

Road Freight OD Inference

Methodology and pilot study using Tasmania data
Final Report

Authors

Wei Sun
Research Fellow
Adelaide University
wei.sun@adelaide.edu.au

Akshay Vij
Professor
Adelaide University
akshay.vij@adelaide.edu.au

Nic Kaliszewski
Research Adjunct
Adelaide University
nicholas.kaliszewski@adelaide.edu.au

Acknowledgements

This study was funded by the Australian Government Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts (DITRDCA) and the iMOVE Cooperative Research Centre (CRC). The project was managed by David Mitchell and Joe O'Sullivan, Bureau of Infrastructure and Transport Research Economics (BITRE).

EXECUTIVE SUMMARY

The challenge: Planning in a data deficit

Road freight is the backbone of the Australian economy, yet the evidence base for planning its future is increasingly fragmented. The last comprehensive national snapshot — the ABS Road Freight Movements Survey (RFMS) — was conducted in 2014. In the decade since, the Australian freight task has evolved significantly, driven by population growth, supply chain reorganization, and the rise of e-commerce. Transport authorities currently face a critical dilemma: they must make multi-billion-dollar infrastructure decisions based on either outdated structural data (2014) or piecemeal observational data such as traffic counts that lack the origin-destination (OD) logic required for strategic planning.

The solution: An integrated inference framework

This project delivers a robust, data-driven solution to this "data deficit." We have developed and piloted an integrated inference framework that systematically fuses the structural legacy of the 2014 survey with the continuous reality of modern traffic monitoring.

The proposed framework adopts a transparent, interpretable architecture containing:

- **Structural Learning:** It extracts the fundamental geo-economic logic (the relationship between land use, employment, freight generation and distribution) from the 2014 survey.
- **Baseline Projection:** It projects this logic forward using current socio-demographic data to create a "Baseline origin-destination matrix (ODM)."
- **Empirical Calibration:** It uses a rigorous optimization process to adjust this baseline against observed traffic counts, generating a "Calibrated ODM" that respects both economic fundamentals and actual road network usage.

Key findings from the Tasmania Pilot

Using Tasmania as a pilot region, the framework demonstrated high efficacy in generating validated 2024 freight matrices:

- **Validity:** The 2024 Calibrated ODM achieved near-perfect symmetry between origin and destination flows, a key marker of internal validity.
- **Consistency:** For the vast majority of regions, the estimated freight generation remained statistically consistent with historical error bands, confirming the stability of the model.
- **Structural decay:** Crucially, longitudinal analysis (2015–2024) revealed a measurable, clear "drift" between the economic baseline and observed reality. This confirms that the 2014 structural knowledge is slowly losing its predictive power, with the deviation growing at a consistent annual rate.

Strategic value for transport agencies

The framework provides the Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts (DITRDCA), and state and territory transport agencies and planning authorities, with the following immediate strategic capabilities:

- **Continuous estimation:** The ability to generate "best-estimate" origin-destination (OD) matrices for any year where traffic count data exists, removing the blind spots between major survey cycles.
- **Long-term projection:** Because the framework is anchored in structural economic drivers (employment and land use), rather than just historical trends, it serves as a robust forecasting engine. It enables transport agencies to model future freight scenarios (e.g., 2030, 2040) based on population growth and land-use changes, independent of current traffic constraints.
- **Evidence-based survey planning:** By quantifying the structural decay rate, transport agencies can move away from arbitrary survey schedules. The framework can also provide an early-warning system, identifying exactly when the 2014 baseline estimates have become too unreliable, thereby providing a robust business case for commissioning the next generation of freight surveys.

Conclusion and next steps

This study proves that it is possible to bridge the gap between historical surveys and modern data streams. The framework is modular, scalable, and designed to ingest emerging data sources (such as GPS telematics and Bluetooth) to further improve accuracy. We recommend scaling this methodology to mainland Australia to establish a continuous, national-scale freight monitoring capability.

TABLE OF CONTENTS

1. Introduction	1
2. Data exploration	2
2.1. ABS Road Freight Movement Survey, 2014	3
2.2. Tasmania traffic counter and counts	4
2.3. Transport Network Strategic Investment Tool by CSIRO	5
2.4. Tasmania Freight Survey, FY 2021-22	6
3. Literature review	6
3.1. Overview	7
3.2. Structural Projection Methods	7
3.3. Observational Inference Methods	8
3.4. Integrated Inference Frameworks	10
3.5. Current Industry Practices	11
3.6. Findings from the literature review	12
4. Methodology	13
4.1. Framework overview	14
4.2. Structural Learning Module	16
4.3. Baseline ODM Projection Module	20
4.4. Route Choice Modelling Module	22
4.5. Location Matching Module	24
4.6. Calibrated ODM Estimation Module	24
4.7. Value delivery	26
5. Tasmania case study	26
5.1. Input data	27
5.1.1. Census data	27
5.1.2. Geography	29
5.1.3. Major road network	32
5.1.4. Traffic counters	34
5.2. Estimated structural model	34
5.2.1. Land use model estimation results	34
5.2.2. Gravity model estimation results	37
5.3. Route choice modelling	38
5.4. Structural model effectiveness check	41
5.5. ODM inference for 2024	44
5.6. Freight movement survey power monitoring	47
6. Discussions	50
6.1. Bridging the data deficit with hybrid inference	51
6.2. Monitoring structural decay	51
6.3. The Role of Hierarchical Regularization	51

6.4.	<i>Limitations and future directions</i>	52
7.	Conclusions	52
7.1.	<i>Methodological contribution</i>	53
7.2.	<i>Key findings from the Tasmania pilot</i>	53
7.3.	<i>Strategic implications for BITRE</i>	53
7.4.	<i>Future outlook</i>	54
	Reference	55

1. INTRODUCTION

Road freight is the lifeblood of the Australian economy. In 2021-22 alone, road transport moved an estimated 234.6 billion tonne-kilometres of freight, representing nearly 30% of the total domestic freight task (BITRE, 2022). With volumes projected to grow by 77% between 2020 and 2050, the need for robust, evidence-based planning tools has never been greater. Yet, despite the magnitude of this challenge, Australia's freight modelling capabilities face a critical bottleneck: a "data deficit" characterized by the disconnect between outdated structural intelligence and fragmented modern observation (McHugh et al., 2021).

Effective transport planning requires two distinct types of intelligence, the structural knowledge and the observational reality. The structural knowledge helps understand why freight moves, which include the economic relationships between industries, land use, and supply chains that determine origin-destination (OD) demand. On the other hand, the observational reality helps understand where freight is moving now, the actual traffic flows on the network.

In Australia, these two information streams are temporally decoupled. The last comprehensive source of structural knowledge, the ABS Road Freight Movements Survey (RFMS), was discontinued after 2014. Conversely, observational data (traffic counts) is abundant and current but lacks the OD logic required for strategic planning. This leaves transport agencies and planning authorities with a forced choice: rely on a decade-old structural snapshot that misses recent shifts in transport dynamics, or rely on traffic counts that provide sparse traffic volume without context.

This project addresses this impasse by developing an integrated ODM inference framework. Rather than viewing the 2014 survey as "obsolete," we treat it as a foundational structural prior, representing a repository of economic logic that can be extracted, projected forward, and calibrated against observational traffic data. Building on the single-level optimization principles established in our previous work (Sun et al., 2023), this framework fuses three critical components including structural learning, baseline projection and empirical calibration.

The primary aim of this report is to demonstrate that valid, actionable freight intelligence can be generated in a data-sparse environment without immediate recourse to expensive new surveys. Specifically, this project delivers on three objectives:

- **Data Robustness Analysis:** Systematically evaluating the utility of publicly available datasets (ABS RFMS, NFDH traffic counts) and complementary sources (CSIRO TraNSIT, State Government surveys) for network-wide OD inference.
- **Framework development and pilot case study:** Developing and testing a scalable inference framework using Tasmania as a closed-system pilot region to validate the methodology.
- **Strategic monitoring capability:** Designing a quantitative method to monitor the "shelf life" of survey data, providing transport agencies with an evidence-based trigger for determining when new data collection cycles are required.

The outcomes of this project provide DITRDCSA, and other transport agencies, with a comprehensive, three-pillar strategic toolset which enables operational estimation, long-term forecasting and survey cycle optimization.

Following sections of this report is structured in chapters below. In Chapter 2, we explore the datasets mentioned above. In Chapter 3, we conduct a comprehensive literature review on

the ODM inference, with a focus on the road freight. In Chapter 4, we describe the proposed methodology for the road freight ODM inference. In Chapter 5, we describe the pilot case study using Tasmania data and summarize the findings. In Chapter 6, we discuss the limitations and future steps. Finally, Chapter 7 concludes this report.

2. DATA EXPLORATION


In this section, we explore some of the most relevant datasets to understand road freight movements in the Australian context. This exercise helps us to form the high-level understanding of the inter- and intra-zonal freight flow, and to pick Tasmania as the pilot region to develop and test the road freight ODM inference methodology.

2.1. ABS Road Freight Movement Survey, 2014

As introduced on ABS's webpage¹, the most recent road freight movement survey conducted by the ABS at the national scale lasted for 12 months and ended at the end of October 2014, after that, this surveyed has been discontinued. While the survey is over 10 years old, it remains one of the most relevant data sets to understand the total volume and geographic distribution of road freight movements in Australia. With a specific focus on the freight trip numbers, we have explored the surveyed ODM and below is the summary of some key insights.

The 2014 survey road freight ODM estimated a total of 415 million road freight trips nationwide. The state-level distribution of trip production and attraction are shown in **Figure 1**. The chart shows that New South Wales produced and attracted the most road freight trips, followed by Victoria and Queensland. Also, it is clear to see that at the state level, zonal trip productions are similar to zonal trip attractions.

Figure 1- State-level road freight trips in 2014 Road Freight Movements Survey



In terms of road freight trip distribution, **Figure 2** shows that at the state level, most trips are intra-zonal trips. Numerically, about 98.3 road freight trips are intra-zonal at the state level (i.e. start and end within the same state). The proportion of intra-zonal trips is still dominant at 65% if we zoom in to the SA4 level. The distribution chart also shows that there are no inter-state 'road freight' trips directly between Tasmania and other states or territories – freight movements between Tasmania and the mainland must travel by sea or air. This makes Tasmania a smaller and independent transport system, compared to the interconnected system in the mainland Australia, and a smaller, more tractable geographical scope on which to test the ODM inference methodology and workflow before scaling it up to cover the entire

¹ <https://www.abs.gov.au/ausstats/abs@.nsf/mf/9223.0>

country. Therefore, this study used Tasmania as the pilot region for developing and testing the freight ODM inference methodology.

Figure 2 - State level road freight trips distribution



2.2. Tasmanian traffic counters and counts


As at the end of August 2025, there were 43 permanent traffic counters (noting that one traffic counter could cover multiple lanes) on the Tasmanian mainland which could differentiate heavy vehicles (likely road freight vehicles) from small vehicles such as private cars, and provide the proportion of heavy vehicles (% HV) as the percentage of the Average Annual Daily Traffic (AADT)². Therefore, we resorted to using traffic count data reported on the Tasmanian Department of State Growth's traffic count website³. The traffic counter locations are marked as circles in Figure 3. The counter locations represent a good coverage, and in 2024, the average annual daily traffic counts for heavy vehicles ranged from 79 to 6689 vehicles per day. Among the 43 available traffic counters, 26 of them were available for 10 years from 2015 to 2024, and these were valuable for understanding how the knowledge obtained from the 2014 survey decays over time. The methodology section will cover this point in detail.

One particular point to mention is that the traffic counts are used at the traffic counter level, instead of the link level. This is because while we know that each traffic counter covers at least 2 lanes, representing two directions, we do not have good enough meta data to assign traffic counts to individual links in the road network graph.

² Please note that initial attempts to use traffic count data reported on the National Freight Data Hub (NFDH) API service (see: https://aueprod01ckanstg.blob.core.windows.net/public-catalogue/9ee49e54-f2dd-4dfc-b5a4-28182d658e9f/Harmonised%20Traffic%20Counts%20dataset%20API%20User%20Guide_July2023.pdf), provided less than ideal data resolution for the ODM methods.

³ <https://tasmaniatrafficdata.drakewell.com/publicmultinodemap.asp>

Figure 3 - Permanent traffic counter locations in Tasmania, 2025



2.3. CSIRO Transport Network Strategic Investment Tool

Another potential data source for understanding Australian road freight movements is the Transport Network Strategic Investment Tool (TraNSIT), developed by the Commonwealth Scientific and Industrial Research Organisation (CSIRO). As described on its webpage⁴, TraNSIT is a computer model that creates detailed maps of Australia's supply chains, freight movements and costings across road, rail, ship and air freight networks, and it has been used since 2012 to provide such information to improve supply chain efficiency and help identify options for freight transport and logistics investment.

