conforming dimensions for cross-process analysis, surrogate keys for historical tracking (where relevant), and clear grain in fact design to preserve additivity and drill-down/up consistency. Evaluation follows an artefact-centric logic: utility, effectiveness, and interpretability of the dimensional model and dashboards are assessed against explicit criteria and stakeholder-oriented tasks.

The result is a set of validated dashboards that demonstrate how citizen-reported incident data can be transformed into actionable insights for municipal operations (Schiff, 2025). Practical contributions include a replicable BI solution that supports transparency, responsiveness, and informed resource allocation; theoretical contributions include strengthening the discourse on dashboard design and BI in the public sector, especially in the underexplored niche of citizen-reporting platforms and reflecting on the transformational dimension through which information technology yields organizational value (Gregor et al., 2006).

The remainder of this dissertation is structured as follows. Chapter 2 reviews the relevant literature on citizen-reporting platforms, decision-support dashboards in the public sector, baseline-based monitoring and anomaly detection, and dimensional data warehousing, establishing the conceptual and methodological foundations of the study. Chapter 3 presents the research design and methodological stance, detailing the Kimball-based dimensional architecture, the Medallion data-engineering pipeline in Microsoft Fabric, and the semantic model implemented in Power BI. Chapter 4 describes the dashboard artefact, including its layout, interaction logic, and alignment with the defined business needs. Chapter 5 reports the empirical evaluation through scenario-based walkthroughs that demonstrate coverage, interpretability, traceability, and decision value. Finally, Chapter 6 concludes by summarising the contributions, discussing limitations, and outlining directions for future research.

2. LITERATURE REVIEW

2.1 PURPOSE OF THE APPROACH

This review positions the dissertation at the intersection of three strands: (i) dashboard-centric decision support for municipal operations; (ii) monitoring-oriented time series analysis using seasonal baselines to contextualise variation at daily and weekly decision cadences; and (iii) prescriptive alerting that translates baseline deviations and anomaly signals into actionable managerial attention. The aim is to critically synthesize contributions, tensions, and gaps that directly inform the design of a dimensional model and the surrounding data architecture in Microsoft Fabric and Power BI. Following the program's guidance, the review is thematic and explicitly connects the state of the art to this study's research question, objectives, and evaluation choices (grain definition, conformance, and operational criteria). This chapter, therefore, prioritizes syntheses over summaries, establishes why specific design requirements are necessary, and concludes with a research gap that motivates the artefact.

2.2 CITIZEN REPORTING PLATFORMS AND OPEN GOVERNMENT DATA (OGD)

Citizen reporting platforms such as FixMyStreet demonstrate the promise of open government data for transparency, civic engagement, and more responsive municipal services (Pak et al., 2017). However, the literature consistently documents a translation gap between OGD (Open Government Data) and its routine operational use, arising from reporting biases, heterogeneous taxonomies, duplicate records, time lags in updates, and uneven geographic coverage. These characteristics complicate the construction of comparable and auditable indicators and underscore the need for careful data architecture with grain clarity, conformed dimensions, and traceable lineage from source to decision (Matheus et al., 2018; OECD, 2019). Empirical evidence from Brussels points to both high volumes and persistent data quality challenges in FixMyStreet datasets, reinforcing the argument for robust governance and standardization before analytics (Brussels-Capital Region, 2024). In Lisbon, Na Minha Rua LX has published incident data since 2017 (Câmara Municipal de Lisboa, 2026) at the statistical subsector level with thematic typologies; this provides a strong empirical setting for municipal analytics and public dashboards, and it enables grain-consistent integration with exogenous sources (e.g., daily weather) for spatio-temporal analysis and workload monitoring. The program's methods guidance further supports explicitly linking such platform characteristics to dimensional modelling choices (e.g., the need for a stable Location dimension with surrogate keys) (Matheus et al., 2018; OECD, 2019).

2.3 DASHBOARDS FOR DECISION SUPPORT IN THE PUBLIC SECTOR

Studies on public sector dashboards show that transparency and accountability gains materialize when solutions are grounded in clear measurement logic: (i) seasonal baselines that contextualize variation; (ii) interpretable anomaly criteria; (iii) drill-downs that preserve context across time and space; and (iv) data lineage to trace metrics

back to sources (Matheus et al., 2018). Policy frameworks for a data-driven public sector likewise emphasize data governance, quality management, and reuse to create public value (OECD, 2019). For this dissertation, these prescriptions translate into requirements for the semantic layer: expose week consistent baselines; document anomaly rules in the metric layer; and preserve end-to-end lineage from Power BI visuals to Fabric dataflows. This design stance aligns with writing guidance on clarity, concision, and explicit linkage between literature, methods, and evaluation criteria (Matheus et al., 2018; OECD, 2019).

2.4 BASELINE MONITORING AND ANOMALY DETECTION FOR MUNICIPAL WORKLOADS

At daily grain, municipal incident workloads typically exhibit weekly seasonality, episodic bursts during adverse weather, and non-stationary behaviour driven by infrastructure, reporting practices, and policy changes. In operational settings, a key requirement is not to predict future counts, but to contextualise observed workloads against transparent expectations that decision makers can audit and discuss. For this reason, the analysis emphasises seasonal baselines such as same-week references and rolling averages that exclude the current period as interpretable monitoring benchmarks (Matheus et al., 2018). These baselines enable managers to distinguish routine seasonal peaks from genuinely abnormal surges, while preserving traceability from dashboards to the underlying incident facts.

Weather is treated as contextual information rather than as a predictive regressor. In an urban infrastructure setting, precipitation may plausibly relate to incident generation (e.g., drainage, surface deterioration, and illegal dumping patterns), but the dissertation intends to support explanation and prioritisation rather than forecast accuracy. Accordingly, weather is incorporated through transparent comparative indicators (e.g., rainy versus dry incidence rates and “rainy lift”) that help stakeholders interpret workload changes under comparable baselines (Arifani et al., 2025).

Finally, anomaly detection is framed as monitoring-based alerting: anomalies are operationalised as deviations from defined baselines (percentage deviation and Z-scores) and used to guide attention to days and parishes where workloads persistently exceed expectations. In this dissertation, the value of anomaly signals is therefore assessed through their interpretability, traceability, and usefulness within the scenario-based walkthroughs, rather than through predictive performance metrics (Matheus et al., 2018).

2.5 KIMBALL AS THE METHODOLOGICAL LENS FOR DATA ARCHITECTURE

This study adopts Kimball’s dimensional data warehousing as the foundation for analytical design, drawing on the principles articulated by (Kimball & Ross, 2013). Four principles structure the literature to design translation:

The bus matrix enumerates business processes, fact tables, and conformed dimensions, enforcing consistency across analytical scopes (e.g., Incidents Events) and enabling reuse as new processes are added. In a citizen incident context, it formalizes shared Date, Location, Incident Type, Priority, and Channel, each with a stable semantic role and common keys.

Star Schemas with clear grain. Fact design uses daily grain for Incident analytics, aligning with the decision cadence and the resolution of weather data. Grain clarity preserves additivity, prevents double-counting in drill-down/up, and supports standard-week baselines and year-over-year (YoY) comparators without ad hoc logic.

Conformed Dimensions, Surrogate Keys, and SCD. Conformed dimensions ensure comparability across processes and time. Surrogate keys decouple analytics from source volatility; Slowly Changing Dimensions (SCD) capture legitimate historical changes (e.g., boundary shifts, typology reclassifications) without corrupting prior aggregates.

ETL in Microsoft Fabric prior to the semantic layer. Transformations run in ETL pipelines so that Power BI consumes a documented, governed model with explainable measures, reliable time intelligence, and lineage from visual to source expectations consistent with public sector accountability and with scientific writing guidance on explicit method and evaluation alignment.

2.6 OPERATIONALIZING KIMBALL WITH THE MEDALLION ARCHITECTURE IN MICROSOFT FABRIC

To operationalize the dimensional design, the dissertation adopts the Medallion Architecture (Bronze → Silver → Gold) as the data engineering pipeline in Microsoft Fabric (Databricks, 2025):

- Bronze (Ingestion & Fidelity Preservation): Raw ingestion of Na Minha Rua LX and daily weather with minimal transformation beyond parsing and schema conformance, preserving provenance for audit and re-processing.
- Silver (Cleaning, Standardization & Conformance): Standardization and conformance: typology normalization, de-duplication, geocoding harmonization, and creation of surrogate keys for time, location, and incident type. Where appropriate, SCD policies retain historical truth (e.g., statistical subsector changes). Weather is aligned with daily grain and enriched with derived features (e.g., daily precipitation, temperature range).
- Gold (Dimensional Models & Semantic Layer): Star schemas with conformed dimensions exposed to Power BI. The semantic model provides baseline measures (e.g., same-week averages), workload KPIs (Key Performance Indicators), and alert metrics (anomaly flags and thresholds), along with documentation for each measure and end-to-end lineage for auditability.

The pairing Kimball ↔ Medallion cleanly separates concerns: Kimball specifies analytical logic (facts, dimensions, conformance, grain), while Medallion governs the lifecycle (ingest → refine → publish) with reproducibility and governance (Kimball & Ross, 2013). This approach aligns with international guidance on a data-driven public sector, which stresses quality management, reusability, and traceability from data to decision. The chapter's structure and writing choices follow the program's instructions, explicitly outlining methods, measures, and scope (OECD, 2019).

3. METHODOLOGY

This chapter explains how the dissertation translates raw Na Minha Rua LX and weather data into a Business Intelligence artefact. It first sets out the research design and methodological stance, then derives business needs and questions, and finally summarises the dimensional architecture, ELT pipelines, and analytical measures.

Full technical artefacts (notebooks, dataflows, semantic model, and report files) are published in a dedicated Git repository that has the Fabric workspace structure (01_Lakehouse, 02_Bronze, 03_Silver, 04_Gold, 99_Master). The text below, therefore, focuses on the main design decisions rather than on line-by-line implementation.

3.1 RESEARCH DESIGN AND METHODOLOGICAL STANCE

The study follows a Design Science Research (DSR) strategy (Peppers et al., 2007): it builds and evaluates a BI artefact rather than estimating causal effects. The artefact is a set of Power BI dashboards built on a dimensional model in Microsoft Fabric that support routine monitoring and exploration of citizen-reported incidents in Lisbon. Evaluation in Chapters 4 and 5 is artefact-centric: it tests whether the dashboards answer the business questions defined below and whether indicators remain transparent and interpretable for municipal managers. Within this research strategy, Kimball's dimensional modelling is used as the design method to structure the engineering of the data model and the semantic layer (e.g., grain definition, conformed dimensions, and traceability).

Operational evaluation criteria. Following the view that artefact evaluation requires explicit criteria and metrics, the dissertation assesses the dashboard artefact through five operational criteria:

- **C1 – Coverage (BN/BQ coverage):** The extent to which the semantic model and dashboard pages provide measures and visuals that answer the business questions listed in Section 3.2. Coverage is evidenced by a mapping between BQs and the implemented visuals/measures used in the scenarios.
- **C2 – Interpretability (auditable logic):** The extent to which indicators can be understood and defended by municipal stakeholders, privileging seasonal baselines, deviations, and simple anomaly criteria over opaque models. Interpretability is supported by measure documentation and by the step-by-step reasoning demonstrated in the scenario walkthroughs (Matheus et al., 2018).
- **C3 – Traceability (lineage and reproducibility):** The extent to which results can be traced from a visual to a DAX measure and back to the underlying fact rows and transformations in Fabric, consistent with governance and public sector accountability principles (Kimball & Ross, 2013; OECD, 2019).

- **C4 – Semantic consistency (single source of truth):** The extent to which the same filter context and conformed dimensions yield comparable results across pages, enabling a coherent analytic path from overview to drill-down without contradictory numbers (Kimball & Ross, 2013)
- **C5 – Actionability (prioritisation support):** The extent to which the dashboards support ranking, hotspot identification, and territory shortlisting under capacity constraints, particularly through pressure indicators and weather-aware contextualisation in BN4–BN5.

These criteria are applied in Chapters 4 and 5 through scenario-based walkthroughs that demonstrate how a manager moves from questions to evidence and decision-relevant outputs.

Kimball’s dimensional data warehousing method provides the main methodological lens (Kimball & Ross, 2013). It is chosen because it foregrounds grain clarity, conformed dimensions, surrogate keys, and slowly changing dimensions (SCDs) as pre-conditions for additivity, drill-down consistency, and traceable indicators. These properties are critical when citizen-generated open data are reused for decision support and public accountability.

Operationally, the Kimball stance is implemented on Microsoft Fabric using the Medallion pattern (Bronze → Silver → Gold). Bronze tables preserve raw incidents and weather with minimal transformation; Silver tables clean, standardise, and conform them into shared dimensions and fact tables; Gold tables expose the final star schemas to a Direct Lake Power BI semantic model. All transformations, quality checks, and lineage are implemented upstream in Fabric, so that Power BI visuals only consume governed dimensions and documented DAX measures.

Data collection is purely secondary. Incident data come from Na Minha Rua LX open data exports supplied by Lisboa Aberta for 2020–2024. Weather data is retrieved from the Open-Meteo archive API for the same period and is aggregated by parish centroid. No personal identifiers are present in either dataset. The analytical focus is descriptive and diagnostic: the artefact characterises workloads, baselines, anomalies, territorial pressure, and weather sensitivity, rather than forecasting or impact evaluation.

3.2 BUSINESS NEEDS AND QUESTIONS

The seven business needs (BN1–BN7) are operationalised through the following business questions (BQs). These questions define what municipal stakeholders want to know from Na Minha Rua LX and weather data (Zippenfenig, 2023) and later drive the dimensional design, measures, and dashboards.

BN1 – Standardised space–time visibility of incident workloads

- **BQ1.1 – Temporal workload profile.** What is the evolution of incident workloads per day, week, and month over the historical period?

- **BQ1.2 – Spatial workload distribution.** How are incidents distributed across parishes, and which territories show systematically higher or lower workloads?
- **BQ1.3 – Thematic composition.** How are incidents distributed across thematic areas and types, at the city and parish levels, and how stable is this composition over time?

BN2 – Contextualisation of variation through simple seasonal baselines

- **BQ2.1 – Typical levels.** What is the “typical” number of incidents per weekday, week, and month, at the city and parish levels?
- **BQ2.2 – Seasonal patterns.** Are there clear seasonal patterns (e.g., winter peaks, holiday dips) in workloads by area and type?
- **BQ2.3 – Deviations from expected levels.** For a given day, week, or month, is the observed workload above, below, or close to its seasonal baseline?
- **BQ2.4 – Structural shifts.** Do baselines themselves shift over the years, indicating sustained workload growth or change after major interventions?

BN3 – Early identification of abnormal surges and spatial hotspots

- **BQ3.1 – Temporal anomalies.** On which days or weeks do incident counts significantly exceed or fall below baseline?
- **BQ3.2 – Spatial hotspots.** Which parishes or subsectors exhibit unusually high incident densities relative to their own history and to the rest of the city?
- **BQ3.3 – Persistence of anomalies.** Are anomalies isolated spikes or sustained periods of elevated reporting, and what is their typical duration?
- **BQ3.4 – Category-specific surges.** Are anomalies concentrated in particular incident areas/types or spread across multiple categories?

BN4 – Territory-focused prioritisation under capacity constraints

- **BQ4.1 – Territorial ranking.** Which parishes and parish–area combinations rank highest in workload and incident density over time?
- **BQ4.2 – Recurring incident patterns.** Where do similar incident types recur, suggesting structural or persistent problems?
- **BQ4.3 – Joint workload–anomaly priority.** Which territories combine a high workload with frequent anomalies and therefore merit priority attention?
- **BQ4.4 – Perceived pressure and equity.** Do some parishes appear systematically more exposed to citizen-reported problems, raising equity considerations?

BN5 – Pragmatic integration of weather as an exogenous factor

- **BQ5.1 – Co-variation with weather.** How do incident workloads, especially for weather-sensitive categories, co-vary with daily precipitation and temperature extremes?
- **BQ5.2 – Weather and anomalies.** When anomalies occur, do they coincide with heavy rain or unusual temperatures, or also under normal weather?
- **BQ5.3 – Incremental value of weather.** Does adding weather information materially improve the explanation or prediction of workloads compared with seasonal baselines alone?
- **BQ5.4 – Differential sensitivity.** Are some parishes, infrastructures, or incident types more sensitive to weather, indicating local vulnerabilities?

BN6 – Transparent and traceable indicators from dashboards back to raw data

- **BQ6.1 – Record-level traceability.** For any indicator, how can staff trace it back to the exact subset of Na Minha Rua LX and weather records used?
- **BQ6.2 – Impact of data quality.** How do duplicates, inconsistent typologies, or missing locations affect indicator reliability?
- **BQ6.3 – Minimal documentation.** What is the minimal set of transformation and aggregation steps that must be documented for auditors and citizens?
- **BQ6.4 – Limits of current data.** Which types of indicators (e.g., SLA compliance) cannot legitimately be derived without workflow logs, and how should this be communicated?

BN7 – Self-service exploration by different municipal roles

- **BQ7.1 – Strategic questions.** Which high-level questions about trends, territorial disparities, and priorities should strategic managers be able to answer self-service?
- **BQ7.2 – Operational questions.** Which short-term questions about recent surges by parish or area should operational coordinators address?
- **BQ7.3 – Analytical questions.** Which methodological or exploratory questions (e.g., alternative anomaly thresholds, new external drivers) are relevant for analysts and data teams?
- **BQ7.4 – Shared “single source of truth”.** Which definitions, filters, and groupings must be common across roles so that independent explorations still rely on the same underlying truth?

These BQs close the gap between high-level needs and concrete analytical requirements and are later used as a checklist when designing and evaluating the dashboards.

3.3 DATA DESCRIPTION

This dissertation is grounded on two raw, purely secondary datasets: (i) citizen-reported urban incident data from Lisbon's Na Minha Rua LX (NMRLX) open-data exports and (ii) daily weather observations retrieved through the Open-Meteo archive API for the same study window. The empirical intent is decision-support oriented: the artefact translates raw incident logs into interpretable monitoring signals (workloads, variation, surges, territorial differences, and weather context).

Incident data are obtained from Na Minha Rua LX open data exports supplied by Lisboa Aberta, covering 2020–2024, and include approximately 943,000 incident records at a daily grain. NMRLX functions as a municipal citizen-reporting channel through which residents submit reports on public-space issues (e.g., potholes, street lighting, waste, and related maintenance problems). Therefore, at the raw stage, the incident dataset should be interpreted as a citizen-generated reporting signal: it captures what residents choose to report, when and where, rather than the full universe of urban infrastructure faults.

At the ingestion stage, the project treats the NMRLX exports as authoritative raw extracts and preserves their original encodings. The exports provide the minimal informational content required for later dashboard objectives: a record reference (enabling traceability), a registration date/time attribute (supporting temporal analysis), a typology/classification attribute (supporting thematic analysis), and a spatial attribute that supports territorial grouping at the municipal scale. Any additional descriptive fields (when present) are treated as optional and are only leveraged later if their completeness and stability are sufficient for reliable interpretation.

To complement the incident logs with an exogenous contextual layer, the project also collects daily weather data from the Open-Meteo archive API for 2020–2024. Weather is included to support explanatory inspection of co-variation between incident pressure and precipitation, especially in business questions explicitly focused on “weather context”. Operationally, the raw extraction is performed per parish centroid, aligning the weather observations with the same territorial scope used for incident monitoring.

Two scope boundaries apply and must be made explicit at the raw-data stage because they constrain what the artefact can legitimately claim. First, only incident and weather data are used; workflow and resolution events are out of scope, meaning the raw data cannot support service performance indicators such as time-to-resolution or SLA compliance. Second, because NMRLX is citizen-reported, the observed spatial and temporal patterns may reflect reporting behaviour (awareness, accessibility, propensity to report) as well as underlying urban conditions; as such, the artefact is

positioned as a monitoring and prioritisation instrument based on reported demand, not as a causal model of urban deterioration.

From an ethics and privacy perspective, the data strategy remains low-risk and consistent with the programme's ethics expectations: the dissertation uses only open, non-personal data from Na Minha Rua LX and Open-Meteo, makes no attempt to identify individuals, and reports analyses at aggregated territorial levels (parish or city).

3.4 DIMENSIONAL DESIGN WITH KIMBALL AND MEDALLION

The analytical backbone follows Kimball's nine-step method, adapted to the Na Minha Rua LX context.

First, the business process is defined as urban incident reporting and monitoring in Lisbon, with daily weather as an exogenous context.

Second, the grain is declared as *Day × Parish × IncidentType* for incidents and *Day × Parish* for weather. This daily grain matches the temporal resolution of the weather source and the intended decision cadence (daily and weekly monitoring), while still supporting monthly and yearly aggregations.

Third, three conformed dimensions are implemented: DimDate, DimLocation, and DimIncidentType. DimDate provides a reusable calendar hierarchy (Year → Semester → Trimester → Month → Day) with flags that support time-intelligence logic. DimLocation represents parishes with centroid coordinates for mapping. DimIncidentType encodes the Na Minha Rua LX taxonomy into Areas and Types. The detailed attribute specifications are provided in Appendix A (Tables A-1 to A-3).

Fourth, two additive fact tables are defined: FactIncidentDaily and FactWeatherDaily. FactIncidentDaily stores one row per incident occurrence at the daily grain, linked to all three dimensions and providing the basis for the Incidents measure. FactWeatherDaily stores one row per parish–day with precipitation, temperature, wind speed, and a Boolean Is_Rainy_Day flag, aligned to the same Date and Location keys. Appendix A (Tables A-4 to A-7) contains the description of their schemas.

Fifth, a bus matrix in Appendix A (Table A-8) documents which dimensions are shared across the implemented processes (Incident Events and Weather) and a conceptual future process (Service Workflows). This matrix acts as the contract for future extensions: new facts must reuse the same DimDate and DimLocation, and, where applicable, DimIncidentType.

Sixth, all dimensions and facts use surrogate keys. DimDate's sk_date is a yyyyymmdd integer and does not require SCD treatment. DimLocation and DimIncidentType are prepared for Type-2 SCD if parish boundaries or taxonomies change; new versions would add rows rather than overwriting history. Late-arriving facts are handled via "unknown" keys and monitored quality thresholds in the silver and gold layers.

Seventh to ninth, the ELT system, semantic layer, and quality governance are implemented via the Medallion pipelines. Bronze pipelines ingest incident Excel exports and weather API responses into `tbl_bronze_incidents_raw` and `tbl_bronze_weather_raw`. Silver dataflows standardise and conform these into `dim_date_silver`, `dim_location_silver`, `dim_incident_type_silver`, `fact_incident_daily_silver`, and `fact_weather_daily_silver`. Gold copy activities publish the final dimensional tables (`dim_date`, `dim_location`, `dim_incident_type`, `fact_incident_daily`, `fact_weather_daily`) and trigger a refresh of the Direct Lake semantic model. A master pipeline orchestrates the sequence Bronze → Silver → Gold to avoid partial updates. Data quality checks (schema conformance, duplicate removal, and unknown key rates) and referential integrity tests are logged at Silver and Gold levels.

All notebooks, dataflows, and pipeline definitions are stored in the Git repository alongside this dissertation, providing full technical documentation of Kimball's nine steps in practice.

3.5 MEASURE DESIGN, ANALYSIS, VALIDITY, AND ETHICS

Analytical logic is encapsulated in DAX measures defined over the gold star schema. Measures are grouped into display folders that mirror the business needs:

- **Core workload** (Incidents, Avg Incidents) provides standardised counts and descriptive statistics that underpin all pages (BN1).
- **Time-intelligence (Incidents YTD, Incidents (LY), Incidents YoY %)** provides year-to-date and year-on-year comparisons using the DimDate hierarchy (BN1–BN2).
- **Baselines** (Incidents Baseline Daily, Incidents Baseline Monthly, and their standard deviations, plus adaptive versions and Incidents vs Baseline %) provide simple daily and monthly seasonal references for contextualising variation (BN2).
- **Anomalies** (daily and monthly z-scores and flags, Anomaly Days Count, Anomaly Days %, Avg Anomaly Z-Score Daily) support the detection of temporal surges and spatial hotspots (BN3).
- **Territorial pressure** (Avg Incidents per Parish, Pressure Index) quantifies relative workload compared to a city-wide per-parish average (BN4).
- **Weather sensitivity** (rainy/dry splits, incident rates per parish day, Rainy Lift, Rainy Share %, Rainy Sensitivity Category, and Precipitation) quantifies how workloads relate to rainy conditions (BN5).

Complex measures are systematically composed from validated intermediate measures rather than raw columns, and naming conventions follow descriptive, traceable standards. The full technical artefacts, including DAX definitions, semantic

model metadata, dimensional schemas, and data-pipeline configurations, and the functional Power BI project file, are documented and version-controlled in the project's GitHub repository (João Ramos, 2026), supporting transparency, reproducibility, and reuse in line with BN6 and BN7.

Validity is addressed by aligning all measures with the declared grain and conformed dimensions, reducing the risk of double-counting and incomparable filters. Reliability is supported by automated pipeline runs, schema checks, and referential integrity tests; failures prevent publication to Gold. The evaluation in Chapters 4 and 5 further examines usefulness and interpretability through scenario-based walkthroughs.

Ethically, the study uses only open, non-personal data from Na Minha Rua LX and Open-Meteo. No attempt is made to identify individuals, and analyses are reported at the aggregated parish or city level. The project follows the programme's ethical approval checklist, with additional discussion of data limitations and biases provided in Chapter 6.

4. EMPIRICAL STUDY

4.1 OVERVIEW OF THE DASHBOARD ARTEFACT AND SHARED LAYOUT

The empirical component of the dissertation materialises the Kimball Medallion architecture described in Chapter 3 into a set of interactive dashboards built on a single Power BI semantic model. The dashboards are conceived as a family of views over FactIncidentDaily and its conformed dimensions (Date, Location, IncidentType), rather than as independent reports, so that all pages reuse the same measure definitions, hierarchies, and governance mechanisms. This design choice follows guidance on public sector dashboards (OECD, 2019), which emphasises a consistent semantic layer and explicit measurement logic as prerequisites for transparency, interpretability, and reuse.

Visually, the artefact adopts a stable layout across all pages:

1. a 16:9 canvas optimised for both desktop viewing and projected presentations
2. a fixed left-hand panel containing institutional branding, global KPIs, and slicers
3. a main analytical area on the right, where page-specific visuals are displayed
4. a navigation bar in Power BI that allows users to move between the homepage and thematic pages (BN1–BN5, weather analysis, etc.).

This layout was intentionally kept simple and repeatable, in line with dashboard design literature that recommends predictable visual structures, clear affordances, and minimal cognitive overhead for frequent users.

4.1.1 Shared filter panel and semantic consistency

The left-hand panel is common to all pages and now contains four slicers:

1. **Date Hierarchy** – a hierarchical slicer built from DimDate’s *Year Hierarchy* (Year → Semester → Trimester → Month → Day). In the report, it is labelled “Date Hierarchy” so that end-users perceive it as a generic calendar control rather than a technical object.
2. **Month** – a complementary slicer on the Month column from DimDate, used to focus the analysis on a specific month while keeping the Date Hierarchy available for higher-level navigation (e.g., selecting a semester or trimester).
3. **Parish Name** – slicer on DimLocation, allowing restriction of all pages to a specific parish or to the full municipality.
4. **Area / Type** – slicer on DimIncidentType’s Area/Type hierarchy, enabling the dashboards to be restricted to particular thematic categories such as road surface, lighting, or sanitation.

These slicers apply a single filter context across all visuals and pages because they operate directly on conformed dimensions in the semantic model. This design enforces the “single source of truth” principle: when a manager filters to, for example, *2024 – 2nd Semester – Arroios – Iluminação Pública*, all pages answer business questions under precisely the same subset of Na Minha Rua LX incidents. In doing so, the artefact addresses recurring concerns in the literature about non-comparable and ad-hoc indicators in open-data dashboards.

KPI strip and time-intelligence measures

Immediately below the municipal logo, a strip of three KPI cards summarises key workload indicators in the current filter context, followed by a fourth card expressing year-on-year change:

- **Incidents** – total number of incidents, defined as `[Incidents] = COUNTROWS(fact_incident_daily)`.
- **Incidents (LY)** – incident count for the corresponding period in the previous year, implemented via `SAMEPERIODLASTYEAR` over `DimDate[Registration Date]` so that it respects whatever date window is selected in the slicers.
- **Incidents YTD** – year to date workload using the standard `DATESYTD` pattern on `DimDate`, aligned with the organisation’s definition of the operational year.
- **Incidents YoY %** – year on year percentage change, computed as the relative difference between *Incidents YTD* and the same measure shifted one year back.

These measures are defined once in the semantic layer and reused across every page, which is consistent with the methodological stance of performing all transformations upstream of the visuals and exposing documented, auditable measures to Power BI. From an evaluation perspective, the KPI strip provides an immediate, standardised view of workload levels and trends that can be interpreted consistently across the empirical study and the subsequent results and discussion chapters.

4.2 HOMEPAGE - OVERVIEW AND NAVIGATION

The homepage serves as the entry point to the artefact and provides orientation before the user moves into the specialised analytical pages linked to BN1–BN5. Rather than presenting additional charts or maps, this page explains *what* each thematic page does and offers a consistent place to apply global filters. Conceptually, it is not tied to a single business need; instead, it situates the dashboards within the broader context of the research question and the methodological architecture.

The homepage reuses the shared left-hand panel (logo, KPIs, and slicers) described in Section 4.1.1, so that users immediately see the overall workload indicators for the

current filter context (e.g., city-wide, or restricted to a parish or incident area). The main area on the right is reserved for navigation tiles that mirror the business needs BN1–BN5.

4.2.1 Objectives of the homepage

The homepage has four main objectives:

1. **Contextualisation** - Provide an immediate sense of scale through the KPI strip (Incidents, Incidents YTD, Incidents YoY %), under the same filter logic that will be used on all other pages. This provides a shared understanding of the overall workload for subsequent analyses.
2. **Pre-filtering entry point** - Allow users to set Year, Month, Parish, and Area/Type filters before navigating to any analytical page. Because the slicers act on conformed dimensions in the semantic model, the chosen context is preserved when moving from the homepage to BN1–BN5, enforcing a single source of truth.
3. **Explanation of analytical scopes** - Summarise, in plain language, what each thematic page does and which business needs and questions it addresses. This makes the empirical study's structure explicit, even for users unfamiliar with BN1–BN5 as research constructs.
4. **Navigation hub** - Offer clear, clickable tiles that take the user directly to the corresponding dashboard pages: BN1 – Workload & Spatial Distribution, BN2 – Seasonality & Baselines, BN3 – Anomalies & Hotspots, BN4 – Territory Prioritisation: Categories & Hotspots, BN5 – Weather Context.