Ideally, we would like to incorporate such data into the ODM inference process, as we have demonstrated such methods in our previous work (Sun et al., 2023), or use it for externally validating our inferred ODM. With the help from the BITRE team, we obtained an extract of TraNSIT freight movement data for Tasmania from January to August 2025. The most relevant part of the movement data is the total estimated freight travel distances between SA3 level origin-destination pairs in Tasmania. Using the estimated distance matrix that we have estimated in Tasmania at the SA3 level (we will describe this in detail in the methodology section), we have converted the obtained Tasmania freight total travel distance matrix into the approximated freight trips ODM. The comparison between the derived TraNSIT ODM and the

⁴ <https://www.csiro.au/en/research/technology-space/it/transport-logistics-transit>

Tasmania road freight ODM from 2014 survey reveals stark differences. In terms of the total number of road freight trips, while the surveyed ODM suggests 7.55 million trips in 2014, the derived TraNSIT ODM only captures about 0.6 million trips in the first 8 months in 2025, accounts for about 8% of the surveyed trips. Structurally, the cosine distance between two ODMs (trips, zonal trip production, and zonal trip attraction) fall into 0.4 ~ 0.5 range, indicating structural differences.

The differences are significant yet understandable, as TraNSIT mainly adopts a bottom-up approach and collects data by engaging with a selected set of stakeholders, makes it different from the “complete picture” approach in the ABS survey.

2.4. Tasmania Freight Survey, FY 2021-22

The Tasmanian Freight Survey (TFS), a regular survey undertaken by the Tasmanian Department of State Growth (DSG), is another potential highly relevant data source⁵. Previous surveys have been collected every 3 years from 2002-03, and the latest round was collected in 2021-22. With the help from the BITRE team, we also obtained the 2021-22 TFS data. The data was provided at the Suburbs and Locality (SAL) level, and we used the 2021 SAL and SA3 digital boundaries⁶ to remap and convert it into the SA3 level ODM for direct comparison to the 2014 surveyed ODM by ABS.

The findings from such a comparison are similar to those from the TraNSIT derived ODM. The TFS only suggests 0.81 million road freight trips in FY 2021-22, accounting for about 11% of trips captured in the 2014 survey, and the structural differences are about 0.5, mainly because the TFS is also based on the bottom-up approach that collects freight data from a selected set of stakeholders. Moreover, the DSG noted that the TFS collection does not obtain data from every freight operator in the state and the raw survey estimates are not further weighted to derive a state-wide freight volume.⁷

In conclusion, the data exploration suggests that the 2014 road freight movements survey remains the gold standard for understanding the volume and structure of freight flows, and the traffic counts, as they reflect the reality of freight vehicle movements, could be used for calibrating the prior knowledge of the freight ODM. On the other hand, the features of TraNSIT and TFS ODMs make them less directly comparable to the full-scope road freight ODM that we are aiming at (as we will follow the scope of the 2014 surveyed ODM). In fact, this finding further emphasizes the need for a full-scope road freight ODM inference framework.

⁵ https://www.transport.tas.gov.au/transport-policy-and-planning/transport-policy-and-planning-small-tiles/management/tasmanian_freight_survey

⁶ <https://www.abs.gov.au/statistics/standards/australian-statistical-geography-standard-asgs-edition-3/jul2021-jun2026/access-and-downloads/digital-boundary-files>

⁷

https://www.stategrowth.tas.gov.au/__data/assets/pdf_file/0004/139504/Tasmanian_Freight_Survey_Report_FINAL_2014_25.pdf

3. LITERATURE REVIEW

3.1. Overview

This section provides a comprehensive literature review on methods for estimating or inferring road freight Origin-Destination Matrices (ODM). It surveys the established academic research papers and grey reports published by major transport planning authorities.

The estimation of Origin-Destination (OD) matrices remains a foundational yet complex challenge in transport engineering, serving as the critical input for traffic assignment, infrastructure planning, and economic forecasting. While methodologies for passenger travel are well-established, truck freight introduces unique stochastic behaviours, such as tour-based trip chaining and heterogeneous vehicle classes, that defy simple extrapolation from standard commuter models.

This challenge is particularly acute in the Australian context, where the freight task is characterized by vast inter-urban distances, a heavy reliance on road transport, and a distinct "data deficit." The national landscape currently lacks a continuous, disaggregate observation of freight movements, relying instead on the discontinued ABS Road Freight Movements Survey in 2014 as the last comprehensive structural snapshot. Consequently, the problem of estimating current flows is not merely one of observation, but of inference: how to project outdated behavioural priors onto modern, aggregate traffic counts.

This section synthesizes the academic literature to inform a robust estimation framework suitable for this data-constrained environment. It is structured to trace the methodological evolution through three distinct paradigms:

- **Structural Projection Methods:** Classical models such as gravity and land-use models, and Input-Output-Based models that infer flow from economic activity.
- **Observational Inference Methods:** Statistical techniques that update matrices using observed traffic counts and emerging telematics.
- **Integrated Inference Frameworks:** Contemporary methods, such as the works based on advanced machine learning and deep learning techniques, Bayesian approach and the single-level optimization approach proposed in our previous work (Sun et al., 2023), which integrate heterogeneous data sources to overcome the limitations and potential biases of any single data source.

3.2. Structural projection methods

Before the widespread availability of digital tracking data, freight OD estimation relied on "synthetic" methods that inferred flows from the spatial distribution of economic activity. These model-driven approaches exploit the strong correlation between land-use attributes (e.g., industrial GVA, warehouse floor space) and freight generation to construct a "prior" matrix—a structural skeleton of demand that remains essential for modern Bayesian updating frameworks.

The dominant model in this category is the gravity model. The gravity model serves as the seminal framework for trip distribution. It postulates that freight flow between two zones is proportional to their relative economic "masses" (productions and attractions) and inversely proportional to the impedance (cost or time) separating them. While originally reliant on extensive interview surveys, Tamin and Willumsen (1989) revolutionized the application of gravity models by demonstrating that model parameters could be calibrated using observed

traffic counts alone. This bridged the gap between theoretical demand synthesis and the data constraints of real-world logistics.

A critical component of the gravity model is the deterrence function, which quantifies the "friction" of distance. The literature proposes several functional forms:

- Exponential Function: Commonly used in urban passenger modelling (Entropy Maximization) where trip costs have a characteristic scale.
- Power Function: Theoretical literature often identifies the Power function as particularly robust for inter-regional and long-haul analyses due to its scale-invariant properties, making it highly relevant for the vast spatial distances of the Australian freight network.
- Gamma (Tanner) Function: Suggested by Duanmu (2012) for capturing complex non-linear sensitivities in specific commodity markets.

The primary critique of gravity formulations, noted by Holguín-Veras and Thorson (2000), is the independence assumption. Gravity models treat trips as isolated events, failing to capture the "trip chaining" or tour-based behaviour characteristic of urban delivery rounds. However, for strategic, long-haul interstate freight—the dominant task in the Australian context—the direct O-D approximation remains structurally sound.

Another type of model is the Input-Output-Based models. These models approach OD estimation from a macroeconomic perspective, as they link the physical flow of goods to the monetary transactions flow between industrial sectors. The concept could be traced back to studies published decades ago (Barna, 1963; Leontief, 1936). The method is typically in three stages including presenting the economic structure with an Input-Output table, generating trade flows given the demand change, and finally performing the value-to-volume conversion to project the physical freight movements (Li et al., 2023; Pompigna and Mauro, 2020).

3.3. Observational inference methods

As traffic monitoring technology has evolved, the literature has shifted toward methods that "update" or "calibrate" prior matrices using observed network data. These approaches aim to solve the classic inverse problem: identifying the OD matrix that best reproduces the observed traffic patterns, while retaining the structural logic of the prior information. It is worth noting that by incorporating the observed traffic data, the task shifts from "synthesizing" demand (theory) to "inferring" the underlying road freight movements from empirical signals. With more data sources becoming available, a rich literature has emerged to demonstrate that specific kind of data is useful for calibrating the prior knowledge of the underlying ODM (either from survey or from the structural model projection).

The first kind of observable data which records the vehicle movement is the traffic counts from pneumatic tubes or induction loops; these represent the most ubiquitous data source in freight modelling. Since counts provide only link volumes (or just the site volumes) with no trip origins or destinations attached, inferring an OD matrix from them is an under-determined problem requiring statistical regularization. Two dominant frameworks have been developed to solve this under-determined problem: Entropy Maximization (EM) and Generalized Least Squares (GLS).

The EM approach was firstly formalized by Van Zuylen and Willumsen (1980), EM identifies the most probable matrix consistent with observed counts by maximizing the number of micro-states (individual vehicle arrangements) that satisfy the macro-constraints. While classically applied to static problems, recent research has revitalized EM for modern, data-constrained freight contexts. The GLS approach was established by Bell (1991) and Cascetta (1984). Such methods minimize the weighted squared deviations between the estimated and prior flows. Its

strength lies in its ability to explicitly weight the reliability of different data inputs, allowing a modeler to trust "hard" traffic counts over a "soft" outdated survey prior. While EM and GLS methods were established decades ago, they remain the fundamental methods for ODM estimation.

Another development is on the route choice, that most statistical inference assumes a static assignment map, the Bilevel Optimization acknowledges that freight route choice depends on congestion. These models frame estimation as a Stackelberg game: an "upper level" adjusts the OD matrix to match counts, and a "lower level" re-assigns traffic to User Equilibrium (UE). While theoretically robust (Yang, 1995; Yang et al., 1992), the computational intensity of bilevel programs often makes them less practical for large-scale continental networks like Australia's. Also, studies (Fisk, 1989) suggest that the actual vehicle route choice may not follow the "optimal" choices from an equilibrium traffic state. Therefore, a stream of later studies adopted the single-level framework to minimize the dependency on the traffic assignment process and such structure make the inference more efficient for large networks (Behara et al., 2022; Dey et al., 2020; Ma and Qian, 2018; Shen and Wynter, 2012).

Apart from the most ubiquitous traffic counts data, the rise of Intelligent Transport Systems (ITS) has introduced "direct observation" methods, which capture the complete or partial trajectories of vehicle movements. A lot of studies have tried to utilize such data sources into vehicle ODM inference, and some of them are specifically targeted at the freight movements. The summary below covers some most relevant recent studies.

- **GPS and Telematics data** that extracted from fleet management systems, provide high-fidelity "breadcrumbs" of truck movements. Recent studies have demonstrated the use cases for such data. Demissie and Kattan (2022) utilize a heuristic-based algorithm to extract significant truck stops and individual trip segments from raw GPS trajectories that lack explicit activity descriptors, and truck OD flows are inferred using a multinomial logit (MNL) destination choice framework that fuses the disaggregated GPS-derived behaviour with aggregate constraints from existing traditional, survey-based travel demand models. Basso et al. (2022) extract route trajectories from a subset of GPS-tracked trucks to calculate empirical transition probabilities between specific highway entry and exit gates, and use such data alongside Automatic Vehicle Identification (AVI) data and a set of scaling factors to estimate the truck ODM in Chile. Other studies (Yang et al., 2021; Zanjani et al., 2015) also demonstrate similar methods to incorporate GPS (or telematic) data into freight ODM inference.
- **Bluetooth data** collected by Bluetooth detectors. A Brisbane-based study (Michau et al., 2019) presents a methodology for estimating Link-OD matrices by fusing Bluetooth data with traditional traffic counts. While the case study covers general traffic in Brisbane, the methodology of using Bluetooth MAC re-identification to correct OD flows is fundamental to freight applications where Bluetooth samples are sparse but highly accurate for specific paths. Other studies (Blogg et al., 2010; Cipriani et al., 2021) also demonstrated similar approaches, and also suggested that the data may need to be filtered or re-identified to be applied to road freight vehicles.
- **Phone (Cellular) data** has also been used in studies to infer the underlying vehicle or people movements (Bonnell et al., 2018, 2015; De Rolland and Bayart, 2025; Graells-Garrido et al., 2023; Wang and Chen, 2018). However, due to the data availability and precision limitation, differentiating or filtering freight movements from the general movements remains challenging and requires accompanying algorithms or richer meta data.

- **Weigh-in-Motion (WIM) data** is unique because it provides vehicle weight and axle configuration, allowing for precise truck classification. The challenge is "re-identification", which is matching a specific truck's weight signature at Location A to its signature at Location B to form a trip. Existing studies (Basar et al., 2018; Cetin et al., 2011; Gajda et al., 2020) explore the uses of WIM data in understanding truck movements, and suggest the incorporation and correct processing of WIM data could be highly valuable for more accurate freight ODM inference.

3.4. Integrated inference frameworks

With the increasing richness of traffic data, and the fact that most data from a single data source suffers from low penetration level or structural bias, or both, studies in recent years tend to apply integrated inference frameworks to explicitly incorporate different observable traffic data into the ODM inference process to achieve better results. It is worth noting that the classical statistical calibration methods such as the EM and GLS methods described previously also fall into this broad category, as they were developed to fuse the traffic counts with the established prior ODM knowledge for a more realistic estimation on the underlying true ODM.

One stream of integrated inference frameworks is Bayesian-based. The Bayesian approach fuses a "prior belief" (derived from structural gravity models or historical surveys) with "new evidence" (observed traffic counts or sparse trajectories) to generate a "posterior" distribution (Maher, 1983). Later studies further advanced the approach to suit different input data (instead of the typical traffic count when the method was firstly developed) and improve the inference performance (Carvalho, 2011; Castillo et al., 2008; Yao et al., 2021).

Another critical computational breakthrough was the Gradient Approach proposed by Spiess (1990). Spiess demonstrated that the sensitivity of link flows to OD changes could be calculated efficiently using a 'backward pass' on the network, enabling the iterative adjustment of the matrix via steepest descent. This method established the mathematical foundation for scalable OD estimation, serving as the precursor to modern single-level optimization frameworks which extend this gradient-based logic to fuse heterogeneous data sources. Our previous work (Sun et al., 2023) falls into this stream of the integrated inference framework, as it follows the optimization structure introduced by Spiess (1990), and replaces the heavy bi-level equilibrium with the single-level structure (Behara et al., 2022; Dey et al., 2020) using observed route choices, and it is designed to be flexible to explicitly incorporate multiple data sources seamlessly into the optimization.

Another recent development on such integrated inference frameworks is applying advanced machine learning and neural network approaches to the ODM inference questions. As the volume and heterogeneity of data explode, traditional optimization methods, which rely on explicit mathematical formulations of the traffic flow physics, are being supplemented or replaced by Deep Learning (DL) architectures. These models "learn" the complex, non-linear inverse function mapping traffic counts and sparse trajectories to OD matrices. A range of studies emerged to apply a variety of advanced machine learning and (or) neural network-based approaches to help improve the vehicle ODM inference performance. To name a few, a combination of Federated Learning (FL) and Deep Neural Networks (DNN) (Abbas et al., 2024) has been applied to predict the truck destination regions while preserving data privacy by applying a decentralized training architecture. XG-Boost algorithm has been trained using the US truck flow data to explore the effects of population and employment on the heavy freight flows (Uddin et al., 2023). A deep-gravity framework has been applied based on the rich transport datasets in China (Yang et al., 2023), and it reaches better performance than the conventional gravity model in generating heavy truck mobility flows. By acknowledging that the generic deep learning models such as Convolutional Neural Network (CNN) and Recurrent

Neural Network (RNN) often fails to capture the non-Euclidean topology of road networks, graph-based models such as Spatial GNNs have become the state-of-the-art for modelling logistics and freight demand because they explicitly model the "links" and "nodes" of the freight network (Peng et al., 2024).

As a relatively new research stream, there are a lot of vehicle flow-related studies not mentioned in this literature review section, but advanced machine learning and deep learning have shown big potential in improving the ODM projection and inference quality. It is worth noting that while advanced machine learning or deep learning-based methods could achieve very high performance, they are often criticized as data hungry black boxes (noting that some recent studies using Physics-Informed Machine Learning for traffic state estimation, making it less Blackbox (Di et al., 2023)), and such features limit their applications in data deficit situation and adoption by authorities.

3.5. Current industry practices

We also conducted a review on grey literature to understand the current practices adopted by Australian and major international transport planning authorities.

In the Australian context, the Australian Transport Assessment and Planning (ATAP) Guidelines and the Bureau of Infrastructure and Transport Research Economics (BITRE) research reports provide national analysis and instructions on the freight demand.

The ATAP Guidelines provide the primary national framework for urban freight demand modelling. The ATAP T9 Urban Freight Demand Modelling guideline outlines that most Australian jurisdictions use a combination of trip-based models and commodity flow models (*T9 Urban freight demand modelling*, 2021). Key practices include:

- **Four-step modelling approach:** Trip generation, trip distribution, mode choice, and network assignment using gravity models and growth factor methods.
- **Forecasting methodology:** Based on growth factoring using ratios of future year and base year trip ends, incorporating trip rates and future land use data
- **Data sources:** Heavy commercial vehicle (HCV) forecasting relies on GPS traces, establishment surveys, and population/employment forecasts as primary drivers.

BITRE publishes the Australian Aggregate Freight Forecasts, which provides 30-year projections of freight volumes by mode (road, rail, sea, air) between 2020 and 2050. The methodology includes:

- **Linear programming approach:** Models commodity movements incorporating supply chain characteristics, intermodal terminals, and collection/storage facilities (Mitchell and Kurniawan, 2015).
- **Commodity-specific modelling:** Separate modules developed for iron ore, coal, grains, cotton, rice, and sugar with accuracy within 2-3% of industry-reported volumes (Mitchell and Kurniawan, 2015).
- **National-scale projections:** Road freight projected to grow by 77% between 2020 and 2050, with urban freight expected to increase from 47.3 billion tkm (2018) to 76.7 billion tkm by 2040 (*Australian aggregate freight forecasts – 2022 update*, 2022).

Also, we are aware of the TraNSIT model developed by CSIRO to forecast the freight travel demand for certain commodity types. As described in previous sections, its bottom-up nature makes it less compatible with the aim of developing a framework to infer the full scope of road freight movements in Australia.

In terms of the international transport planning authorities, this literature review surveyed the practices adopted by authorities in UK and USA.

In UK, the Transport Analysis Guidance (TAG)⁸ provides comprehensive modelling and forecasting guidance. Key documents include:

- TAG Unit M4: Forecasting and Uncertainty - provides practical guidance for forecasting transport project impacts through mathematical models with different assumption sets (*TAG Unit M4 - Forecasting and Uncertainty*, 2025)
- TAG Unit M1.1: Principles of Modelling and Forecasting - establishes standard structure for modelling transport demand and supply based on economic principles (*TAG UNIT M1.1 - Principles of Modelling and Forecasting*, 2024)
- Multi-modal approach: Includes without-scheme and with-scheme forecasts for multiple forecast years to account for population changes and infrastructure development (*TAG UNIT M1.1 - Principles of Modelling and Forecasting*, 2024)

In USA, the reports published by the Federal Highway Administration (FHWA), especially the Freight Demand Modelling and Data Improvement program⁹ serve as the main source of instructions. Their approach is based on the Behavioural-based national freight demand model. Such model is led by development of components including joint mode choice and shipment size models, firm and establishment synthesis models, and new carrier choice models. The Freight software, is an Open-source implementation available to all states and regions (*Freight Demand Modelling and Data Improvement Handbook*, 2017).

3.6. Findings from the literature review

The literature review provides rich information on the methods that developed in the academic literature to project and infer the road freight movements, and the industrial practices covered by grey literature. Such insights help us to perceive the different focus of academic exploration and industrial practices, as well as to identify the rough shape of the framework that best fit into the Australian context.

Firstly, the grey literature review suggests that the current industrial practice of major transport planning authorities primarily focuses on structural forecasting (projecting movements based on economic activity and land use) rather than statistical inference (deriving movements directly from observed traffic data). While authorities do use real traffic data (e.g., loop counts, screen lines), they typically treat it as a calibration and validation target rather than the primary input for inferring the Origin-Destination (O-D) matrix. This insight suggests that the transport planning authorities in general have not fully exploited the value of the rich traffic data, and utilise it to better understand the underlying freight movement.

Secondly, in the Australian context, the data deficit reality (as we have discussed in the previous sections that the only full-scope knowledge source is the 2014 ABS survey, and the traffic counts serve as the only truth at this stage), limits the methodology choices by the

⁸ <https://www.gov.uk/guidance/transport-analysis-guidance-tag>

⁹ https://ops.fhwa.dot.gov/freight/freight_analysis/fdmdi/index.htm

authorities to develop an integrated freight ODM inference framework. From BITRE's perspective, we would suggest the developed full-scope (as in the 2014 ABS survey) inference framework to have several desired features:

- Methodologically solidity. The framework should be based on established work, makes it more defensible.
- Adaptivity and improvability. The framework should be able to achieve reasonably good performance given the current data deficit reality. Also, it could be able to improve its performance when more and better data sources become available, as indicated by the National Freight and Supply Chain Strategy (*National Freight and Supply Chain Strategy 2025, 2025*).
- Efficiency and scalability. The framework should be efficient and scalable to suit the national-scale modelling and analysis.
- Interpretability. The framework's mechanism and output should be interpretable.

With the above understanding in mind, we have developed an integrated framework for road freight ODM inference.

4. METHODOLOGY

This section describes the inference framework in detail, discusses the methodological and technical choices, and indicates the minimal and desired input data for each module to make the model run and to further improve the performance.