Together, these objectives ensure that the empirical study does not start from a highly specialised perspective, but rather from a neutral, explanatory landing page that clarifies scope, data coverage, and available analytical paths.

4.2.2 Visual composition

The homepage keeps the standard 16:9 canvas and fixed left-hand panel. The right-hand side of the page is organised as a vertical stack of five navigation tiles, one per business need:

Structure of each tile.

1. A **blue card** on the left with the business need label (e.g., “BN1 – Workload & Spatial Distribution”).
2. A **grey description panel** on the right containing 2–3 lines of explanatory text written from the user’s perspective (“What you can do here”), explicitly referencing the corresponding BN and BQ identifiers (e.g., BN1/BQ1.*).

3. A **navigation icon** on the far right is configured with Power BI's *Page navigation* action, which takes the user to the respective report page.
4. **Ordering and alignment.**

Tiles are ordered from top to bottom as BN1 through BN5, mirroring the logical progression of the empirical study: from standardised workloads (BN1) to seasonality (BN2), anomalies (BN3), territory prioritisation (BN4), and weather context (BN5). The layout preserves consistent margins and alignment with the shared sidebar, reinforcing the perception of a single, coherent artefact.
5. **Interaction with filters.**

While the tiles themselves do not display numeric values, the sidebar KPIs respond immediately to any filter changes made on the homepage. Users can therefore, for example, restrict the view to a single parish and incident area, observe the updated KPIs, and then navigate to the BN2 or BN3 pages to analyse seasonality or anomalies under that same context.

By separating orientation (homepage) from detailed analysis (BN1–BN5 pages) and by keeping the semantic model and slicers consistent across all views, the dashboard design aligns with recommendations from the literature on public sector dashboards: make scope and structure explicit, minimise cognitive load at entry, and maintain a clear line of sight from high-level navigation to concrete analytical questions. A full-page screenshot of the homepage is provided in Appendix A (Figure A-1).

4.3 PAGE “WORKLOAD & SPATIAL DISTRIBUTION” (BN1)

4.3.1 Objective and link to BN1

The first thematic page addresses BN1 – Standardised space–time visibility of citizen-reported incident workloads. It provides a consistent view of how many incidents are reported over time and across the territory, and how these workloads decompose by thematic area.

In terms of BQ1.1–BQ1.3, the page is designed to answer questions such as:

- How many incidents are reported per month and per year in *Na Minha Rua LX*?
- How are workloads distributed across parishes?
- For a given time period or parish, what is the composition of incidents by thematic area (e.g., road surface, lighting, waste)?

The page, therefore, serves as the main workload cockpit for *Na Minha Rua LX*. All visuals are driven by the conformed Date, Location, and Incident Type dimensions and by the shared slicers (Year/Semester/Trimester/Month, Month, Parish, Area/Type),

ensuring that managers always read the same definition of “incident” and “period” across the dashboard.

4.3.2 Visual composition and interactions

The shared left-hand panel (logo, KPI strip, and slicers) described in Section 4.1 remains visible and controls the entire page. The main analytical area contains two core visuals, both enriched with report-page tooltips.

Visual 1 – Incidents by Month and Year (temporal workload profile)

- **Type:** Clustered column chart.
- **X-axis:** Date Hierarchy
- **Legend:** Year
- **Y-axis:** [Incidents] (count of incident fact rows).

For the selected filter context, this chart shows monthly workloads side-by-side by year. It directly answers **BQ1.1** by allowing users to compare:

- workloads between months within the same year;
- the same month across different years (seasonal growth or decline).

A **report-page tooltip** is attached to this visual:

- **Tooltip visual:** horizontal bar chart “*Incidents by Area and Year*”.
- **Y-axis:** high-level incident *Area* (aggregated categories from *DimIncidentType*).
- **X-axis:** [Incidents].
- **Filter context:** inherits the hovered Month and Year from the main chart.

When the user hovers a column (e.g., *March 2024*), the tooltip opens and immediately shows how that month’s workload is split across thematic areas. This keeps the main chart simple while still exposing the composition needed for **BQ1.3**.

Visual 2 – Incidents by Parish and Area (spatial distribution)

- **Type:** Bubble map.
- **Location:** parish centroids from *DimLocation* (Latitude / Longitude).
- **Bubble size:** [Incidents] in the current filter context.
- **Legend:** Parish Name, colouring bubbles by parish.

This visual operationalises **BQ1.2** and **BQ1.3** by showing:

- which parishes sustain the largest absolute workloads (bubble size);
- how workloads are distributed geographically across the municipality.

The map is also enhanced with a **report-page tooltip**:

- **Tooltip visual:** horizontal bar chart “*Incidents by Area and Parish Name*”.
- **Y-axis:** incident Area.
- **X-axis:** [Incidents].
- **Filter context:** the hovered parish (and any slicer filters on Year/Month/Area).

Hovering a bubble (e.g., *Alvalade*) reveals how that parish’s incidents are distributed by area (road surface, lighting, waste, etc.). Clicking a parish cross-filters the monthly chart and KPI cards, allowing focused analysis of that territory. A full-page screenshot of the BN1 dashboard page is provided in Appendix A (Figure A-2).

4.3.3 Analytical use

Typical analytical tasks supported by this page include:

- **City-wide monitoring.** With *Parish = All* and a given Year selected, managers can quickly see which months are busier, and which thematic areas dominate each month via the tooltip. This fulfils BN1’s requirement for standardised time series indicators instead of ad-hoc spreadsheets.
- **Territory-focused exploration.** Selecting a parish on the map filters the entire page to that territory. The monthly chart then shows the temporal profile for that parish alone, helping answer questions such as “How did incident volumes evolve in *Parish X*?” or “Is *Parish X* consistently high in a given thematic area?”.
- **Thematic composition analysis.** Filtering by *Area* in the slicer (e.g., *Higiene Urbana*) and/or using the tooltips allows managers to examine whether some parishes systematically generate more reports in a specific area. This provides an initial evidence base for later prioritisation and equity discussions (linking BN1 with BN4).

Overall, the BN1 page combines a temporal profile with a spatial snapshot and contextual tooltips to deliver a coherent, drillable view of workloads by time, territory, and thematic area.

4.4 PAGE “SEASONALITY & BASELINES” (BN2)

4.4.1 Objective and link to BN2

The second thematic page is dedicated to **BN2 – Contextualisation of variation through seasonal baselines**, which emphasises that incident counts should not be interpreted in isolation but always against historical reference levels. This need aligns with the forecasting literature, which recommends simple, parsimonious baselines such as moving averages and same-day comparisons as robust points of reference for noisy municipal workloads.

Concretely, the page is designed to answer BQ2.1–BQ2.3:

- What is the typical number of incidents per day and per week for the selected context?
- Are there seasonal patterns at the monthly level?
- For recent days, is the workload above, below, or close to its expected level?

Whereas Page 1 focuses on absolute volumes, Page 2 shifts to relative interpretation: assessing whether “today”, “this week”, or “this month” is unusually busy given recent history.

4.4.2 Visual composition and interactions

Beyond the shared filter panel and KPIs, the page contains three main visuals.

Incidents vs Baseline (Daily / Monthly)

- Type: combo chart (columns + line).
- X-axis: Registration Date from DimDate, with drill-down from Month to Day.
- Legend: Year, allowing side-by-side comparison of multiple years.
- Columns: [Incidents].
- Line: [Incidents Baseline (Adaptive)].

At the monthly level, the chart shows, for each month, the total incidents by year (e.g., 2023 vs 2024) overlaid on a single 12-month rolling baseline line. This makes it clear which years sit persistently above or below the historical reference. When the user drills down to the daily level, the line switches to the 28-day daily baseline, revealing the typical weekday “saw-tooth” pattern and allowing day-by-day comparison against expected workload. This directly supports BQ2.2 and BQ2.4.

Average Daily Incidents by Day of Week

- Type: bar chart.

- X-axis: Week Day from DimDate.
- Y-axis: average of [Incidents] for the period defined by the slicers.
- Additional design: gradient conditional formatting on bar colour to reflect relative magnitude.

This visual operationalises BQ2.1 by quantifying weekday effects (e.g., Mondays and Tuesdays consistently busier than weekends), providing empirical evidence for resource planning decisions.

Recent Deviations from Seasonal Baseline

- Type: matrix.
- Rows: Registration Date (with optional expansion to incident Area).
- Columns: [Incidents], [Incidents Baseline Daily], and [Incidents vs Baseline % (Daily)].
- Sorting: by deviation in descending order to surface the most exceptional days first.
- Conditional formatting: large positive deviations highlighted in red, moderate ones in amber, and values close to baseline in neutral tones.

This table answers BQ2.3 by listing precisely which recent days were substantially above or below expected levels and by how much in percentage terms. It also contributes to BN6 by making indicator derivation transparent: for any flagged day, users can see the raw count, the baseline, and the deviation side by side. A full-page screenshot of the BN2 dashboard page is provided in Appendix A (Figure A-3).

4.4.3 Analytical use

In practice, the **Seasonality & Baselines** page supports:

- **Daily operational monitoring** – coordinators quickly check whether yesterday or last week were unusually busy before interpreting complaints about “exceptional” workloads.
- **Planning of inspection schedules** – weekday averages and seasonal patterns justify staffing levels and on-call capacity, especially around recurrent high-pressure days.
- **Strategic review** – managers examine whether baselines themselves drift upward or downward over the years, which would indicate structural changes in citizen reporting behaviour or underlying infrastructure conditions.

By combining simple, well-documented baselines with clear visuals, multi-year comparison and drill-down capabilities, this page realizes BN2. It implements

recommendations from the time-series monitoring and dashboard literature to favour interpretable, auditable metrics over opaque modelling in public-sector decision support.

4.5 PAGE “ANOMALIES & HOTSPOTS” (BN3)

4.5.1 Objective and link to BN3

The third thematic page implements **BN3 – Early identification of abnormal surges and spatial hotspots**. Building on the seasonal baselines of BN2, the focus now is on episodes in which workloads depart significantly from what is considered normal, both over time and across space.

In terms of BQ3.1–BQ3.4, the page is designed to support questions such as:

- On which specific days do incident counts strongly exceed their seasonal baseline?
- Which parishes are consistently associated with high anomaly levels?
- For a given anomalous day, which thematic areas (e.g., road surface, lighting, waste) drive the spike?
- Are anomalies isolated, or do they appear in clusters, indicating sustained abnormal periods?

The page thus acts as an “exception radar” for Na Minha Rua LX, helping coordinators distinguish routine variation from genuinely exceptional situations.

4.5.2 Visual composition and interactions

In addition to the shared filter panel and KPI strip, the page contains three coordinated visuals.

Temporal anomaly profile

- Type: combo chart (columns + line).
- X-axis: Registration Date at **daily** grain (no grouping by day-of-month).
- Columns: [Incidents].
- Line: [Incidents Anomaly Z-Score Daily].

The chart shows the evolution of incidents day by day, with the anomaly Z-score line highlighting spikes ($Z \gg 0$) and unusually quiet days ($Z \ll 0$). Users can scroll horizontally through the period using the built-in zoom bar.

A report page tooltip is attached to this visual: when hovering over a bar, a separate bar chart “*Category-specific anomalies*” appears, plotting [Incidents Anomaly Z-Score Daily] by Area for that specific day. This immediately reveals which thematic categories are driving a given spike (e.g., a road surface surge versus a waste-related surge), directly informing BQ3.4.

Spatial hotspots

- Type: bubble map.
- Location: parish centroids from DimLocation.
- Bubble size: [Incidents] for the period defined by the slicers.
- Bubble colour: conditional formatting based on aggregated [Incidents Anomaly Z-Score Daily] for each parish (e.g., red for high anomalies, orange for moderate, green for near-baseline).

This visual operationalises BQ3.2 by revealing parishes that, in the selected time window, experience systematically higher anomaly levels. Selecting a parish filters the temporal chart and the anomaly table to that parish, allowing users to zoom into localised episodes.

Top Anomaly Days (Detail)

- Type: matrix.
- Rows: Registration Date, with drill-down to incident Area when expanded.
- Columns: [Incidents], [Incidents Baseline Daily], [Incidents Anomaly % (Daily)], and [Incidents Anomaly Z-Score Daily].
- Sorting: descending by Z-score to surface the strongest anomalies at the top.
- Formatting: heatmap-style conditional formatting on the Z-score (red for very high values, amber for moderate, neutral tones near baseline).

This table answers BQ3.1 and BQ3.3 by listing the most anomalous days, quantifying how far above baseline they sit, and showing whether anomalies persist over adjacent days. When a row is selected, the map and temporal chart are cross-filtered to that date, closing the loop between tabular detail, spatial pattern, and temporal context. A full-page screenshot of the BN3 dashboard page is provided in Appendix A (Figure A-4).

4.5.3 Analytical use

Typical analytical workflows supported by the **Anomalies & Hotspots** page include:

- **Rapid anomaly screening** – managers scan the temporal chart and matrix for days with exceptionally high Z-scores to confirm whether perceived “crisis” periods are statistically exceptional.
- **Root-cause exploration by category** – once a spike is identified, the tooltip chart shows which incident Areas dominate the anomaly, guiding conversations with the relevant operational teams (e.g., road maintenance versus sanitation).
- **Spatial prioritisation** – the hotspot map highlights parishes with recurrent high anomalies, providing an evidence base for targeted inspections, local campaigns, or infrastructure review.
- **Episode characterisation** – combining the three visuals helps distinguish isolated spikes (e.g., a single storm-related day) from longer anomalous periods, which may suggest structural service problems.

By making anomalies explicit, explainable, and navigable across time, space, and category, this page operationalises BN3 and prepares the ground for BN4’s territory-focused prioritisation of municipal attention.

4.6 PAGE “TERRITORY PRIORITIZATION CATEGORIES & HOTSPOTS” (BN4)

4.6.1 Objective and link to BN4

This page operationalises BN4 – Support for territory-focused prioritisation under capacity constraints. While BN1 and BN2 show “how much” and “how unusual” workloads are, BN4 asks where scarce inspection and planning effort should go first, given only citizen-reported incidents. Concretely it targets BQ4.1–BQ4.4:

- **BQ4.1 – Territorial ranking by workload and concentration**
- **BQ4.2 – Recurring incident types and locations**
- **BQ4.3 – Combined workload–anomaly prioritisation**
- **BQ4.4 – Perceived pressure versus geographic equity**

The page, therefore, serves as a *prioritisation cockpit*: it ranks parishes, highlights the most recurrent thematic areas, and combines workload with anomaly metrics to suggest where attention is most urgent.

As on all pages, the shared left filter panel (Year / Month / Parish / Area) defines a single context for every visual.

4.6.2 Visual composition and interactions

Visual 1 – Parish workload ranking (top-left)

- Type: horizontal bar chart.
- Axis: *Parish Name* (DimLocation).
- Value: [Incidents].
- Purpose: Implements **BQ4.1** by ranking parishes by total incident volume in the current filter context (e.g., full year, or a specific area such as “Higiene Urbana”). Scrolling allows inspection of all parishes. Selecting one bar cross-filters all other visuals to that parish.

Visual 2 – Workload vs anomaly (prioritisation quadrant) (top-right)

- Type: scatter plot.
- X-axis: [Incidents].
- Y-axis: [Avg Incidents Anomaly Z-Score Daily] (average anomaly intensity).
- Legend: *Area* (thematic category).
- Purpose: Addresses **BQ4.3** by combining absolute workload and anomaly behaviour.
 - Points in the upper-right represent parishes that are both busy and frequently above baseline (high-priority).
 - Points with high Z-scores but moderate workload suggest volatile territories, while high workload with low Z-scores indicate *consistently busy but predictable* areas. This gives managers an interpretable “quadrant” view for balancing routine load against exceptional behaviour.

Visual 3 – Recurring incident types by parish (bottom-left)

- Type: horizontal bar chart.
- Axis: *Area* (incident thematic class).
- Value: [Incidents].
- Purpose: Implements **BQ4.2** by highlighting, for the selected parish (or for the whole city if no parish is selected), which incident types recur most often. When a parish is selected in Visual 1, this chart becomes a detailed profile of that territory’s recurring problems, supporting more targeted planning (e.g., repeated “Higiene Urbana” vs “Saneamento”).

Visual 4 – Perceived pressure index by parish (bottom-right)

- Type: horizontal bar chart with reference line at Pressure Index = 1.
- Axis: *Parish Name*.

- Value: [Pressure Index]; bars coloured by conditional formatting (e.g., >1.5 red, 1.1–1.5 orange, <1 green).
- Purpose: Directly addresses **BQ4.4** by quantifying whether some parishes appear systematically more exposed to citizen-reported problems than would be expected under equal geographic treatment. Values above 1 indicate above-average pressure; below 1 indicate comparatively lighter loads. A full-page screenshot of the BN4 dashboard page is provided in Appendix A (Figure A-5).

4.6.3 Analytical use

In practice, this page supports several prioritisation workflows:

- **Ranking and shortlisting territories** – managers first use the workload ranking and pressure index to identify the top parishes by volume and equity-adjusted pressure.
- **Understanding what drives pressure** – for a shortlisted parish, the recurring types chart reveals which categories dominate the workload, guiding specialised interventions (e.g., drainage vs street-surface maintenance).
- **Balancing routine vs exceptional problems** – the workload–anomaly scatter helps distinguish parishes that are consistently busy from those with frequent exceptional peaks, supporting different management strategies (continuous resourcing vs contingency planning).

By combining rankings, anomaly-aware quadrants, and equity indicators, the BN4 page concretely realizes the business need of territory-focused prioritisation under capacity constraints, while remaining fully traceable to the same open incident data and baseline logic established in BN1–BN3.

4.7 PAGE “WEATHER CONTEXT” (BN5)

4.7.1 Objective and link to BN5

This page addresses **BN5 – Pragmatic integration of weather as an exogenous explanatory factor**. Rather than treating weather as a complex predictive model, the goal is to explain incident workloads using simple, auditable weather flags that primarily distinguish rainy days from dry days at the parish–day level.

The page is designed to support **BQ5.1–BQ5.4**:

- **BQ5.1** – How do daily incident counts co-vary with precipitation extremes?
- **BQ5.2** – When surges occur, do they coincide with days of heavy rain or also appear under normal weather?

- **BQ5.3** – Does including weather add explanatory power compared to purely seasonal baselines?
- **BQ5.4** – Are some incident categories or parishes more sensitive to weather than others?

Internally, the semantic model links incidents and weather through a (sk_date, sk_location) bridge: rainy parish–days are identified in fact_weather_daily, and measures such as *Incidents (Rainy)*, *Incidents (Dry)*, and *Rainy Lift* are computed by re-filtering fact_incident_daily on those parish–date pairs. This keeps weather logic transparent and fully traceable to the raw open data.

4.7.2 Visual composition and interactions

Beyond the shared filter panel and KPI strip, the Weather Context page contains four visuals.

Weather-aware temporal profile

- **Type:** combo chart (columns + line)
- **X-axis:** dim_date[Registration Date]
- **Columns:** [Incidents]
- **Line (secondary axis):** daily precipitation from fact_weather_daily

This visual directly addresses **BQ5.1** by showing, day by day, how incident volume and rainfall move together in the selected context (year, parish, area). Conditional formatting based on *Rainy Lift* can be applied to the columns so that days whose workloads are unusually concentrated on rainy parish–days stand out immediately.

Weather-sensitive categories

- **Type:** horizontal bar chart
- **Axis:** incident Area
- **Measure:** [Rainy Lift] = (Incidents on rainy parish–days) / (Incidents on dry parish–days)
- **Colour:** conditional on the derived [Rainy Sensitivity Category]
 - *High* (\approx lift ≥ 1.20) – substantially more incidents on rainy days;
 - *Moderate / Neutral* (around 1.0);
 - *Lower on rainy days* (lift < 1).

This chart operationalises **BQ5.3 and BQ5.4** at the **category level**, ranking which thematic areas show the strongest proportional increase on rainy days (e.g., *Estradas e Ciclovias* and *Saneamento*) versus those that do not.

Rainy lift by parish (ranking)

- **Type:** horizontal bar chart
- **Axis:** dim_location[Parish Name]
- **Measure:** [Rainy Lift]
- **Colour:** same rules as the category chart (High / Moderate / Neutral / Lower).
- **Tooltip:** displays [Incidents (Rainy)], [Incidents (Dry)], *Rainy Share %*, and *Rainy Sensitivity Category*.

This visual tackles **BQ5.4** in territorial terms. It provides a clear ranking of parishes by weather sensitivity without any Top-N cut, which is feasible given Lisbon’s limited number of parishes. Selecting a parish here cross-filters the temporal profile and the map, making it easy to move from “who is most sensitive” to “when did that parish experience rain-related surges”.

Weather sensitivity by parish (map)

- **Type:** bubble map
- **Location:** parish centroids from DimLocation
- **Size:** [Incidents] in the current filter context
- **Colour:** [Rainy Sensitivity Category] or a gradient based on [Rainy Lift]

The map complements the ranking by showing the spatial distribution of weather sensitivity. Users can see whether high lift parishes are clustered (e.g., around specific drainage basins or coastal areas) or dispersed across the city. Clicking on a parish highlights it in the ranking chart and filters the temporal profile to its parish–day history.

Together, the ranking and map pair reproduce the pattern used in BN4 (workload and perceived pressure), but now focused on rain-related vulnerability rather than general workload. A full-page screenshot of the BN5 dashboard is provided in Appendix A (Figure A-6).

4.7.3 Analytical use

Typical analytical questions supported by the Weather Context page include:

- **Event-level interpretation (BQ5.1–BQ5.2):**
When a spike is observed in BN2/BN3, coordinators can check BN5 to see whether it coincides with intense rainfall in the temporal profile. The

combination of columns and the precipitation line clarifies whether the surge is plausibly weather-driven.

- **Category-level sensitivity (BQ5.3):**

By reading the *Weather-sensitive categories* chart, managers can see which thematic areas react most strongly to rain, supporting arguments such as “on rainy days, road-surface and sanitation incidents increase disproportionately compared to other categories”.

- **Territorial vulnerability (BQ5.4):**

The *Rainy lift by parish* ranking and the *Weather sensitivity by parish* map allow planners to identify parishes where infrastructure appears particularly vulnerable to rain, and to contrast them with parishes where workloads are actually lower on rainy days. This informs targeted maintenance, drainage investments, or communication campaigns.

- **Cross-page interpretation:**

BN5 is intended to be read in conjunction with BN2 (Seasonality & Baselines) and BN3 (Anomalies & Hotspots). When an anomalous day or hotspot is identified there, BN5 provides a concise, evidence-based answer to the follow-up question: “was this surge mainly driven by weather, or is something else happening?”

By keeping the weather logic simple, transparent, and grounded in conformed Date and Location dimensions, this page satisfies BN5 while remaining consistent with the dissertation’s broader emphasis on interpretable, auditable decision support rather than opaque predictive models.

5. RESULTS AND DISCUSSION

5.1 SCENARIO-BASED EVALUATION

To explore how the artefact supports BN1–BN5 in realistic decision situations, four scenarios were constructed using 2024 data (Table 5-1). Each scenario simulates a typical question raised by municipal managers and shows how different dashboard pages, filters, and interactions are combined to obtain an answer.

Table 5-1 - Scenario-based walkthroughs: mapping business questions, evaluation criteria, and observable evidence (BN1–BN5).

Scenario	Scope (BN/pages)	Main BQs addressed	Criteria evidenced	Observable outputs/evidence
City-wide workload in 2024	BN1	BQ1.1– BQ1.3	C1, C2, C4	Monthly profile + parish distribution + thematic composition via hover/drill interactions
Is 2024 unusually busy?	BN2	BQ2.1– BQ2.4	C1, C2, C4	Cross-year comparison + adaptive baselines + deviation matrix enabling identification of exceptional days and responsible categories
Second-semester anomaly wave	BN3	BQ3.1– BQ3.4	C1, C2, C5	Temporal anomaly profile + hotspot map + top anomaly days breakdown supporting localisation and diagnosis of sustained surges
Prioritising SANEAMENTO and adding weather context	BN4 & BN5	BQ4.1– BQ4.3 and BQ5.1– BQ5.4	C1, C2, C5 (and C4 via preserved filters)	Territory rankings + workload–anomaly quadrant + pressure index + Rainy Lift by category/parish supporting shortlisting and contextual explanation

Reproducibility note (filter protocol). All scenarios use the shared slicer panel (Date Hierarchy / Month / Parish Name / Area/Type) to define the analysis scope, ensuring a single filter context across visuals and pages. When a two-year scope (2023–2024) is required to populate YoY indicators, year-specific interpretation is ensured by selecting the target year directly on the temporal chart (Incidents by Month and Year). This interaction constrains the remaining visuals—especially spatial aggregation—to a

single year, preventing misleading accumulation across years. The year selection is therefore an explicit step of the scenario protocol, not an accidental cross-filter.

5.1.1 Scenario 1 – City-wide workload in 2024 (BN1)

Scenario 1 (BN1) starts by positioning the analysis under a comparative temporal context while preserving the interpretability of the spatial view (Figure 5-1). On the BN1 – Workload & Spatial Distribution page, the Date Hierarchy slicer is set to a 2-year window (2023 and 2024), with Month, Parish Name, and Area/Type left at All. This multi-selection enables the comparative KPIs to be populated (Incidents (LY) and Incidents YoY %), providing an immediate high-level signal of change between the current year in focus and the prior-year reference. The year displayed in the visuals is then controlled interactively through the temporal chart (Incidents by Month and Year): selecting the target year (e.g., 2024) in the chart constrains the remaining visuals, most importantly the bubble map to that specific year, avoiding the misleading effect of summing workloads across years in the spatial aggregation. In this way, the page supports both comparison (via two-year context and YoY/LY indicators) and year-specific spatial interpretation (via explicit year selection on the chart) within a single, traceable filter context.

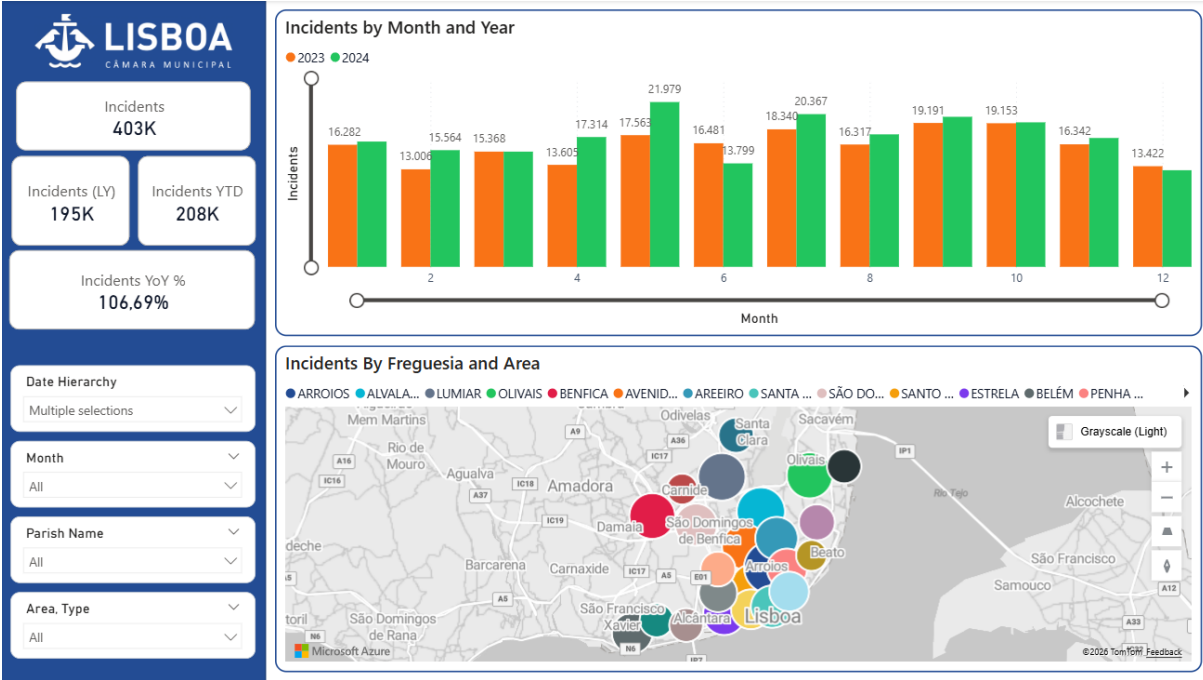


Figure 5-1 - BN1 - Workload & Spatial Distribution

The top chart, Incidents by Month and Year, displays monthly totals for both years side by side, enabling rapid identification of within-year patterns and cross-year differences. The 2024 bars reveal pronounced intra-annual variation, with peaks above 20k incidents in spring and late summer months, while other months remain closer to the mid-teens. Compared with 2023, several months show visibly higher totals, which is consistent with the YoY KPI signal. This directly addresses BN1’s temporal workload

component by showing that demand is neither flat within the year nor stable across years, implying that operational planning should anticipate recurring peaks and adjust resourcing accordingly.

The spatial perspective is provided by the Incidents by Freguesia and Area bubble map. With no parish filter applied, bubble size encodes the total incident volume per parish and colour encodes the thematic Area. Under the two-year slicer window, spatial aggregation would otherwise reflect combined totals; therefore, the walkthrough fixes the spatial interpretation by selecting the target year in the temporal chart before reading the map. This produces a year-specific spatial snapshot (e.g., 2024) that supports the identification of high-load parishes and allows the coordinator to move from city-wide demand to territorial concentration without conflating years.

Finally, thematic composition is exposed through a tooltip interaction that supports drill-down without altering the global slicer state. In the complementary tooltip screenshot, the cursor hovers over Olivais, triggering a tooltip chart (Incidents by Area and Parish Name) that decomposes the parish workload by thematic Area under the active year focus. The example (Figure 5-2) shows a dominant contribution from Higiene Urbana, followed by Iluminação Pública and Árvores e Espaços Verdes, with smaller contributions from categories such as Estradas e Sinalização and Saneamento. This interaction completes the BN1 decision path by linking (i) overall volume and year-over-year signal, (ii) monthly workload pattern, (iii) parish-level spatial concentration, and (iv) thematic composition, while keeping the filtering logic transparent and auditable through explicit year selection and on-hover evidence.