4.1. Framework overview

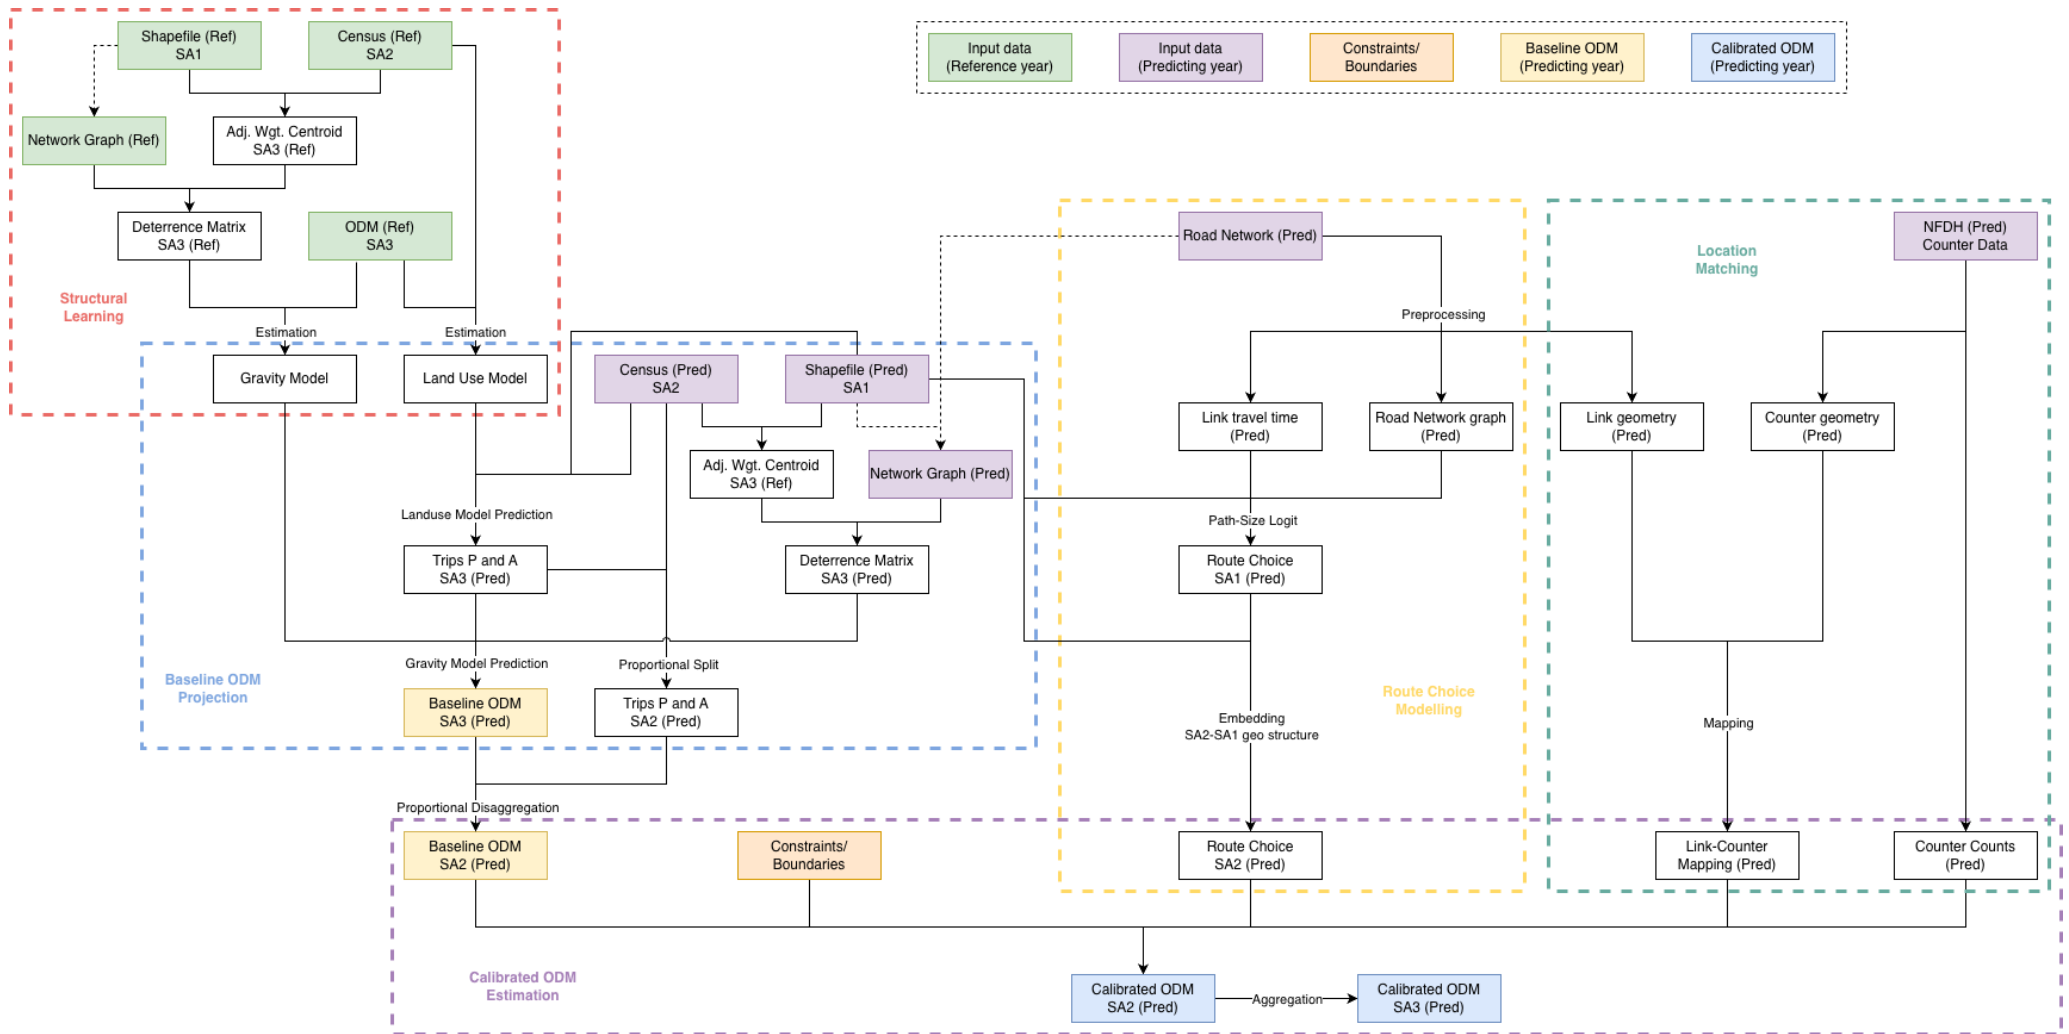
The major challenge of this project is the data deficit, that we only have a few reliable data sources, namely the 2014 ABS Road Freight Movements Survey (RFMS) for road freight trip volume and structure, and traffic counts data for freight traffic reality, but there are many moving parts involved, in order to infer the underlying road freight movements, the ODM. Therefore, acknowledging the current data limitations and anticipating future data availability, the inference framework is designed to be flexible enough to work with the minimum level of data requirements to achieve reasonably good results, and to improve its performance by incorporating new data sources once they become available.

Figure 4 below presents the structural design for the complete framework. The final purpose of the framework is to generate the best guess of the underlying road freight movements at the SA3 resolution (the same as in the 2014 ABS RFMS) with the available data. At the high level, the framework is designed to have 5 modules, and each module can work independently, and can be adjusted or improved with new data and more sophisticated algorithms.

The modules contained in the framework and their purposes are:

- **Structural Learning:** This module aims to estimate a structural model from the 2014 ABS RFMS, alongside the geographical and socio-demographic data that we obtained from the ABS.
- **Baseline ODM Projection:** This module combines the structural model (the product of the Structural Learning module) with the socio-demographic data and other data sets (e.g., the deterrence matrix) in the future year, to generate the projected baseline ODM for the future year. This baseline ODM represents the best guess of the underlying road freight ODM, assuming the structural information learnt from the 2014 ABS RFMS hold stable.
- **Route Choice Modelling:** This module aims to generate the representative route choice between each pair of SA2 zones (the reason for using the SA2 level route choice will be discussed in detail in the corresponding subsection), and this route choice can be model-based, empirical data-based, or ideally, a hybrid between the two.
- **Location Matching:** This module aims to match the locations of traffic data records (e.g., traffic counters, Bluetooth detectors) to the links in the road network. Only by matching the data, the road network (represented by nodes and links), route choices (represented by the probability of choosing a link) and the observed traffic data (such as traffic counts) could work together in the final ODM inference process.
- **Calibrate ODM Estimation:** This is the final module that unites the products of previous modules together to calibrate the baseline ODM using the observed traffic realities, and finally to generate the best guess of the underlying road freight ODM at the SA3 level.

Figure 4 - Framework Overview



The following subsections go through each module in the framework and provide more detailed descriptions on the methodological/technical choices, the rationale behind the decision and the detailed steps and tasks. Please note that while the modularized framework itself is deliberately designed to be flexible enough to host specific methods/algorithms that are paired with different level of data richness, as a working prototype, the actual method/algorithm adopted by the pilot region (Tasmania) modelling exercise is constrained by the current level of data availability. Therefore, in the following subsections, both baseline methods and potential improvements will be described, accompanied by the necessary data requirements.

4.2. Structural learning module

As briefly described in the framework overview, this module aims to learn the structural information from the previous 2014 ABS RFMS (as it is currently the only credible source of the full-scope road freight movements) and pack such structural information into a model. Then, assuming such structural information holds relatively stable in the future, we could use the model, along with other data (e.g., socio-demo data) in the target year, to project a baseline ODM which is largely reasonable for the target year.

Acknowledging that the ultimate purpose of the inference framework is to infer the underlying road freight ODM with the best guess, having this structural learning process provides two advantages:

- **The ability of future baseline projection.** Unlike many studies, including our previous work (Sun et al., 2023), that treat the outdated survey as a static historic snapshot, in this framework, the precious survey data is treated as a living asset which provides the structural anchor to facilitate understanding future road freight demand and distribution. As long as the structural model is proved effective enough to preserve the structural information, it could be used to generate the future baseline of road freight demand and trip distribution to a reasonably good approximation.
- **Conceptually faster calibration speed and better results.** Given that the underlying road freight movements is expected to drift away from the initial surveyed ODM as time passes, from an optimization perspective, the baseline ODM is likely to be closer to the optimal solution. Therefore, using the projected baseline ODM as the seed ODM to initiate the optimization process could conceptually make it faster to reach the optimal solution. Also, due to the complexity of the final optimization, it is possible that the solution space contains multiple local optima, therefore, using the projected baseline ODM as the seed could conceptually make it more likely to reach the global optimum, or at least reach a better local optimum.

As mentioned previously, as a working prototype, the methodological choices are constrained by the data availability in the pilot region, Tasmania. The readily available data for estimating the structural model include:

- The geographical boundary for Statistical Areas (SA1, SA2, SA3 and SA4) in 2011 (corresponding with the 2014 ABS RFMS), 2016 and 2021.
- The Tasmania part of ODM in the 2014 ABS RFMS at the SA3 level.
- The 2011, 2016 and 2021 Census data which contain population and employment information at the SA2 level.
- The historical and current road network obtained from OpenStreetMap (OSM).

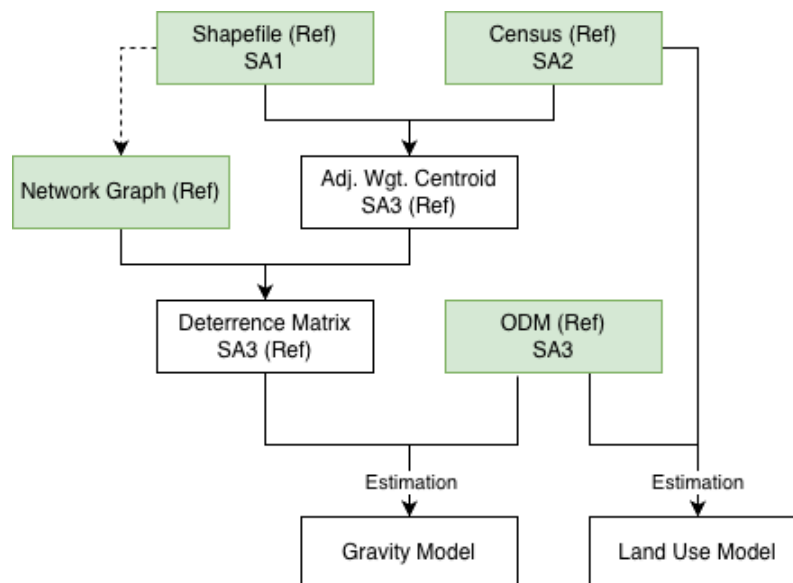
Also, it is worth mentioning that there are only 15 SA3 zones in Tasmania (containing 98 SA2 zones), and that makes 225 SA3 OD pairs. Therefore, not only are the credible data sources

limited, the number of observations at the SA3 level that we could use to learn from the 2014 RFMS are also limited.

It is worth mentioning that we are aware of recent studies, such as the one conducted by researchers in the National Transportation Research Centre, USA (Uddin et al., 2023) that adopt an advanced machine learning model (e.g., Extreme Gradient Boosting, also known as XGBoost) to estimate a structural model that directly predicting the number of total truck trips between OD pairs. However, such models tend to require much more data (e.g., over 200 thousand total trip numbers within and between 583 FHWA zones), and are thus incompatible with the Tasmanian context.

Given the current data availability, the structural model for the Tasmanian context needs to be simple so that can be trained on the very limited data and effective enough to capture the structural pattern in the 2014 RFMS. Therefore, we decided to adopt the land use and gravity model combination to estimate the structural model. **Figure 5** below presents the workflow inside the structural learning module.

Figure 5 - Structural Learning module workflow



There are two main tasks of the structural learning module:

- To estimate a **land use model** using the SA3 level road freight trip production and attraction, respectively, using the SA3 level ODM from the 2014 RFMS, and the (adjusted) SA3 level socio-demo data from the ABS Census.
- To estimate a **gravity model** using a SA3 level deterrence matrix (estimated using a network graph and a set of adjusted SA3 centroid, both will be discussed soon) and the SA3 level ODM from the 2014 RFMS.

Let's start with main steps for estimating the land use models for trip production and attraction:

- (1) **Preprocessing geographical data.** Read the SA1 level ASGS shapefile in 2011 (which fits the 2014 ABS RFMS), for each SA1, find its adjacent neighbour SA1 zones based on the zonal digital boundaries. Then, use this SA1 level connectivity information to construct an **SA1 adjacency graph**. Next, only keep the largest

connected component in the graph, as this represents the scope of connected road network (in the Tasmanian context, it is the road network in the “mainland” Tasmania). Get rid of the other SA1 zones, as they are disconnected to the main road network (in the Tasmanian context, they are offshore SA1 zones such as in the King Island and the Flinders Island). The number of road freight trips are expected to be minimal and are unlikely to be counted into the road freight ODM (as they are offshore). For each SA2 zone, based on the number of its containing SA1 zones in the main SA1 adjacency graph, calculate the **SA2 connected ratio** for later use.