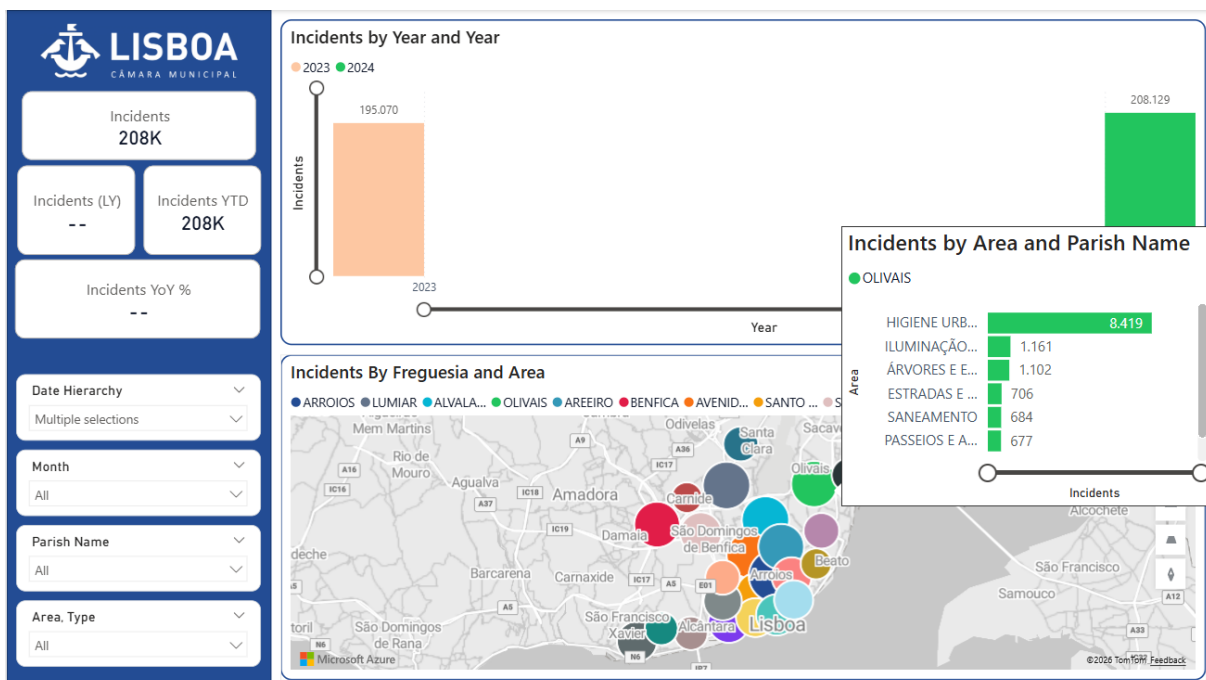


Figure 5-2 - BN1 - Workload & Spatial Distribution Tooltip

5.1.2 Scenario 2 – Is 2024 unusually busy? (BN2)

The second scenario asks whether 2024 workloads are genuinely exceptional or simply reflect normal seasonal variation (Figure 5-3). This is addressed on the BN2 – Seasonality & Baselines page, which operationalises workload interpretation by using explicit seasonal baselines and deviations rather than relying solely on raw incident counts.

In the first step, the Date Hierarchy slicer is set to a 2-year window (2023 and 2024), keeping Month, Parish Name, and Area/Type set to All. This multi-year context enables a direct comparison between years in the Incidents vs Baseline (Daily / Monthly) combo chart, where columns represent monthly incident totals (distinguished by Year in the legend) and the baseline line provides an adaptive reference level at the same aggregation grain. Visual inspection shows that several months in 2024 sit above both the corresponding 2023 levels and the baseline, suggesting that higher volumes are not merely sporadic noise but reflect a sustained elevation relative to recent historical expectations.

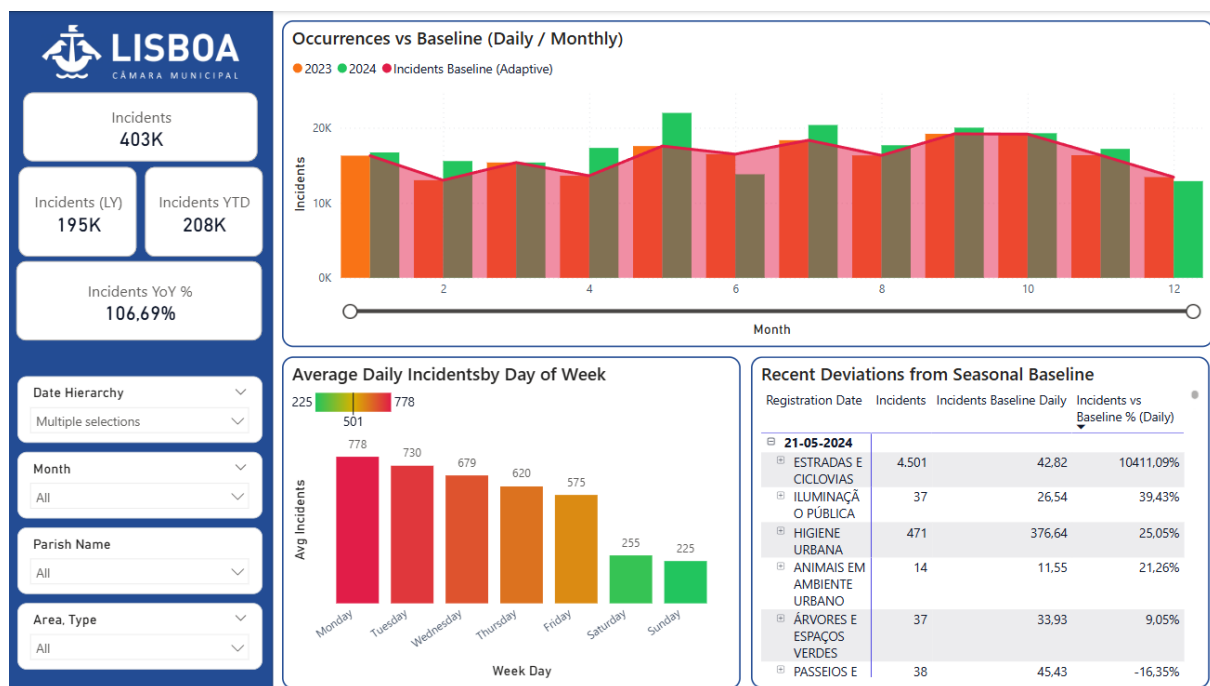


Figure 5-3 - BN2 - Seasonality & Baselines

To move from global patterns to concrete operational signals, a second interaction is performed: the user clicks the May 2024 column in the combo chart. This selection constrains the analysis window and cross-filters the lower visuals to the selected month (Figure 5-4). The Average Daily Incidents by Day of Week chart now summarises the weekday profile for May 2024, showing whether the month’s workload is driven by a general weekday uplift or by specific high-pressure days. In parallel, the Recent

Deviations from Seasonal Baseline matrix lists only dates in the selected context and quantifies, side by side, the observed incidents, the daily baseline reference, and the percentage deviation from baseline.




Figure 5-4 - BN2 - Seasonality & Baselines selected

Within the matrix, 21-05-2024 appears as an extreme outlier, with incident volume exceeding the daily baseline by a wide margin. Expanding the hierarchy for that date reveals which thematic Areas drive the deviation, with Estradas e Ciclovias dominating the spike. This sequence from cross-year monthly comparison to a selected month, to weekday profile, to a specific day and responsible category demonstrates how BN2 transforms raw counts into interpretable, auditable signals. Managers can therefore substantiate claims of “exceptional days” by explicitly comparing them to expected values, distinguishing genuine anomalies from ordinary seasonal peaks, and identifying the operational domain most likely to require follow-up.

5.1.3 Scenario 3 – Second-semester anomaly wave (BN3)

The third scenario investigates whether the second semester of 2024 contains a period of elevated anomaly signals in citizen reporting, and where (spatially) and why (thematically) those anomalies concentrate (Figure 5-5). This is addressed on the BN3 – Anomalies & Hotspots page.

In the first step, the Date Hierarchy slicer is set to 2024 (Year) + 2nd Semester, with Month, Parish Name, and Area/Type set to All. The KPI strip updates to summarise

the workload under this half-year window, establishing the baseline context for the subsequent anomaly interpretation.

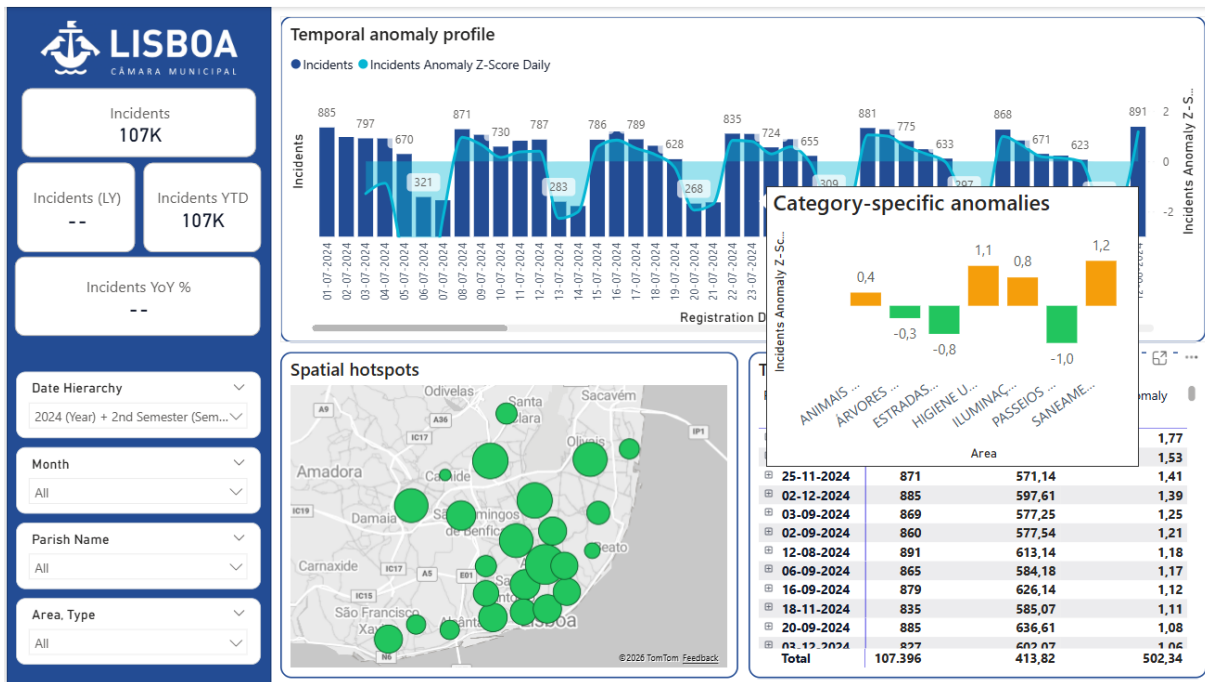


Figure 5-5 - BN3 - Anomalies & Hotspots

The Temporal anomaly profile chart displays daily incidents as bars and the corresponding Incidents Anomaly Z-Score Daily as a line. The Z-score operationalises anomaly intensity as the deviation from a rolling 28-day daily baseline, expressed in standard-deviation units; sustained positive values therefore indicate a cluster of days with incident volumes above the expected seasonal level rather than isolated noise (Figure 5-6). The chart indicates a clear mid-summer cluster: around late July, multiple days show elevated Z-scores relative to the baseline, consistent with an “anomaly episode” rather than a single outlier. Selecting 22-07-2024 enables the coordinator to anchor the investigation on a single representative peak day and propagate that context across the remaining visuals.

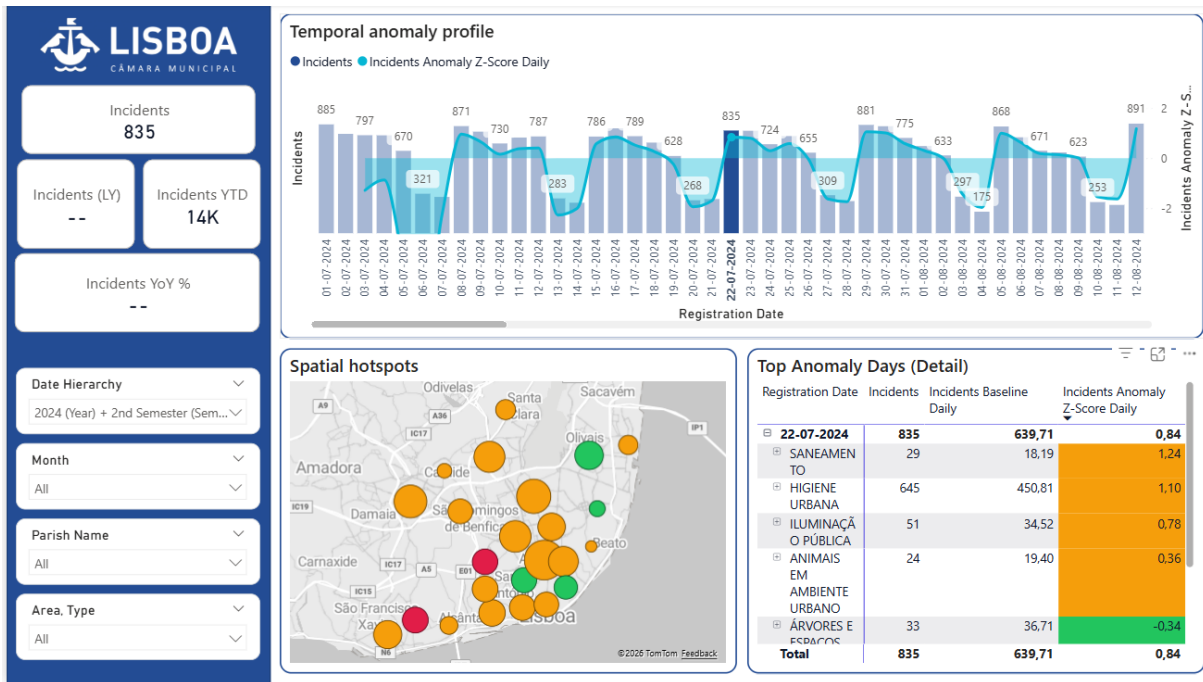


Figure 5-6 - BN3 - Anomalies & Hotspots Selected

In parallel, the Spatial hotspots map shows that the anomaly episode is not geographically uniform. Under the selected context, some parishes exhibit consistently higher anomaly intensity (warm colours), while others remain closer to baseline (cool/neutral colours). This supports an operational insight that the event driving the anomaly is likely localised—either concentrated in specific territories or driven by geographically clustered conditions—rather than a city-wide reporting surge. As a managerial implication, this helps focus field verification and resource allocation on the subset of parishes with the strongest evidence of abnormal pressure.

The Top Anomaly Days (Detail) matrix closes the loop by linking anomalous dates to attributable thematic drivers. Sorted by anomaly intensity, it lists 22-07-2024 among the top days and, when expanded, reveals the breakdown by Area. The conditional formatting is explicitly rule-based on Incidents Anomaly Z-Score Daily ($Z \geq 2$ in red; $0 \leq Z < 2$ in yellow; $-100 \leq Z < 0$ in green), making the anomaly criterion transparent and auditable. In the selected example, the decomposition indicates that the elevated workload is not evenly distributed across categories: a small subset of Areas contributes disproportionately to the anomaly signal, while other Areas remain near baseline or below. This produces a concrete investigative shortlist: (i) verify whether the dominant contributing Areas correspond to a known operational disruption or external event, and (ii) check whether the same Areas repeatedly dominate across adjacent anomaly days, which would indicate a stable underlying driver rather than a one-off reporting artifact.

Overall, the scenario yields three decision-relevant insights: (1) temporal evidence of an anomaly episode (clustered above baseline days rather than random fluctuation); (2) spatial evidence of hotspot concentration (anomaly pressure is uneven across parishes); and (3) thematic evidence of attribution (specific incident domains drive the anomaly). Together, these insights allow municipal managers to move from detection to prioritised response: they can justify escalation based on standardised anomaly criteria, target interventions geographically, and assign follow-up to the most relevant operational departments.