- (2) **Preprocessing population data.** Read the SA2 level population data in the 2011 ABS Census, adjust the SA2 population by the SA2 connected ratio (calculated in step 1). The rationale for this adjustment is to estimate the population covered by the scope of the connected road network, assuming the population in each SA1 zone within a SA2 zone falls into similar range¹⁰. Aggregate the adjusted SA2 population into the adjusted SA3 population. Next, calculate the population density for each SA3 zone, under a prevalent Australian criterion¹¹, the SA3 zones that has population density higher than 200 persons per square kilometre are tagged as “**Urban**”. Repeat the process using the population data in the 2016 Census. Interpolate (assuming linear change) to get the **adjusted SA3 population in May 2014**, which is conceptually consistent to the 2014 RFMS (late 2013 to late 2014).
- (3) **Preprocessing employment data.** Read the SA2 level employment by industry in the 2011 ABS Census, following a prevalent Australian freight-related industry list (BITRE, 2022) (Agriculture, Forestry and Fishing; Manufacturing; Electricity, Mining; Construction; Transport, Postal and Warehousing; Wholesale Trade; Retail Trade), calculate the SA2 level zonal freight-related zonal employment (FRE). Then, adjust the SA2 FRE by the SA2 connected ratio, and aggregate into the adjusted SA3 FRE. Repeat the process using the employment data in the 2016 Census, finally, interpolate to get the **adjusted SA3 FRE in 2014**.
- (4) **Estimating the land use model.** Finally, use the SA3 trip production, attraction, adjusted population, adjusted FRE and “Urban” dummy to estimate the regression-form land use model for trip production and attraction, respectively. Also, to better capture the other factors that impact trip production and attraction but not included in the land use model, we followed a standard industry practice (Cohen et al., 2008; *Quick Response Freight Manual II*, 2007) to calculate a set of **SA3 level scaling factors** (also known as k-factors) based on the ratio between surveyed and predicted 2014 zonal trip generation. This set of zonal scaling factors will be applied to future year projection, assuming such factors are structurally stable. The actual model estimated for Tasmania will be discussed in the case study section.

Now, let’s move to the main steps for estimating the gravity model for trip distribution:

- (5) **Representing the SA3 centroid.** Zonal centroids are typically used as trip starting and ending points in transport planning. However, in the Australian context, the

¹⁰

[https://www.abs.gov.au/ausstats/abs@.nsf/Lookup/by%20Subject/1270.0.55.001~July%202016~Main%20Features~Statistical%20Area%20Level%201%20\(SA1\)~10013#:~:text=Statistical%20Areas%20Level%201%20\(SA1,size%20of%20approximately%20400%20people.](https://www.abs.gov.au/ausstats/abs@.nsf/Lookup/by%20Subject/1270.0.55.001~July%202016~Main%20Features~Statistical%20Area%20Level%201%20(SA1)~10013#:~:text=Statistical%20Areas%20Level%201%20(SA1,size%20of%20approximately%20400%20people.)

¹¹ <https://www.abs.gov.au/statistics/standards/australian-statistical-geography-standard-asgs-edition-3/jul2021-jun2026/significant-urban-areas-urban-centres-and-localities-section-state/urban-centres-and-localities>

geographical centroid is likely to be in the middle of nowhere and does not reflect the representative trip starting and ending points. Therefore, to have more realistic representative trip starting and ending points at the SA3 level, we calculate the **adjusted SA3 centroid** by weighting the underlying adjusted SA2 centroid (by removing the unconnected SA1 zones) using the adjusted SA2 FRE. The rationale is that the literature suggests that the FRE is the main driver of the freight trips (Holguín-Veras et al., 2014, 2011).

(6) **Estimating the SA3 deterrence matrix.** The deterrence matrix sits at the centre of a gravity model. Without any readily available matrix to borrow, choices to construct such matrix include using distance, using travel time, and using a generalised travel cost which contains distance, travel time and unit cost. Also, ideally, the deterrence matrix should be based on the road network and traffic conditions in different years. However, as mentioned previously, as a working prototype, the actual methodology choice is constrained by data availability. In the Tasmanian context, what's available is the historical and current road network which contains roads and their corresponding speed limits. This brings extra challenges that: not only the travel distance and travel time will be based on the choice of network granularity level (e.g., do we include tertiary roads or not), but also the slight network topological differences between years could make the distance and travel time estimation unstable (we have actually done the experiment). Therefore, given such considerations, we decided to use the estimated travel distance based on the **SA1 adjacency graph** to estimate the travel distance between each SA3 zone to form the SA3 deterrence matrix. To be more specific, we firstly overlay the **adjusted SA3 centroids** onto the **SA1 adjacency graph**. Then, the estimated travel distance between each pair of SA3 zones is defined and the summation of:

- a. The Euclidean distance between the centroid of origin SA3 (SA3_O) and its nearest SA1 node in the adjacency graph (SA1_O).
- b. The Euclidean distance between the centroid of destination SA3 (SA3_D) and its nearest SA1 node in the adjacency graph (SA1_D).
- c. The shortest distance from SA1_O to SA1_D based on the SA1 adjacency graph in which the distance between neighbouring SA1 zones is the Euclidean distance between them.

We have randomly compared the Google Map-based distance and the SA1 adjacency graph-based distance between 20 SA3 OD pairs, and found the results to be very close, with the graph-based distance systematically shorter. We acknowledge that using an SA1-adjacency graph is an abstraction of the physical road network, especially, as it only considers the travel distance. However, as suggested by recent studies (Uddin et al., 2023), the travel distance remains the most significant factor that impacts road freight distribution. Therefore, we could expect a distance-only deterrence will still be sufficiently representative.

On the other hand, this approach offers two distinct methodological advantages for a longitudinal inference framework:

Firstly, temporal stability: Unlike OSM or state road networks, which suffer from topological changes and definition shifts over years (making 'distance' an unstable variable), SA1 centroids are geostationary. This ensures that the change in calculated

impedance is driven purely by demand shifts, not by artifacts in the road network dataset.

Secondly, internal consistency: The gravity model parameters are calibrated specifically on this adjacency metric. Consequently, the model “learns” the impedance sensitivity relative to this graph structure, effectively compensating for the systematic underestimation of physical travel distance.

Given such considerations, as long as the gravity model can effectively learn the trip distribution pattern from the 2014 ABS RFMS (will show that in the case study section), it could be confidently used for the structural model estimation and later, for the baseline ODM projection.

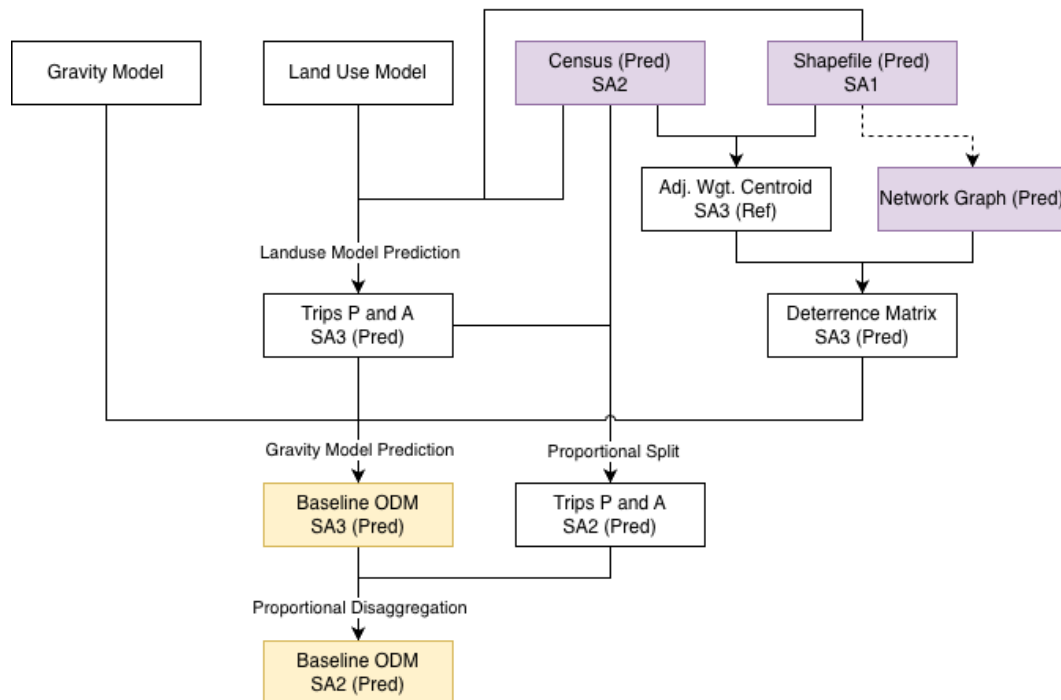
Furthermore, the use of the SA1 adjacency graph decouples the demand generation from the transport supply. By defining impedance based on topological structure rather than network travel time, our baseline projections could be driven exclusively by socio-demographic forecasts (population and employment). This avoids the methodological circularity of requiring a future network model to predict future demand, and ensures that the baseline reflects pure “unconstrained” demand growth, which is the standard requirement for strategic needs assessment.

- (7) **Fitting the gravity model.** Conceptually, between the exponential-form and power-form deterrence function, the power-form is suggested to better suit road freight (Pitfield, 1978), because it better captures the long-tail behaviour of freight movements. So, we will fit a doubly-constrained gravity model using the power-form deterrence function.

4.3. Baseline ODM projection module

The Baseline ODM Projection Module is straight forward, as it combines the estimated structural model with the deterrence matrix and socio-demographic data in the targeting year to project the zonal trip production and attraction, and further projects the trip distribution to generate the baseline ODM for the targeting year. **Figure 6** below presents the workflow inside this module.

Figure 6 - Baseline ODM Projection module workflow



While this module is straightforward in nature, there are some details worth explaining.

- (1) **The socio-demographic data.** In the Australian context, the most credible socio-demo data source is the Census conducted by the ABS every 5 years. The available Census are 2011, 2016 and 2021. Therefore, the same as how we approximate the 2014 population and employment at the SA2 level, for other years between two Census years, we use such linear interpolation to get the approximated SA2 zonal population and employment. For the years after 2021, in the absence of the 2026 Census data, the current interpolation adopts the linear changing rate between 2016 and 2021 Census. We recognize this introduces uncertainty, particularly if growth rates changed post-2021 due to COVID-19 recovery dynamics or other structural shifts. However, such change should be captured by the later ODM calibration process. In the case of inconsistent SA2 boundaries (e.g., We observed a boundary change in Tasmania, regarding the Flinders Island, from 2011 to 2016 Census), we adjust the statistics based on the SA2 zonal area (as the SA1 level socio-demo data is not available) to make the SA2 zones consistent before and after such change.
- (2) **The change in the deterrence matrix.** As described in the Structural Learning module, the deterrence matrix is designed to be distance-based on the SA1 adjacency network. This approach will make the deterrence matrix free from the topology or definition changes in the actual road network and become stable. However, it is worth noting that the year-to-year matrix still changes slightly, this is because the corresponding adjusted SA3 centroids are expected to drift slightly between years, due to the approximately linear change in zonal employment.
- (3) **SA3 level trip production and attraction prediction and adjustment.** With the estimated land use model in the reference year (2014 in this case), and the approximated zonal population and FRE in the targeting year, we firstly predict the SA3

level trip generation, then, adjust them using the scaling factors that we have calculated based on the 2014 data.

(4) **The interpolated SA2 level trip generation and distribution.** This is another point that needs explanation. Because we will use the SA2 level route choice to connect the zonal trip flow to the underlying road network and counter locations (we explain why the SA2 level route choice is chosen in the following Route Choice Modelling module subsection), we need the corresponding SA2 level road freight ODM to facilitate reconstructing the traffic reality (e.g., the observed traffic counts). This auxiliary SA2 ODM is generated in two steps:

- a. Generating SA2 zonal trip production and attraction. For each SA3, the trip production and attraction for its containing SA2 zones is approximated by proportionally splitting the SA3 zonal trip production and attraction by the FRE number in each SA2 zone.
- b. Distributing SA2 level trips. Once we have the SA2 level trip distribution and attraction, we use the following formula to proportionally distribute (it is essentially a naïve gravity model which assumes uniform impedance between SA2 OD pairs under the SA3 OD pair) the SA3 level baseline ODM into the SA2 level, that:

$$Flow_{i_n j_m} = Flow_{i,j} * \frac{P_{i_n}}{P_i} * \frac{A_{j_m}}{A_j}$$

Where:

i stands for the SA3_O and i_n stands for the SA2_O within the SA3_O;

j stands for the SA3_D and j_m stands for the SA2_D within the SA3_D;

P stands for the zonal trip production and A stands for the zonal trip attraction.

This approach is essentially a simplified gravity model at the SA3 pair level, it allows the approximated SA2 level ODM to aggregate to the SA3 level ODM in a consistent manner and it is in line with existing studies (Horowitz, 2010; Viswanathan et al., 2008). While we will only report the SA3 level baseline ODM as part of results of the integrated framework, the necessity of the underlying SA2 level ODM projection will be further discussed in the later ODM calibration module.