5.1.4 Scenario 4 – Prioritising Saneamento and Adding Weather Context (BN4 & BN5)

The fourth scenario illustrates how the artefact supports prioritisation under capacity constraints by combining (i) territory-focused workload and anomaly evidence (BN4) with (ii) a pragmatic, explanatory integration of weather signals (BN5). The analysis deliberately keeps a consistent filter context across pages via the shared slicer panel (Year / Month / Parish / Area), ensuring that BN4 and BN5 are comparable views over the same semantic scope.

4A – Pressure for SANEAMENTO in 2024 (BN4)

On the BN4 – Territory Prioritisation (Categories & Hotspots) page, the filter panel is set to Year = 2024, with Month = All, Parish Name = All, and Area/Type = SANEAMENTO (*Figure 5-7*). In this configuration, the page operationalises a “prioritisation cockpit” by combining ranking, anomaly-aware comparison, and equity-oriented workload normalization (*Figure 5-7*).

First, the Parish workload ranking bar chart provides an immediate shortlist of territories that concentrate sanitation demand under the selected year and category. In the documented filtered state, Olivais, Lumiar, and Arroios appear among the highest volume parishes, indicating where sanitation-related reporting is structurally concentrated. This addresses the territorial ranking logic of BN4 by making the volume dimension explicit and easy to compare across parishes.

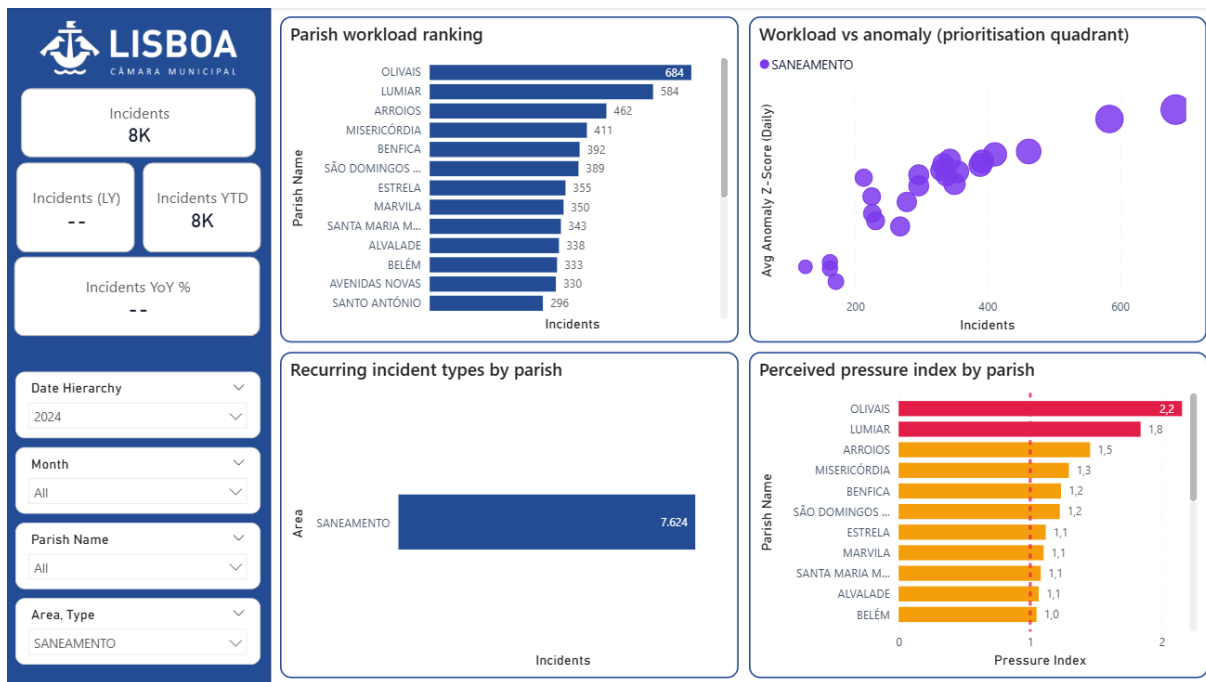


Figure 5-7 - BN4 - Territory Prioritisation Categories & Hotspots

Second, the Workload vs anomaly (prioritisation quadrant) scatterplot complements raw volume by adding behavioural abnormality: it relates parish workload ([Incidents]) to average anomaly intensity (Avg Incidents Anomaly Z-Score Daily). Parishes located toward the upper-right of the plot are interpreted as candidates for sustained operational pressure territories that are not only busy but also tend to exceed baseline expectations more frequently. This supports BN4’s intent to balance routine workload against persistent above baseline behaviour, rather than treating all high volume territories as equally urgent.

Third, the Perceived pressure index by parish provides an equity-adjusted lens by normalising incident workload relative to an average reference (Pressure Index = 1). The visual, therefore, separates “busy because the city is busy everywhere” from “systematically above average exposure,” allowing managers to flag high-pressure territories for sanitation under limited inspection capacity. Importantly, BN4’s value here is not merely descriptive: it produces an interpretable prioritisation shortlist that can be justified using auditable indicators rather than subjective impressions.

4B – Weather sensitivity of SANEAMENTO (BN5)

To test whether sanitation pressure is plausibly weather-driven (as opposed to purely seasonal baselines or unrelated operational factors), the analysis moves to the BN5 – Weather Context page while keeping exactly the same filter context (Year = 2024; Area/Type = SANEAMENTO; Month = All; Parish = All) (Figure 5-8). BN5 is explicitly designed as an explanatory layer: it integrates weather using simple, auditable rainy-versus-dry logic at the parish–day level, rather than as a predictive forecasting model (Figure 5-8).

At the event level, the Weather-aware temporal profile juxtaposes daily incident volume with precipitation on a secondary axis, enabling direct inspection of co-variation within the selected category. In the documented window, several workload peaks occur close to, or following, precipitation peaks—an interpretable pattern suggesting that part of the sanitation workload may be sensitive to rainfall conditions. This answers BN5’s core diagnostic question (“is a surge plausibly weather-related?”) using visible, traceable evidence.

At the category level, the Weather-sensitive categories chart confirms the interpretation through a proportional indicator: Rainy Lift, defined as the ratio of incidents on rainy parish–days to incidents on dry parish–days. With the analysis scoped to SANEAMENTO, the chart directly communicates that sanitation exhibits a Rainy Lift above 1, i.e., incidents occur more frequently on rainy days than on dry days under the same semantic conditions. This operationalises the claim that adding weather provides explanatory power beyond “counts only,” while keeping the logic transparent and reproducible.

Finally, BN5 adds a territorial vulnerability layer that BN4 alone cannot provide: the Rainy lift by parish ranking and the Weather sensitivity by parish map reveal that rainfall sensitivity is not uniform across Lisbon. In the documented state, parishes such as Avenidas Novas, Belém, and São Vicente appear among the most rain-sensitive territories for sanitation (high Rainy Lift), while others sit closer to neutral response. This produces a second, complementary prioritisation shortlist: not necessarily “where sanitation is highest,” but “where sanitation is most likely to surge when rain occurs.”

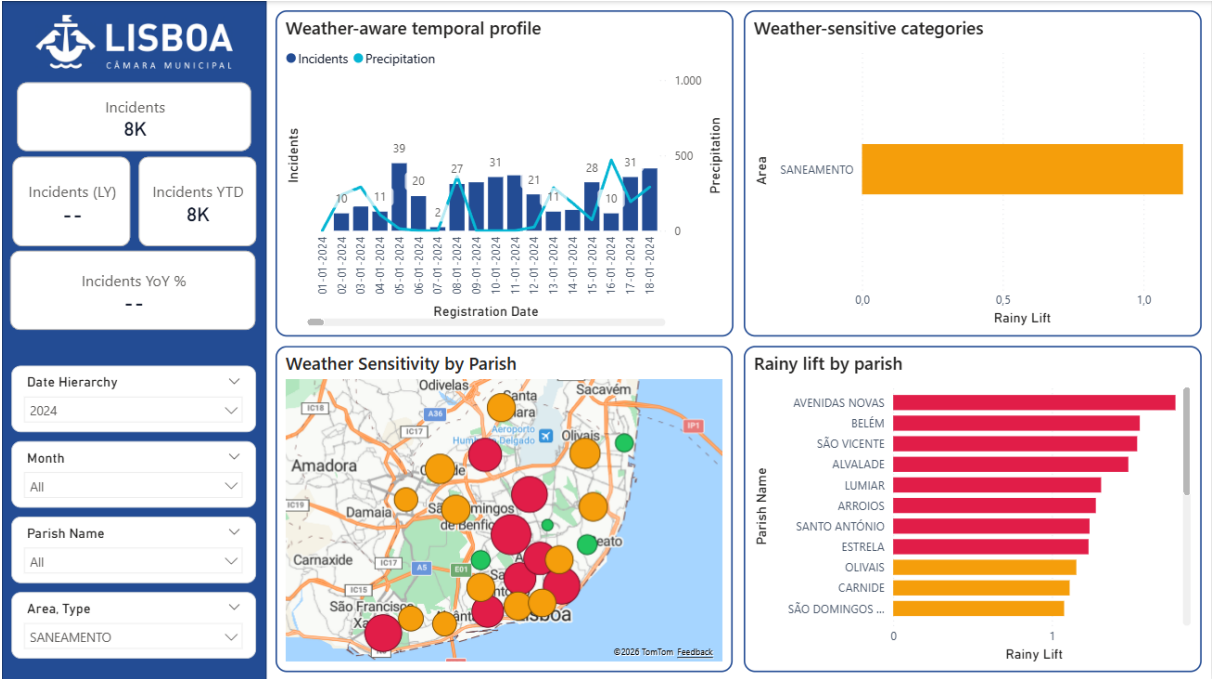


Figure 5-8 - BN5 - Weather Context

Integrated decision insight (BN4 + BN5). Taken together, the two pages support an evidence-based prioritisation narrative: (i) sanitation workload and perceived pressure are structurally concentrated in a subset of parishes (BN4); and (ii) within sanitation, some territories are disproportionately sensitive to rainfall (BN5). This combination justifies targeted actions under constraints—e.g., focusing inspections and preventive maintenance on high-pressure parishes while also using the rain-sensitivity shortlist to strengthen contingency planning during rainfall events, without resorting to opaque forecasting claims.

5.2 PRINCIPAL PATTERNS IN NA MINHA RUA LX INCIDENT REPORTING

The scenario-based evaluation not only demonstrates the functional adequacy of the dashboards but also reveals consistent empirical patterns in Lisbon’s urban incident reporting dynamics. When analysed across time, territory, and incident typologies, several structured regularities emerge.

5.2.1 Temporal Patterns

Incident volumes exhibit stable seasonal variation across the analysed period (2020–2024). Daily and monthly baselines indicate cyclical fluctuations rather than random dispersion. Year-on-year comparisons show relative stability in overall workload magnitude, with episodic peaks exceeding expected seasonal baselines.

Anomaly flags highlight concentrated surges in specific periods, distinguishing structural seasonality from exceptional events. This differentiation is analytically significant because it prevents overreaction to predictable cycles while enabling early identification of abnormal increases in workload.

These temporal regularities suggest that citizen reporting behaviour follows patterned rhythms rather than erratic dynamics.

5.2.2 Territorial Distribution

Spatial analysis at the parish level reveals uneven reporting intensity across Lisbon. Certain parishes consistently register higher average daily incidents, even when normalised by city-wide averages.

The Pressure Index reinforces this asymmetry, showing persistent territorial concentration rather than homogeneous distribution. This indicates that urban incident reporting reflects localised structural conditions rather than uniform city-wide exposure.

Importantly, the dashboards enable proportional reasoning rather than absolute ranking, mitigating the risk of simplistic comparisons between parishes of different sizes or characteristics.

5.2.3 Typological Concentration

Incident categories display heterogeneous frequencies. Some thematic Areas and Types dominate reporting volumes, while others remain episodic.

This concentration suggests that specific categories represent recurrent operational pressure points. The ability to filter by Area and Type allows managers to distinguish between structurally recurrent issues and isolated events.

The typological breakdown thus provides a foundation for targeted prioritisation rather than reactive allocation.

5.2.4 Weather Sensitivity

Weather integration, particularly precipitation data aligned at daily grain, indicates that selected incident categories exhibit higher reporting frequency during rainy conditions. While no causal inference is made, rainy/dry splits and Rainy Lift measures reveal directional associations consistent with plausible environmental effects.

This contextual sensitivity enhances interpretability by distinguishing endogenous urban patterns from exogenous environmental influences.

5.2.5 Synthesis of Empirical Patterns

Across scenarios, four structured regularities emerge:

1. Stable seasonality with episodic surges;
2. Persistent territorial concentration;
3. Category-level workload asymmetry;
4. Weather-sensitive reporting variation.

These patterns demonstrate that the dataset contains structured analytical signals rather than unorganised noise. The dimensional architecture enables these signals to be consistently detected, contextualised, and interpreted.

To synthesise the main empirical regularities identified across scenarios, Table 5-2 summarises the principal patterns observed in Na Minha Rua LX incident reporting. The table consolidates temporal, territorial, typological, and environmental dimensions, linking each observed pattern to its analytical relevance and corresponding decision-support implications. This structured overview clarifies how the dimensional architecture transforms raw reporting data into interpretable analytical signals that can inform prioritisation and monitoring in municipal governance.

Table 5-2 - Summary of main empirical findings.

Dimension	Observed Pattern	Analytical Relevance	Decision-Support Implication
Time	Seasonal cyclicity with episodic anomalies	Distinguishes expected variation from abnormal surges	Supports proactive monitoring and surge detection
Territory	Uneven parish-level concentration	Reveals structural spatial asymmetries	Enables proportional prioritisation
Typology	Dominance of specific incident categories	Identifies recurrent pressure points	Supports targeted resource allocation
Weather	Higher incident rates during rainy days (selected categories)	Suggests environmental sensitivity	Improves contextual interpretation

5.3 DISCUSSION

The empirical analysis demonstrates that citizen-reported urban incident data in Lisbon exhibit structured regularities across time, territory, typology, and environmental context. These patterns are neither random nor purely subjective; they reveal stable seasonal rhythms, persistent territorial asymmetries, recurrent thematic concentrations, and weather-sensitive variations. Such regularities reinforce the argument that citizen-generated urban data, when systematically structured, can support meaningful governance insights, a point increasingly highlighted in the smart city and urban analytics literature (Chamusca, 2025). Moreover, research on civic reporting platforms suggests that structured citizen input can evolve into a strategic urban intelligence resource when embedded within appropriate analytical infrastructures (Soares & Coutinho, 2024).

However, the central contribution of this dissertation does not lie merely in identifying these patterns. Rather, it lies in demonstrating how disciplined dimensional modelling enables such patterns to be consistently detected, interpreted, and compared transparently.

The research question asked how Power BI dashboards built from Na Minha Rua LX data can improve decision support in managing urban incident reporting. The findings suggest that improvement arises primarily from structural clarity rather than analytical complexity. By enforcing a clear analytical grain and implementing conformed dimensions across time, location, incident type, and reporting channel, the artefact reduces ambiguity in interpretation and ensures that measures remain comparable across dashboards and scenarios. This finding resonates with service design research in e-government, which argues that specific structural design characteristics, rather than generic adoption constructs, drive perceptions of service quality and satisfaction (Chan et al., 2021). Likewise, systematic reviews of information quality highlight that

system-level structuring mechanisms are central to ensuring reliability and interpretability across contexts (Abd Aziz et al., 2024).