4.4. Route choice modelling module

As introduced in existing studies based on the single-level data fusion frameworks (Behara et al., 2022; Ma and Qian, 2018; Sun et al., 2023), a route choice table is required to connect zonal trip flows to the underlying road network, and such route choice could be either derived algorithmically (Dey et al., 2020), or from the empirical observation (Sun et al., 2023). In the Tasmanian context, in the absence of empirical route choice, we adopt the classical Path-Size Logit model (Ben-Akiva and Bierlaire, 1999; Ramming, 2002) to estimate the route choice. Also, it is worth mentioning that in the absence of the officially maintained road network data, we use the road network downloaded from OpenStreetMap to derive the route choice. More specifically, we use the road network as at the end of August 2025, and the type of roads included in the road network are Motorway, Motorway Link, Trunk, Trunk link, Primary, Primary link and Secondary. Here we have made two decisions:

- (1) **The choice of a static snapshot of road network.** As mentioned in previous sections, ideally, we would like to use the year-specific road network to derive the route choice for each year (actually, if to pursue the maximum level of accuracy, one should use the

day-by-day road network to reflect the actual availability of roads), however, after experimentation (e.g., we have downloaded the Tasmanian road network from 2014 to 2025, one snapshot for each year), we found that the road topology and travel time estimation tends to be slightly inconsistent across years, and such inconsistency could introduce unwanted instability to the route choice model. Therefore, we chose to use a snapshot of the road network in 2025 for deriving the route choice, given the side-by-side eyeball comparison suggests the minimal road network change between 2014 and 2025 in the Tasmanian context.

- (2) **The choice of included road classes.** To reflect the fact that freight vehicles may be restricted from using certain roads, we resorted to the OSM Australian Road Tagging Guidelines¹² to only include the major roads from motor ways to secondary roads. Other roads tagged as tertiary or lower hierarchies are excluded as they are mainly minor through roads.

We acknowledge that this road network configuration is a simplification over the temporally changed actual road network for freight, it is accessed to be a solid starting point to test the inference framework with the limited available data sources, and we could easily replace it with the officially maintained actual road network data once it becomes available.

After settling the road network configuration, we adopted the following steps to derived route choices, based on the Path-Size Logit (PSL) method (Ben-Akiva and Bierlaire, 1999; Ramming, 2002).

- (1) **Route Choice Set Generation:** Generating a behaviourally realistic choice set is a prerequisite for PSL, yet it presents a significant challenge in large-scale networks due to the computational explosion of potential paths. Simple k-shortest path algorithms often fail in dense urban grids, as they produce sets dominated by minor link-level variations rather than strategically distinct corridors (Bovy, 2009). To address this, we adopted a stochastic route generation method (Fiorenzo-Catalano, 2007). Origins and destinations were defined at the Statistical Area Level 1 (SA1) to ensure the richness of route choices at a more aggregated level (e.g., SA2 and SA3 levels). Centroids were mapped to the nearest network nodes. For each OD pair, we performed repeated shortest-path searches on a perturbed network. In each iteration, link travel times were modified by a link-classification-specific penalty factor (Bekhor et al., 2006; Bovy and Fiorenzo-Catalano, 2007). By applying higher variance to major arterials and motorways (e.g., A1 vs. M31), we effectively forced the algorithm to explore distinct regional corridors, successfully filtering out the 'noise' of minor variations in local city streets.
- (2) **Path size calculation.** For each SA1 OD pair, for each shortest travel time path recorded, we follow the standard PSL method calculate its path-size using the formula below:

$$PS_i = \sum_{links\ a \in i} \left(\frac{length\ of\ link\ a}{length\ of\ route\ i} \right) \times \frac{1}{Number\ of\ routes\ using\ link\ a}$$

- (3) **Calculating the probability of selecting each path.** Based on the calculated path-size and the road network data, in the absence of locally calibrated parameters, we adopted the formula in an established study (Hess et al., 2015) to estimate the fuel and other costs by taking the path. And then, using the formula below to combine cost,

¹² https://wiki.openstreetmap.org/wiki/Australian_Tagging_Guidelines/Roads

travel time and path size for each path into the utility of taking the path (beta parameters also borrowed):

$$V_i = \beta_{cost} * \ln(C_i) + \beta_{time} * \ln(T_i) + \beta_{pathsize} * PathSize_i$$

Then, we use a standard Multinomial Logit (MNL) formula to calculate the probability of taking a path for each SA1 OD pair.

- (4) **Aggregating the route choice into SA2 level.** While route choice sets were generated at the granular SA1 level to capture local accessibility, using SA1-level OD matrices for national-scale calibration is computationally prohibitive. Therefore, we aggregate the resulting route choice probabilities to the Statistical Area Level 2 (SA2). This aggregation relies on the assumption of intra-zonal spatial homogeneity: that within a given SA2, the route choice profiles and Freight Related Employment (FRE) of constituent SA1s are sufficiently similar. Therefore, such SA1 number-weighted aggregation will be conceptually equivalent to the FRE-weighted aggregation. Such aggregation reduces the optimization dimensionality significantly while retaining better route fidelity than if paths were generated directly from large SA2 centroids.

4.5. Location matching module

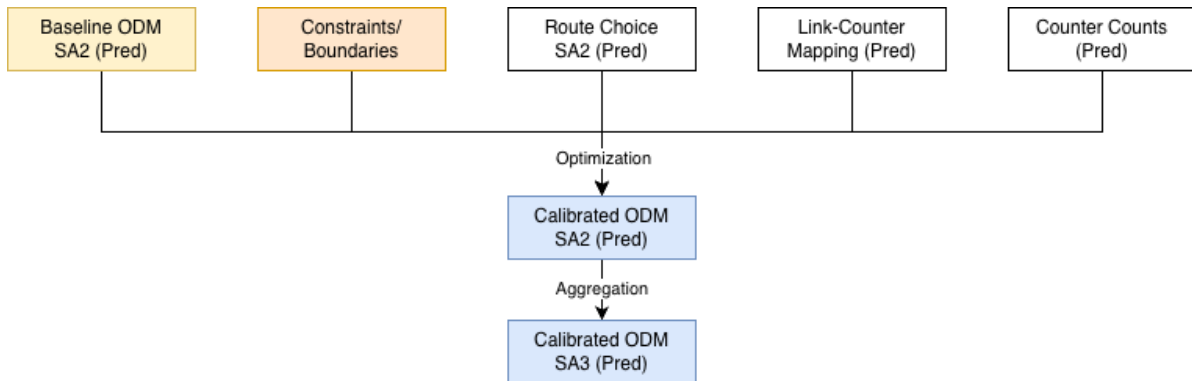
As mentioned previously, it is of absolute necessity to represent the locations of traffic detectors (e.g., traffic counters, Bluetooth readers) to the links in the road network, to connect O-D flows, route choices and traffic realities into a system. Ideally, with sufficient and consistent meta data between road network and traffic detectors (such that a traffic counter is associated with road X, and such road X could be identified uniquely in the road network data), such mapping process could be straightforward and reliable. However, in reality, we may not have such consistency between datasets, and the matching process becomes a task by itself.

For this experiment in Tasmania, since we only have less than 50 permanent traffic counter locations across the island, we have manually matched them to the links in the major road network to ensure a high-quality matching. However, this manual method will be infeasible when we scale up the framework to cover mainland Australia which could have thousands of traffic counters. Further, since the framework allows incorporation other types of data collectors such as Bluetooth detectors, that means the location-link matching is required for each additional type of data collectors, this makes the manual matching even more infeasible. Therefore, we will develop an automatic location-link matching and checking script in the follow-up project.

4.6. Calibrated ODM estimation module

This is the final step where we unite results from previous modules under the single-level inference framework (Sun et al., 2023) to calibrate the baseline ODM projection. The final product is the most likely underlying ODM (the Calibrated ODM) that matches the traffic realities and best fits with the prior structural knowledge. The input and output of this module is represented in [Figure 7](#) below.

Figure 7 - ODM Calibration Module



While this module seems straightforward, there are several arbitrary designs in the optimization process that need further elaboration:

- (1) **Hierarchical regularization.** As suggested by the module workflow, the final calibration process takes the baseline ODM projection at the SA2 level as the seed SA2 ODM to initiate the optimization process. At the very high level, this process uses the SA2 level OD flows along with the SA2 level route choices to re-generate the observed traffic reality, and this makes it a classical under-determined problem for which there will be multiple SA2 ODM that could match the traffic counts. To achieve a reasonable ODM, we designed a hierarchical regularization to guide the optimization algorithm. The **upper-level regularization** uses the volume, and structural information contained in the baseline SA3 ODM, to make the trip volume match the baseline ODM (as it is from an effective land use model) and minimize the structural difference between inferred SA3 ODM (which is based on the inferred SA2 ODM) and the baseline SA3 ODM. However, such SA3 level regularization is not sufficient, as the underlying SA2 ODM (also they can be aggregated into a reasonable SA3 ODM) could be unrealistic. Therefore, the lower-level regularization focuses on preserving a reasonable SA2 level structure. Specifically, the **lower-level regularization** aims to minimize the structural difference between the inferred SA2 ODM and an inferred SA2 ODM alternative. Such an inferred SA2 ODM alternative is the proportional disaggregated SA2 ODM based on inferred SA3 ODM, by the SA2 zonal FRE, exactly how we derive the baseline SA2 ODM projection in the Baseline ODM Projection Module. There is a relevant nuance that, the guided proportions for disaggregating the inferred SA3 ODM is based on the static SA2 zonal FRE, not the dynamic SA2 zonal trip production and attraction. This is a conscious choice to preserve the underlying logic of the land use model, and use it as an anchor for the ODM calibration process, which enforces internal spatial consistency.
- (2) **Optimization target and constraints.** Given the hierarchical regularization, the optimization target is to find an optimal SA2 ODM that minimises the weighted sum of:
 - a. The structural distance between the inferred SA3 ODM and the baseline SA3 ODM projection, and
 - b. The structural distance between the inferred SA2 ODM and the inferred SA2 ODM alternative.

In terms of the system constraints, they are:

- a. Preserve the total trips in the baseline SA3 ODM projection, as it is the prediction from an effective land use model.
 - b. Preserve the traffic counts with a small percentage variation allowed (e.g., 1%) to reflect the potential data quality issue, as they are the observed traffic reality, the only “ground truth” in the system.
- (3) **Interpretation of results.** The direct product of the calibration (optimization) process is the inferred SA2 ODM. However, given the under-determined nature of OD estimation, where traffic count constraints are sparse relative to the number of SA2 pairs, individual SA2-to-SA2 flows possess a higher degree of uncertainty. (Noting that future availability of GPS telematics data or other trajectory-enabled data would allow direct validation of SA2 flows and potentially relax the hierarchical structure.) Therefore, we report the SA3-level aggregated ODM as the primary output for strategic planning. The SA2-level flows function as intermediate regularizers that ensures the aggregate SA3 flows are spatially plausible (consistent with land-use distribution) and topologically valid (consistent with network routing), but they are not intended to be interpreted as precise point-to-point forecasts. Therefore, the result reported for facilitating decision making is the **calibrated SA3 ODM** which represents the most likely underlying freight flow given the available data. Another set of results worth analysing are the two structural differences (SA3 level and SA2 level). Conceptually, the **SA3 level structural difference** measures how much the calibrated SA3 ODM deviates from the projected baseline SA3 ODM, this indicates how much the power of structural information that we learnt from the original RFMS has decayed. In terms of the **SA2 level structural difference**, it measures how badly the underlying SA2 ODM has to violate the original land use patterns to make the SA3 ODM structurally consistent.

4.7. Value delivery

While this project was originally scoped as the “road freight ODM inference”, the developed integrated inference goes beyond the initial scope and could potentially contribute to BITRE’s transportation planning duty in three ways:

- (1) Enabling long-term road freight movements baseline projection.
- (2) Providing the insight of most likely underlying road freight movements, given the available data. And the accuracy of the calibrated ODM could improve as more data becomes available. This is about maintaining institutional knowledge of freight structure between survey cycles.
- (3) Allowing monitoring how the power of the survey decays over time, and providing evidence to support a new survey, or deferring a new survey.

5. TASMANIA CASE STUDY

This section describes the case study for the pilot region, Tasmania, in detail, including the input data, estimated structural model and output of the framework.

5.1. Input data

5.1.1. Census data

This section describes the ABS Census data regarding population and freight-related employment in Tasmania. In terms of population (Figure 8), the total population in Tasmania increased from 494 thousand person in 2011 Census to 509 thousand persons in 2016, and to 556 thousand persons in 2021 Census, and the interpolated population in 2014 is 503 thousand persons. Figure 9 further plots the population density (population per squared km) in 2014, it is clear that the population density in most Tasmania zone is low, based on the 200 population per squared km criteria, only Hobart Inner (~800), Hobart Northwest (~310) and Launceston (~310) are classified as urban areas.

Figure 8 - Tasmania SA3 population




Figure 9 - Tasmania SA3 population density (2014)





Figure 10 shows the total freight-related employment numbers, it increases from 78 thousand employed persons in Census 2011 to 89 thousand employed persons in Census 2021, and the interpolated figure for 2014 is slightly above 78 thousand employed persons.

Figure 10 - Tasmania SA3 freight related employment



5.1.2. Geography

Figure 11 presents the SA1 adjacency graph based on the 2011 ASGS boundaries at the SA1 level - 2011 boundaries are chosen as it aligns with the 2014 RFMS. In the graph, each node represents a SA1 zone, and the link connecting two neighbouring zones represents the Euclidean distance between their geographic centroids. In total, 1433 out of 1444 SA1 zones in Tasmania are included in the adjacency graph; the excluded zones are mainly located in offshore islands (e.g., King Island, Flinders Island and Cape Barren Island).

Figure 12 presents the FRE-adjusted centroid of each SA3 zone. As described in the methodology section, such weighting is based on the adjusted SA2 zone geography (by excluding the SA1 zones that are not in the SA1 adjacency graph) and the proportionally adjusted FRE (based on the number of connected SA1 zones and the total number of SA1 zones in each SA2 zone).

The approximated SA3 distance matrix for estimating the gravity model is calculated using the SA1 adjacency graph and the adjusted SA3 centroids. The distance matrix for 2014 is shown in **Figure 13**, noting that for each SA3, the approximated intra-zonal distance is set as half of its shortest inter-zonal distance. As shown in the approximated distance matrix, the SA3 O-D pair distances range from 5km (intra-zonal distances in Hobart Inner and Hobart Northwest) to over 300km (inter-zonal distances between west coast and southeast zones). We did some manual checks using Google Maps and found that the approximated distances tend to be shorter by around 10-15% (e.g., the map-based distance between Launceston and Hobart Inner around 195 km, compared to 174 km based on the SA1 adjacency graph, 91 km vs 81 km for between Devonport and Launceston, and 385 km vs 339 km for between West Cost and Southeast Cost), representing a good proxy for zonal distances.

Figure 11 - SA1 adjacency graph



Figure 12 – SA2-FRE-adjusted SA3 centroid

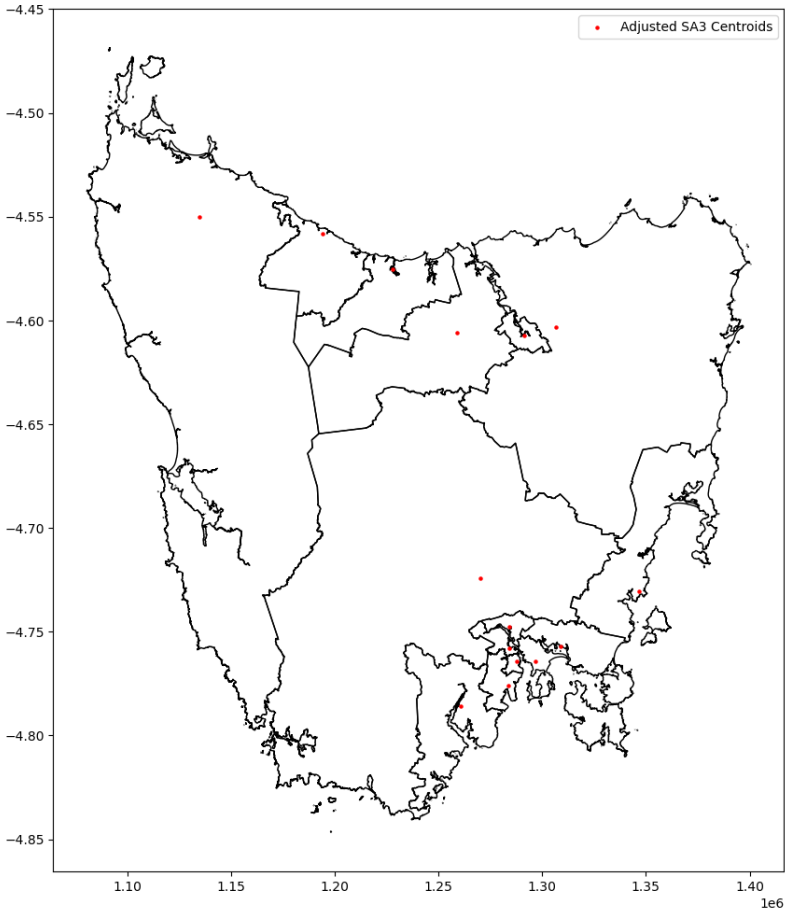



Figure 13 - Approximated SA3 distance matrix (2014)



5.1.3. Major road network

Figure 14 presents the main road network in Tasmania downloaded from the OpenStreetMap using the OSMnx Python package (Boeing, 2025), as at 30 August 2025. As described previously, only the major roads (hierarchically from Motorway to Secondary roads) are kept representing the likely accessible roads to road freight vehicles. Please note that this road network can be replaced with officially-maintained road network data once it is available in the future. The link (the graph theory concept, could be a chunk of a road) count by road classification and the corresponding average speed limit is presented in the table below.

Figure 14 - Main road network in Tasmania

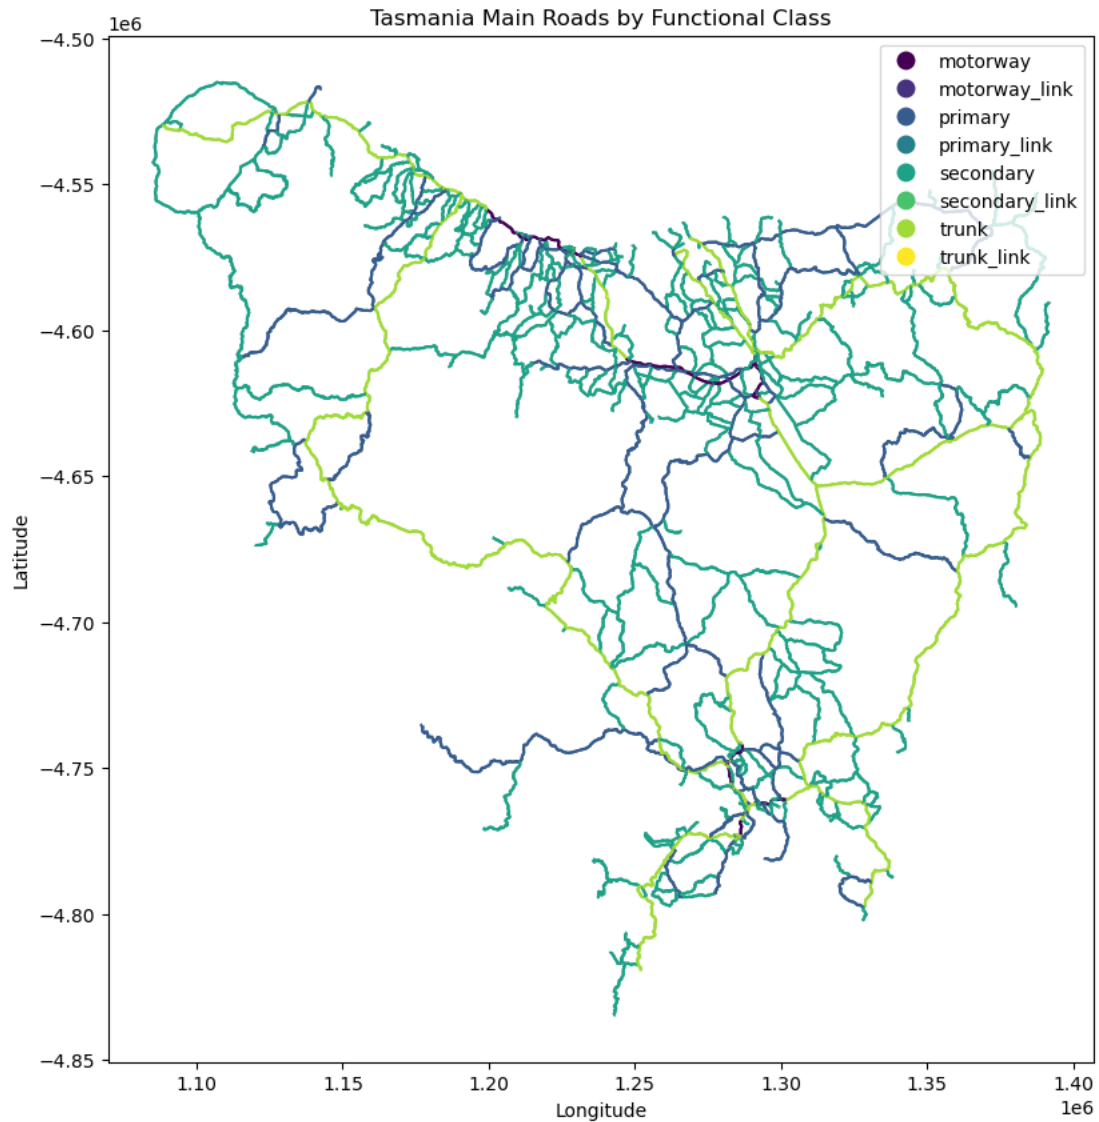


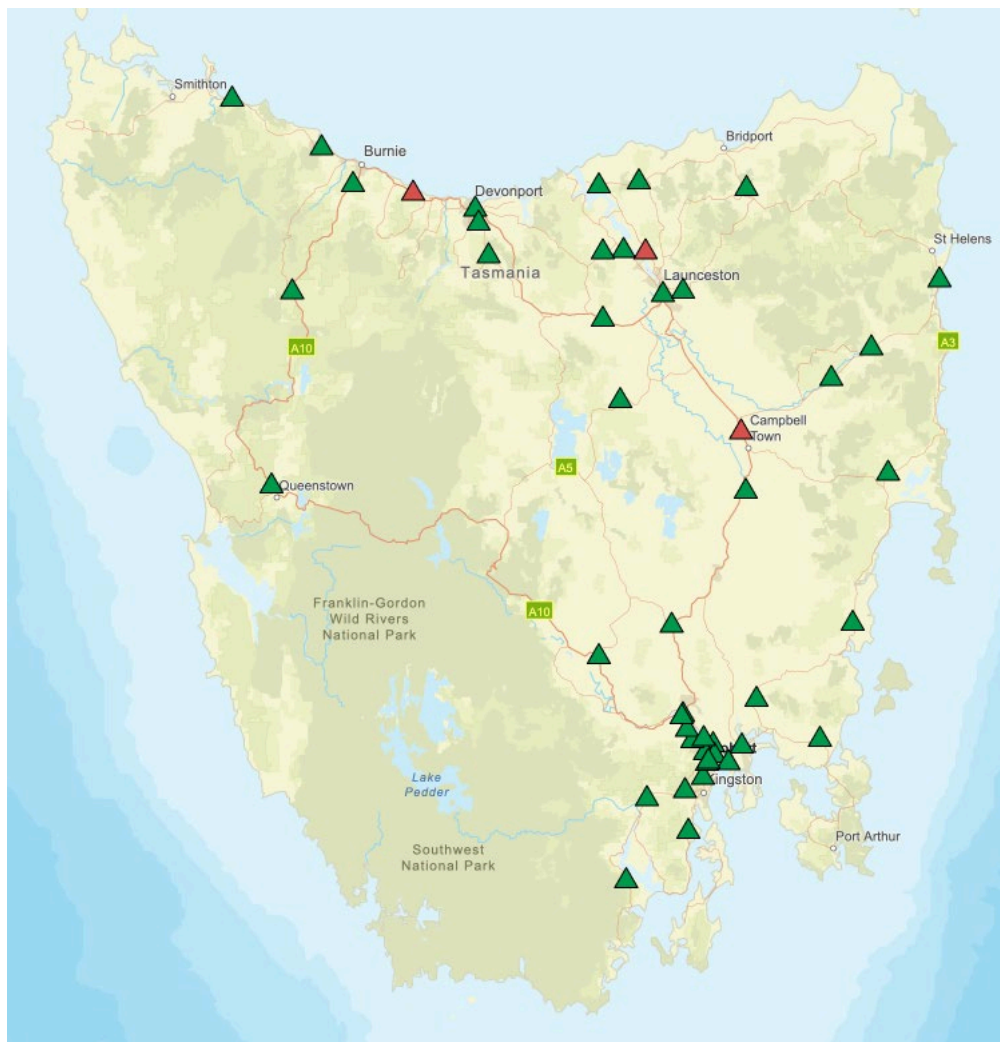
Table 1 - Main road network summary (Tasmania 20250830)

Road Classification	Count	Average speed limit (km/h)
Motorway	209	105
Motorway Link	182	88
Trunk	1256	82
Trunk Link	153	65
Primary	1216	73
Primary Link	69	65
Secondary	2031	64
Secondary Link	97	57
Total	5231	

5.1.4. Traffic counters

Figure 15 presents the permanent traffic counter locations in the main island Tasmania, noting that there were 42 traffic counters available as at the end of 2024, and 25 among them were consistently available from 2015, providing the Average Annual Daily Traffic (AADT) count and corresponding percentage of heavy vehicles (% HV), which we use as a proxy for road freight vehicle counts. The use of consistently available traffic counters is further discussed in the later sub-sections.