Seasonal baselines and anomaly flags illustrate this structural advantage. Instead of reacting to raw fluctuations, decision-makers receive contextualized indicators that distinguish predictable cyclical patterns from exceptional surges. Similarly, territorial pressure metrics move beyond absolute counts and enable proportional reasoning, avoiding simplistic ranking while highlighting persistent spatial concentration. Such contextualisation mechanisms are aligned with public-sector information quality dimensions such as timeliness, comparability, and analytical coherence (Abdulkareem et al., 2022). Even the integration of weather data, treated pragmatically, reinforces interpretability by situating workload variation within broader environmental conditions without introducing opaque modelling assumptions. In line with recent research on public-sector information quality, contextualized indicators enhance interpretability by embedding timeliness and comparability within the analytical design (Chan et al., 2021).

From a Design Science Research perspective, the artefact demonstrates utility within its intended organizational context. The contribution is best understood as an instantiation of dimensional modelling principles in a smart-city citizen-reporting environment. While dimensional warehousing is not new, its disciplined application to open urban incident data—combined with a Medallion architecture and explicit semantic-layer governance—provides a replicable framework for public-sector analytics. This responds to calls in contemporary governance research for stronger integration between data infrastructures and decision-support capabilities in municipal environments (Clausen et al., 2025). Furthermore, recent work on smart-city data ecosystems stresses that analytical value emerges not from data availability alone but from structured modelling practices that ensure traceability and reuse (Idzi & Gomes, 2022). The study, therefore, contributes to design knowledge rather than to algorithmic innovation.

Importantly, the dissertation also highlights a broader insight: in public-sector environments characterized by accountability requirements and heterogeneous stakeholders, interpretability and traceability may be more valuable than predictive sophistication. The artefact shows that structured semantic modelling can enhance transparency and monitoring capacity without resorting to complex machine-learning techniques. This perspective resonates with digital governance scholarship, which argues that trust and legitimacy in public-sector data systems depend on clarity, explainability, and organisational alignment rather than on analytical sophistication alone (Choi & Song, 2020; Larsen et al., 2025; Omol et al., 2025). In this sense, the work bridges the gap between open-data availability and actionable analytical governance, a connection also emphasised in prior work on evaluative frameworks for smart urban initiatives (Neves et al., 2020).

At the same time, the findings must be interpreted within clear boundaries. The dashboards measure reported workload rather than municipal service performance. Incident counts reflect citizen reporting behavior and perceived urban conditions, not operational efficiency or resolution effectiveness. Furthermore, the evaluation remains scenario-based and does not include controlled user testing with municipal staff. The weather analysis identifies associations but does not establish causal relationships. These constraints delimit the scope of the contribution while preserving its validity.

Overall, the study demonstrates that improving decision support in urban incident management does not necessarily require greater analytical sophistication, but rather greater structural discipline. By transforming open citizen-reporting data into a conform, transparent, and reproducible dimensional model, the artefact enhances interpretability, consistency, and prioritization capacity in municipal governance contexts (Neves et al., 2020).

6. CONCLUSIONS AND FUTURE RESEARCH

This dissertation set out to bridge the translation gap between raw citizen-reported incident logs from Na Minha Rua LX and structured decision support for municipal management in Lisbon. Rather than approaching the problem through predictive modelling or algorithmic sophistication, the study adopted a dimensional data warehousing strategy, implemented through a Kimball-style architecture, a Medallion (Bronze–Silver–Gold) pipeline in Microsoft Fabric, and a governed Power BI semantic layer.

The findings indicate that meaningful improvements in decision support arise primarily from structural clarity. By defining a stable analytical grain, enforcing conformed dimensions (Date, Location, Incident Type), and ensuring semantic consistency across dashboards, the artefact transforms heterogeneous reporting logs into interpretable analytical signals. Seasonal baselines, anomaly indicators, territorial pressure metrics, and weather-context measures do not introduce analytical opacity; instead, they contextualise workload patterns to support prioritisation and monitoring without sacrificing transparency.

Empirically, the dashboards reveal structured regularities in Lisbon’s urban incident reporting: stable seasonal cycles punctuated by episodic surges, persistent territorial concentration, category-level asymmetries, and directional associations with rainfall in selected domains. These patterns are not presented as causal findings, but as evidence that citizen-generated data contain meaningful analytical structure when appropriately modelled. The dimensional architecture makes such a structure visible, comparable, and traceable.

From a Design Science Research perspective, the contribution lies in demonstrating how dimensional discipline can be operationalised within a smart-city open-data environment. The work delivers a fully implemented artefact in its target context, a formalised dimensional representation of the problem domain (bus matrix, declared grains, additive fact tables), and a set of reusable design patterns that connect governance principles to concrete implementation choices. In doing so, the dissertation reinforces the argument that interpretability, traceability, and semantic consistency are foundational enablers of reliable decision support in public-sector settings.

Importantly, the study also advances a broader insight: in municipal governance environments characterised by accountability requirements and heterogeneous stakeholders, structural modelling choices may have greater practical impact than increasing analytical sophistication. The artefact demonstrates that transparent baselines and contextual indicators, embedded within a governed semantic layer, can support informed prioritisation without reliance on opaque predictive systems.

The work is subject to clear limitations. The dashboards measure reported incident pressure rather than service performance, as resolution workflows were out of scope.

Evaluation was scenario-based and researcher-driven, without controlled usability testing involving municipal staff. Weather integration identifies descriptive associations but does not establish causal relationships. Furthermore, the solution operates at daily granularity and in batch mode, leaving real-time analytics and broader exogenous drivers for future exploration.

These constraints point to concrete directions for further research. Integrating service workflow data would enable joint analysis of demand and organisational responsiveness. Structured usability studies could assess adoption dynamics and refine interaction design. Expanding the dimensional model with additional contextual datasets and comparing interpretable baselines with more complex anomaly-detection techniques would further clarify the trade-off between analytical sophistication and transparency in municipal decision environments.

Overall, this dissertation demonstrates that improving decision support in urban incident management does not necessarily require greater algorithmic complexity, but rather greater semantic discipline. By operationalising dimensional modelling within an open-data smart-city context, the study contributes a replicable framework for transforming citizen-generated reporting logs into consistent, interpretable, and governance-aligned analytical artefacts.

REFERENCES

- Abd Aziz, A., Nor, R. N. H., Jusoh, Y. Y., Rahman, W. N. W. A., & Ali, N. M. (2024). Factors influencing information quality of information systems: A systematic literature review. *JOIV: International Journal on Informatics Visualization*, 8(3–2), 1923–1931. <https://doi.org/10.62527/joiv.8.3-2.3483>
- Abdulkareem, A. K., Abdulkareem, Z. J., Ishola, A. A., & Akindele, I. T. (2022). Does e-government impact e-participation? The influence of trust in e-government. *International Review of Public Administration*, 27(2), 91–110. <https://doi.org/10.1080/12294659.2022.2071540>
- Arifani, K. B., Pintarko, D., Sari, A. P., & Agussalim, A. (2025). Real-Time Data Integration and Weather Reporting Automation with Cloud Computing-based Interactive Spatial Dashboard for Extreme Weather Risk Analytics in Indonesia. *Journal of Computer Networks, Architecture and High Performance Computing*, 7(3), 816–824. <https://doi.org/10.47709/cnahpc.v7i3.6298>
- Brussels-Capital Region. (2024). *FixMyStreet Brussels: Report incidents in the open space*. <https://fixmystreet.brussels/>
- Câmara Municipal de Lisboa. (2026). *Lisboa Aberta: Dados abertos da cidade*. <https://lisboaaberta.cm-lisboa.pt/index.php/pt/dados/grupos>
- Chamusca, P. (2025). Governance, participation, and development: Local power's perspective in Portugal. *International Planning Studies*, 30(3), 292–318. <https://doi.org/10.1080/13563475.2025.2457656>
- Chan, F. K., Thong, J. Y. L., Brown, S. A., & Venkatesh, V. (2021). Service design and citizen satisfaction with e-government services: A multidimensional perspective. *Public Administration Review*, 81(5), 874–894. <https://doi.org/10.2139/ssrn.3719710>

- Choi, J. C., & Song, C. (2020). Factors explaining why some citizens engage in E-participation, while others do not. *Government Information Quarterly*, 37(4), 101524. <https://doi.org/10.1016/j.giq.2020.101524>
- Clausen, J. B. B., Møller, C., Asmussen, C. B., & Jørgensen, S. L. (2025). Self-service business intelligence systems in production domains: Introducing data democratization and mass customizable mashups. *Digital Business*, 100156. <https://doi.org/10.1016/j.digbus.2025.100156>
- Databricks. (2025). *What is the Medallion Lakehouse Architecture?* <https://www.databricks.com/glossary/medallion-architecture>
- Gonçalves, J. E., Ioannou, I., & Verma, T. (2024). No one-size-fits-all: Multi-actor perspectives on public participation and digital participatory platforms. *Philosophical Transactions A*, 382(2285), 20240111. <https://doi.org/10.1098/rsta.2024.0111>
- Gregor, S., Martin, M., Fernandez, W., Stern, S., & Vitale, M. (2006). The Transformational Dimension in the Realization of Business Value from Information Technology. *The Journal of Strategic Information Systems*, 15(3), 249–270. <https://doi.org/10.1016/j.jsis.2006.04.001>
- Idzi, F. M., & Gomes, R. C. (2022). Digital governance: Government strategies that impact public services. *Global Public Policy and Governance*, 2(4), 427–452. <https://doi.org/10.1007/s43508-022-00055-w>
- João Ramos. (2026). *Improving-Decision-Support-with-Power-BI-Using-Na-Minha-Rua-Lx-Data*. <https://github.com/joaopedrofbz/Improving-Decision-Support-with-Power-BI-Using-Na-Minha-Rua-Lx-Data>
- Kimball, R., & Ross, M. (2013). *The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling* (3rd ed.). Wiley. <https://www.wiley.com/en->

us/The%2BData%2BWarehouse%2BToolkit%3A%2BThe%2BDefinitive%2BG
uide%2Bto%2BDimensional%2BModeling%2C%2B3rd%2BEdition-p-
9781118530801

Larsen, K. R., Lukyanenko, R., Mueller, R. M., Storey, V. C., Parsons, J., VanderMeer, D., & Hovorka, D. S. (2025). Validity in design science. *MIS Quarterly*, 49(4), 1267–1294. <https://doi.org/10.25300/MISQ/2024/18064>

Matheus, R., Janssen, M., & Maheshwari, D. (2018). Data science empowering the public: Data-driven dashboards for transparent and accountable decision-making in smart cities. *Government Information Quarterly*, 37(3), 101284. <https://doi.org/10.1016/j.giq.2018.01.006>

Meijer, A., & Bolívar, M. P. R. (2016). Governing the smart city: A review of the literature on smart urban governance. *International Review of Administrative Sciences*, 82(2), 392–408. <https://doi.org/10.1177/0020852314564308>

Neves, F. T., de Castro Neto, M., & Aparicio, M. (2020). The impacts of open data initiatives on smart cities: A framework for evaluation and monitoring. *Cities*, 106, 102860. <https://doi.org/10.1016/j.cities.2020.102860>

OECD. (2019). *A Data-Driven Public Sector: Enabling the Strategic Use of Data for Productive, Inclusive and Trusted Government*. <https://doi.org/10.1787/09ab162c-en>

Omol, E. J., Mburu, L. W., & Abuonji, P. A. (2025). Digital maturity assessment model (DMAM): Assimilation of design science research (DSR) and capability maturity model integration (CMMI). *Digital Transformation and Society*, 4(2), 128–152. <https://doi.org/10.1108/DTS-04-2024-0049>

- Pak, B., Chua, A., & Vande Moere, A. (2017). FixMyStreet Brussels: Socio-Demographic Inequality in Crowdsourced Civic Participation. *Journal of Urban Technology*, 24(2), 65–87. <https://doi.org/10.1080/10630732.2016.1270047>
- Peppers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A Design Science Research Methodology for Information Systems Research. *Journal of Management Information Systems*, 24(3), 45–77. <https://doi.org/10.2753/MIS0742-1222240302>
- Schiff, K. J. (2025). Does collective citizen input impact government service provision? Evidence from SeeClickFix requests. *Public Administration Review*, 85(1), 32–45. <https://doi.org/10.1111/puar.13747>
- Soares, J., & Coutinho, C. (2024). Urban issue reporting applications towards government 2.0. *2024 8th International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT)*, 1–6. <https://doi.org/10.1109/ISMSIT63511.2024.10757230>
- Volinz, L., Melgaço, L., & Steenhout, I. (2025). From Participatory to Lateral Surveillance: Municipal Apps as Platforms for Digital Informants. *Surveillance & Society*, 23(3), 273–286. <https://doi.org/10.24908/ss.v23i3.18357>
- Zippenfenig, P. (2023). *Open-Meteo.com Weather API* [Computer software]. Zenodo. <https://doi.org/10.5281/zenodo.7970649>

APPENDIX A

Table A-1 - Operational criteria for evaluating “improving decision support.”

Criterion	Operational definition	How is it evidenced in this dissertation
C1 – Coverage (BN/BQ coverage)	The extent to which the dashboards and semantic model answer the defined business questions (BN1–BN5) with implemented measures and visuals.	Demonstrated by the Scenario→BQs mapping (Table 5-1) and by walkthrough outputs showing each scenario resolves the corresponding questions.
C2 – Interpretability (auditable logic)	The extent to which indicators are understandable and defensible by municipal managers, privileging simple baselines, deviations, and transparent anomaly criteria over opaque models.	Demonstrated through documented measure logic (baselines, deviations, anomaly signals, pressure/ranking measures) and the step-by-step reasoning in the scenario walkthroughs.
C3 – Traceability (lineage and reproducibility)	The extent to which results can be traced from a visual to the underlying measure logic and, when needed, to the supporting fact records and transformations.	Demonstrated by the semantic model documentation, drill-down paths, and explicit references to the artefact package (GitHub repository).
C4 – Semantic consistency (single source of truth)	The extent to which conformed dimensions and a shared filter context produce consistent results across pages, enabling coherent navigation from overview to drill-down.	Demonstrated by cross-page analysis under preserved filters and consistent aggregation at the declared grains (Day × Parish × Type).
C5 – Actionability (prioritisation support)	The extent to which dashboards support prioritisation, shortlisting, and hotspot identification under capacity constraints, including contextual explanation (e.g., weather-aware signals).	Demonstrated by scenarios that end with ranked outputs, quadrant positioning (workload vs anomaly), and parish/category shortlists derived from the dashboards.

Table A-2 - DimDate attribute dictionary

Field / hierarchy element	Data type	Role in model	Notes from TMDL
sk_date	int64 (hidden)	Surrogate key (PK)	Key used in relationships from both fact tables; hidden and non-null
Registration Date	dateTime (Underlying DateTimeData Type = Date)	Date attribute used for time intelligence	Main date column referenced in measures (DATESYTD, SAMEPERIODLASTYEAR)
Year	int64	Calendar year	Part of 'Year Hierarchy'
Semester	string	Semester label	Part of 'Year Hierarchy' (note: stored as string in TMDL)
Trimester	int64	Quarter/trimester	Part of 'Year Hierarchy'
Month	int64	Month number	Part of 'Year Hierarchy' and 'Month Hierarchy'
Day	int64	Day number	Part of 'Year Hierarchy'
Week Day	string	Weekday label	Used for weekday effects / baseline interpretation
Is Last Year	boolean	Convenience flag	Present in TMDL; supports filtering scenarios
Is Last Year-1 Until Last Period	boolean	Convenience flag	Present in TMDL; supports comparable period filtering
Hierarchy: Year Hierarchy	—	Navigation hierarchy	Levels: Year → Semester → Trimester → Month → Day
Hierarchy: Month Hierarchy	—	Navigation hierarchy	Levels: Month → Registration Date

Table A-3 - DimLocation attribute dictionary

Field	Data type	Role in model	Notes from TMDL
sk_location	int64 (hidden)	Surrogate key (PK)	Used as FK target from both fact tables
bk_location	string	Business key	Stable identifier for location (parish)
Parish Name	string	Display attribute	Primary slicer/map label
Latitude_Centroid	double	Spatial attribute	Used for mapping and for weather alignment (centroid latitude)
Longitude_Centroid	double	Spatial attribute	Used for mapping and for weather alignment (centroid longitude)

Table A-4 - DimIncidentType attribute dictionary

Field / hierarchy element	Data type	Role in model	Notes from TMDL
sk_incident_type	int64 (hidden, non-null)	Surrogate key (PK)	Joined from fact_incident_daily
bk_incident_type	string	Business key	Stable identifier for type (taxonomy key)
Area	string	Higher-level category	Used for grouping (e.g., Saneamento)
Type	string	Detailed category	Drill-down detail under Area
Hierarchy: Area Hierarchy	—	Navigation hierarchy	Levels: Area → Type

Table A-5 - FactIncidentDaily schema

Field	Data type	Role in model	Notes from TMDL
sk_fact_incident	int64 (hidden, non-null)	Surrogate key (PK)	Hidden; used for internal uniqueness
bk_fact_incident	string	Business identifier (degenerate dimension)	Supports traceability to source incident/report
sk_date	int64 (hidden, non-null)	FK to dim_date[sk_date]	Relationship defined (referential integrity)
sk_location	int64 (hidden, non-null)	FK to dim_location[sk_location]	Relationship defined (referential integrity)
sk_incident_type	int64 (hidden, non-null)	FK to dim_incident_type[sk_incident_type]	Relationship defined (referential integrity)

Table A-6 - FactIncidentDaily Measures

Measure name	Folder	Purpose (summary)
Incidents	Incidents\Core	Core count of incident rows (COUNTRWS)
Avg Incidents	Incidents\Core	Average incidents across dates in context
Incidents (LY)	Incidents\Time Intelligence	Prior-year incidents (SAMEPERIODLASTYEAR)
Incidents YTD	Incidents\Time Intelligence	Year to date incidents (DATESYTD)
Incidents YoY %	Incidents\Time Intelligence	YoY percentage change vs LY
Incidents Baseline Daily	Incidents\Baseline	28-day trailing mean excluding current day (RefDay = Today-1)
Incidents Baseline Daily StdDev	Incidents\Baseline	28-day trailing population std dev
Incidents Baseline Monthly	Incidents\Baseline	Average of monthly totals over last 12 complete months
Incidents Baseline Monthly StdDev	Incidents\Baseline	Std dev for monthly baseline window
Incidents Baseline (Adaptive)	Incidents\Baseline	Switches daily/monthly baseline by visual grain
Incidents Baseline StdDev (Adaptive)	Incidents\Baseline	Adaptive std dev by grain
Incidents vs Baseline % (Daily)	Incidents\Baseline	Daily deviation ratio vs daily baseline
Incidents Anomaly % Daily / Monthly	Incidents\Anomalies	Percent deviation vs baseline (daily/monthly)
Incidents Anomaly Z-Score Daily / Monthly	Incidents\Anomalies	Standardised deviation (daily/monthly)
Incidents Anomaly Z-Score (Adaptive)	Incidents\Anomalies	Adaptive Z-score by grain
Incidents Anomaly Flag (Daily/Monthly/Adaptive)	Incidents\Anomalies\Flags	Flag using Z-score threshold (thr = 2.0)
Anomaly Days Count / %	Incidents\Anomalies	Counts / share of anomalous days in context
Avg Anomaly Z-Score (Daily)	Incidents\Anomalies	Mean daily Z-score across dates
Avg Incidents per Parish	Incidents\Territory	Total incidents / number of parishes (ALL dim_location)
Pressure Index	Incidents\Territory	Incidents / Avg Incidents per Parish (relative workload pressure)

Table A-7 - FactWeatherDaily schema

Field	Data type	Role in model	Notes from TMDL
sk_fact_weather	int64 (hidden)	Surrogate key (PK)	Hidden; internal key
bk_fact_weather	string	Business identifier (degenerate dimension)	Traceability
sk_date	int64 (hidden, non-null)	FK to dim_date[sk_date]	Relationship defined (referential integrity)
sk_location	int64 (hidden, non-null)	FK to dim_location[sk_location]	Relationship defined (referential integrity)
precipitation_sum	int64	Daily precipitation	Used by measure Precipitation (SUM)
temperature_2m_max	double	Daily max temperature	Optional contextual variable
temperature_2m_min	double	Daily min temperature	Optional contextual variable
wind_speed_10m_max	double	Daily max wind speed	Optional contextual variable
Is_Rainy_Day	int64	Rain indicator	1/0 flag (treated as numeric in TMDL)

Table A-8 - FactWeatherDaily Measures

Measure name	Folder	Purpose (summary)
Precipitation	Weather\Core	Total precipitation in context
Incidents (Rainy)	Weather\Rain vs Dry	Incidents restricted to parish days flagged rainy (TREATAS pairs)
Incidents (Dry)	Weather\Rain vs Dry	Incidents restricted to parish days flagged dry (TREATAS pairs)
Parish days (Rainy)	Weather\Rain vs Dry	Count of rainy parish day pairs in context
Parish days (Dry)	Weather\Rain vs Dry	Count of dry parish day pairs in context
Incidents per Parish day (Rainy)	Weather\Rain vs Dry	Incidents (Rainy) / Parish days (Rainy)
Incidents per Parish day (Dry)	Weather\Rain vs Dry	Incidents (Dry) / Parish days (Dry)
Incidents (Total Rainy+Dry)	Weather\Rain vs Dry	Sum of rainy + dry incidents
Rainy Share %	Weather\Rain vs Dry	Share of incidents that occur on rainy parish days
Rainy Lift	Weather\Sensitivity	(Incidents per rainy parish day) / (Incidents per dry parish day)
Rainy Sensitivity Category	Weather\Sensitivity	Categorical interpretation of Rainy Lift (High/Neutral/Lower/Very low)

Table A-9 - Bus matrix (processes × dimensions)

Business process / fact table	DimDate (dim_date)	DimLocation (dim_location)	DimIncidentType (dim_incident_type)	Notes
Incident recording and monitoring (fact_incident_daily)	✓	✓	✓	Core process; daily grain via sk_date; drill-down by Area/Type
Weather context at parish day (fact_weather_daily)	✓	✓	—	Context process; not incident-typed; joined by same conformed date/location
Cross-process monitoring (combined analysis)	✓	✓	✓ (via incident fact)	Enabled through shared conformed dimensions and shared filter context

Figure A-1 - Homepage: Overview and navigation

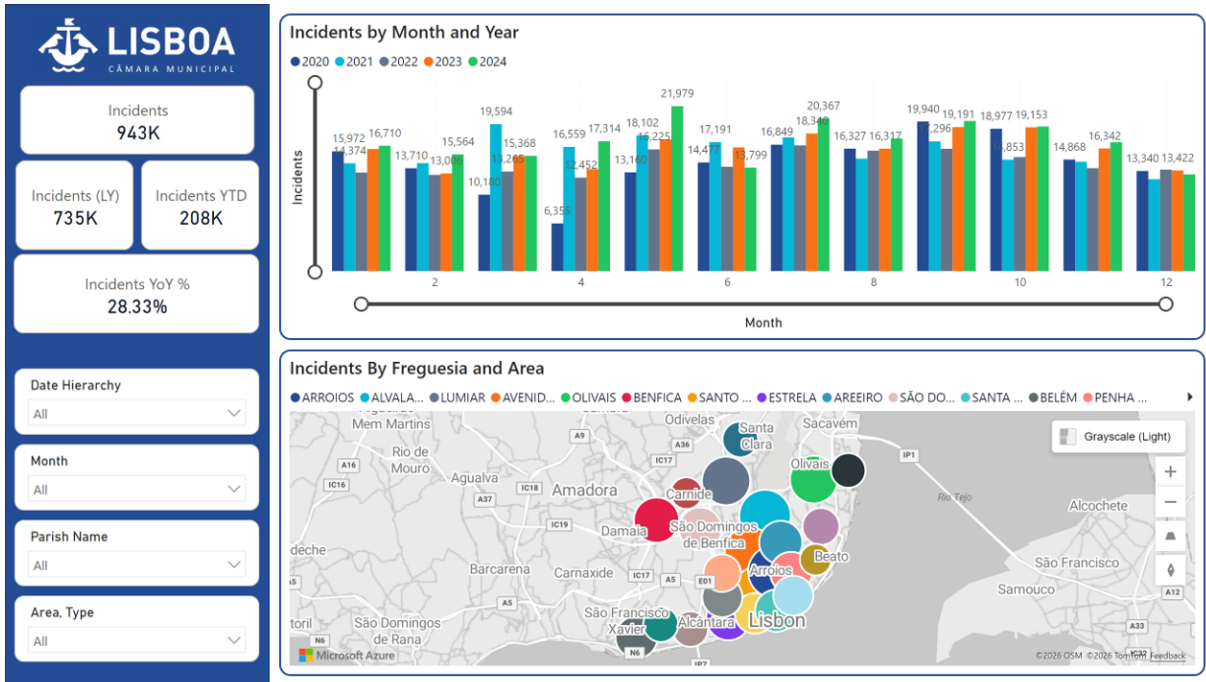


Figure A-2 - BN1 page: Workload and spatial distribution




Figure A-3 - BN2 page: Seasonality and baseline monitoring

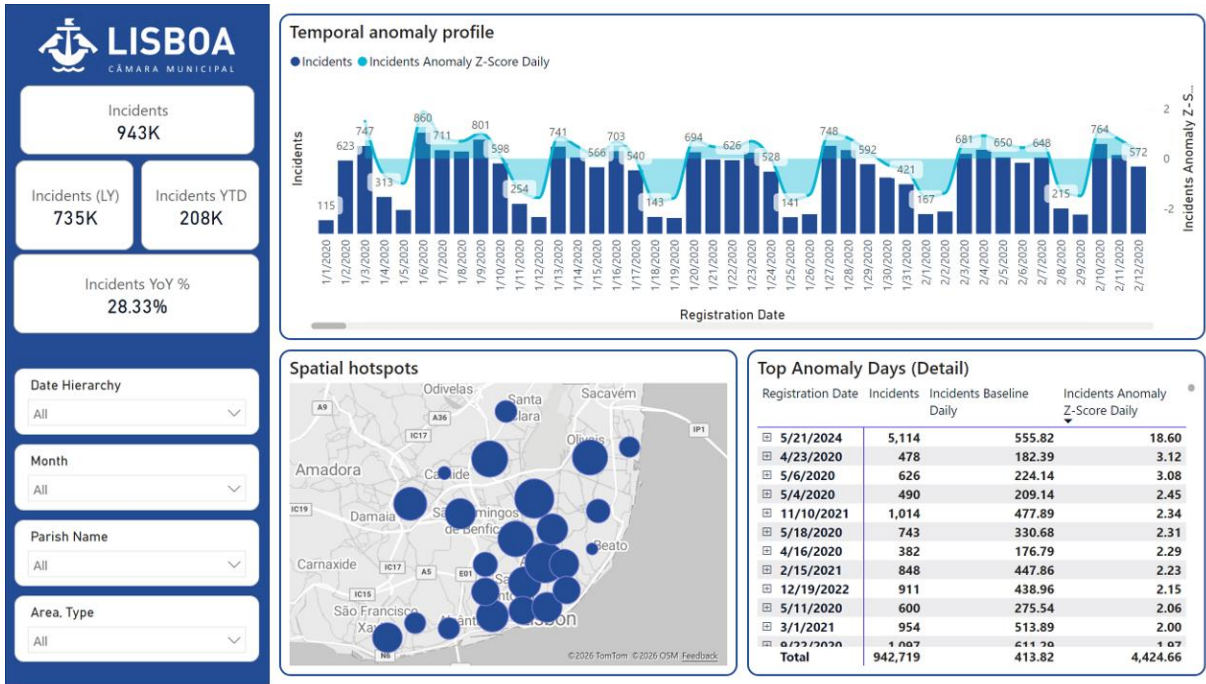


Figure A-4 - BN3 page: Anomalies and hotspots



Figure A-5 - BN4 page: Territory prioritisation (pressure and ranking)

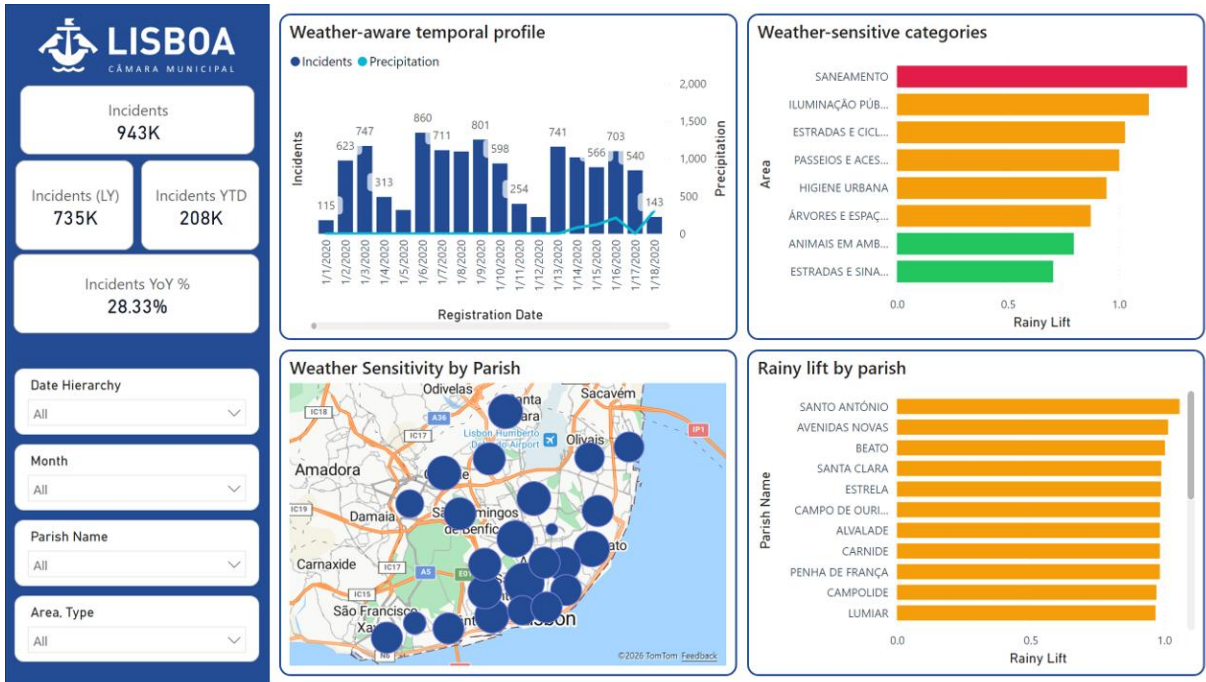


Figure A-6 - BN5 page: Weather context (rain indicators and comparative rates)

Data with Purpose.