Figure 15 - Permanent traffic counters in Tasmania (as at the end of 2024)



5.2. Estimated structural model

5.2.1. Land use model estimation results

Based on the method that is described in the structural model learning section, the final land use model will be in the form of: $Prod_i = Const. + \beta_1 * FRE_i + \beta_2 * Urban_i + \varepsilon$ where i stands for a SA3 zone. The same structure applies to zonal attractions. Noting that this form is very minimal and the zonal population is not included.

Ideally, we want to keep both zonal population and FRE in the formula to better capture the population-driven part of truck freight movement. However, we finally drop it from the model because of the reasons below:

- Zonal population is highly correlated with zonal FRE (with a correlation coefficient above 0.9).
- Zonal population is associated with very high model Variance Inflating Factor (VIF).
- The model dropping population (adj. R^2 about 0.83) significantly outperformed the model dropping FRE (adj. R^2 about 0.62), and this is consistent with the literature that freight movement is primarily driven by employment (Holguín-Veras et al., 2014, 2011).
- We have tried using Principal Component Analysis (PCA) to generate a composite variable which incorporates information from both population and FRE (this variable is called sociodemographic activity level in some literature), and using this new variable in the model. However, it still significantly underperformed the model using FRE.

Therefore, we finally decided to drop zonal population, and only use FRE and urban dummy in the land use model. Figures 16 and 18 present land use model estimation results for trip production and attraction, respectively. In general, we could see that based on Tasmanian RFMS results, and the corresponding interpolated socio-demo data, one additional zonal FRE tends to increase the annual zonal road freight trip production by 112 in rural areas, and such effect is 78 in the urban areas. Similar effects apply to the annual trip attractions. Lastly, we calculated a set of **SA3 level scaling factors** (also known as k-factors) (Cohen et al., 2008; *Quick Response Freight Manual II*, 2007) based on the ratio between surveyed and predicted 2014 zonal trip generation. Figures 17 and 19 show that zonal scaling factors mainly vary between 0.5 and 1.5, except for around 2.7 in Brighton. This set of scaling factors will be carried into the baseline ODM projection for future years, assuming such factors are structurally stable.

Figure 16 - Land use model (Trip production)

OLS Regression Results						
Dep. Variable:	Production	R-squared:	0.855			
Model:	OLS	Adj. R-squared:	0.831			
Method:	Least Squares	F-statistic:	35.51			
Date:	Mon, 22 Dec 2025	Prob (F-statistic):	9.12e-06			
Time:	10:56:28	Log-Likelihood:	-197.74			
No. Observations:	15	AIC:	401.5			
Df Residuals:	12	BIC:	403.6			
Df Model:	2					
Covariance Type:	nonrobust					
	coef	std err	t	P> t	[0.025	0.975]
Intercept	1443.2667	7.29e+04	0.020	0.985	-1.57e+05	1.6e+05
Total_Employment	111.5091	16.225	6.873	0.000	76.159	146.859
Total_Employment:Urban	-33.4634	13.704	-2.442	0.031	-63.321	-3.605
Omnibus:	2.739	Durbin-Watson:	1.681			
Prob(Omnibus):	0.254	Jarque-Bera (JB):	1.540			
Skew:	0.784	Prob(JB):	0.463			
Kurtosis:	2.950	Cond. No.	1.50e+04			

Figure 17 - Trip production comparison

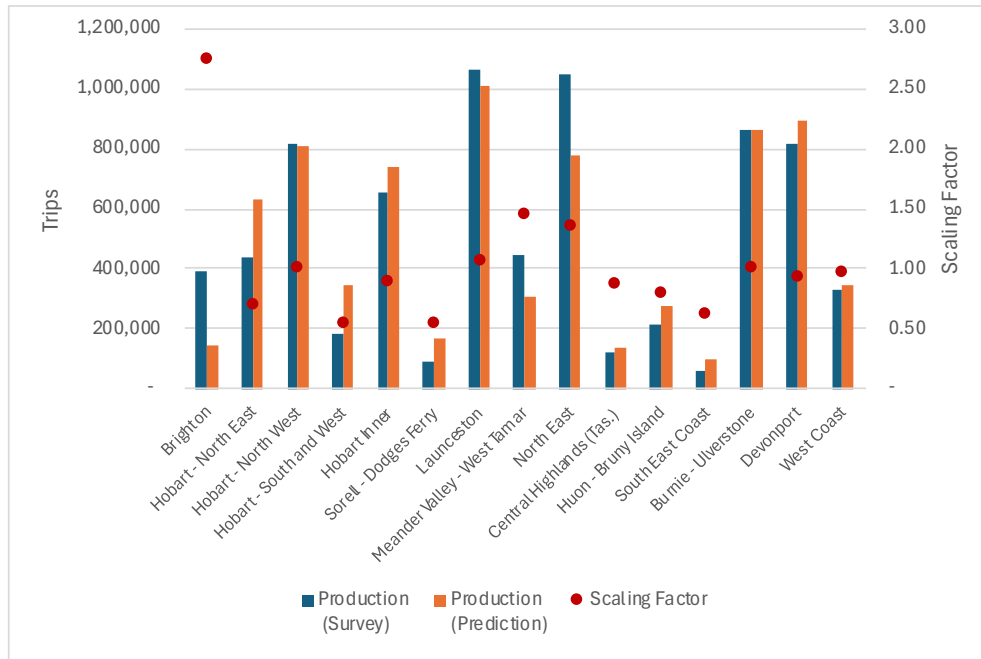


Figure 18 - Land use model (Trip attraction)

OLS Regression Results

```

=====
Dep. Variable:          Attraction      R-squared:                0.857
Model:                  OLS             Adj. R-squared:           0.833
Method:                 Least Squares  F-statistic:              35.94
Date:                   Mon, 22 Dec 2025  Prob (F-statistic):       8.58e-06
Time:                   10:56:30       Log-Likelihood:          -197.72
No. Observations:      15             AIC:                     401.4
Df Residuals:          12             BIC:                     403.6
Df Model:               2
Covariance Type:       nonrobust
=====

```

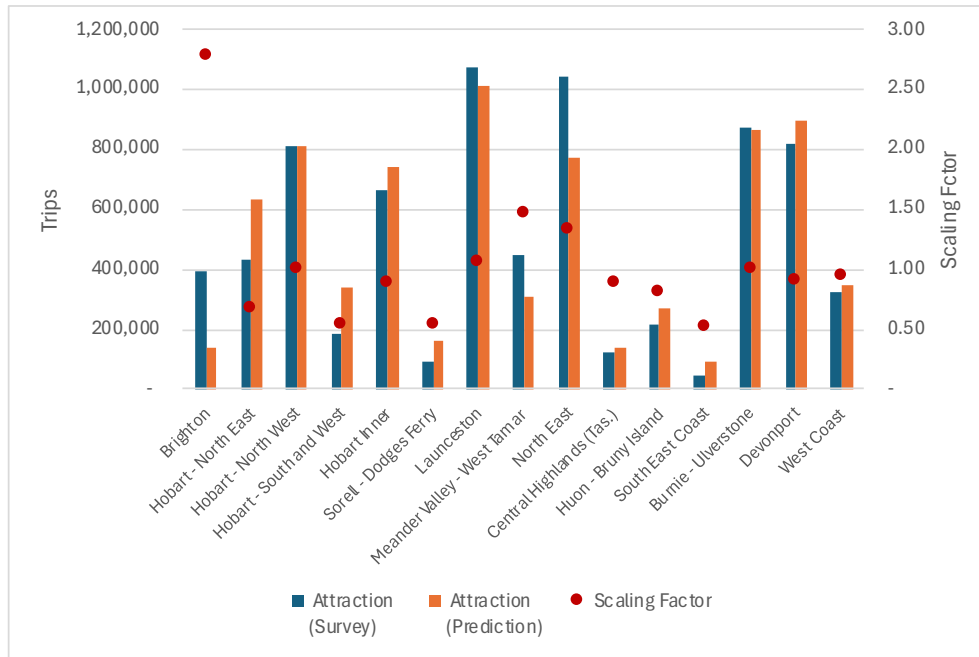
	coef	std err	t	P> t	[0.025	0.975]
Intercept	-580.8126	7.28e+04	-0.008	0.994	-1.59e+05	1.58e+05
Total_Employment	111.8945	16.202	6.906	0.000	76.593	147.196
Total_Employment:Urban	-33.4464	13.685	-2.444	0.031	-63.263	-3.629

```

=====
Omnibus:                2.549      Durbin-Watson:            1.659
Prob(Omnibus):          0.280      Jarque-Bera (JB):         1.443
Skew:                   0.759      Prob(JB):                 0.486
Kurtosis:               2.915      Cond. No.:                1.50e+04
=====

```

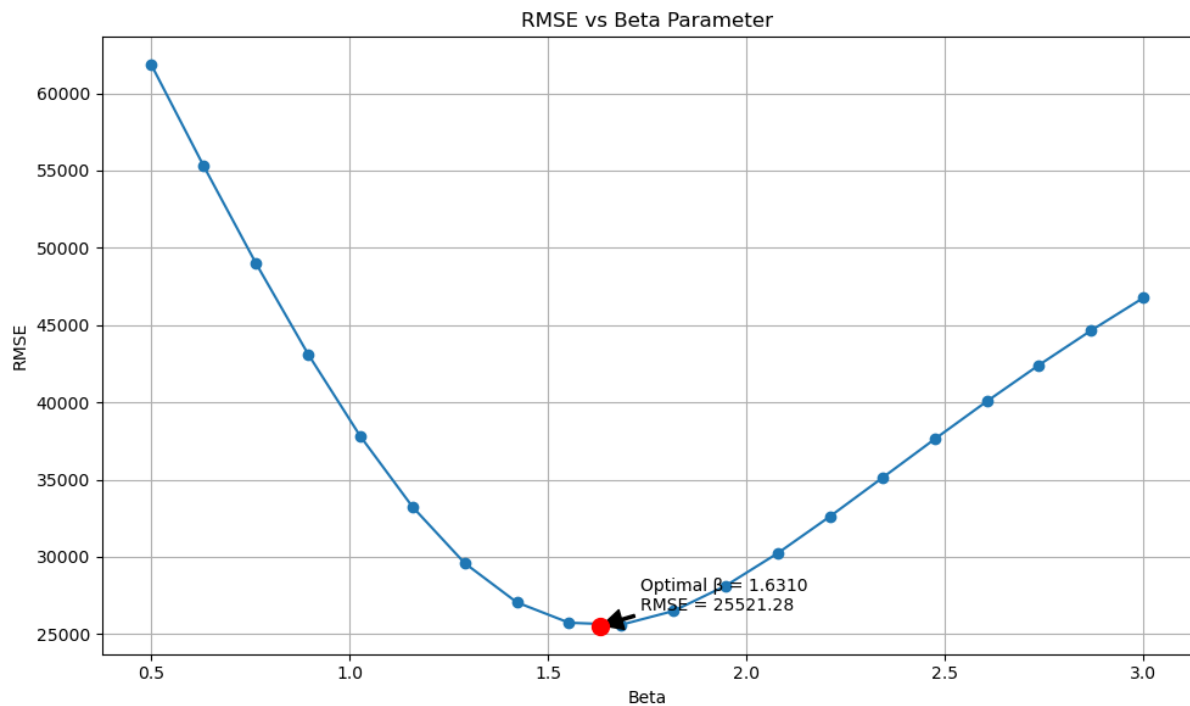
Figure 19 - Trip attraction comparison



5.2.2. Gravity model estimation results

As described in the methodology section, we choose the power deterrence function form that $F_{ij} = d_{ij}^{-\beta}$, and fit it into a doubly-constrained gravity model to find the optimal beta that minimizes the RMSE from predictions. The SA3 level deterrence matrix used in the gravity model estimation is the approximate SA3 level distance matrix, given FRE-adjusted SA3 centroids for 2014 and the underlying SA1 adjacency graph. Figure 20 presents the gravity model fitting result, which shows that the optimal β which minimizes the RMSE between the predicted and surveyed trip distribution is 1.6310, without further data to facilitate further breaking down by commodity types, this β value represents a reasonable combination of typical heavy, long-haul trucks (around 0.8) and light, urban service trucks (around 2.2), and is close to the baseline value of 1.5, as suggested by the textbook (Ortúzar S. and Willumsen, 2024). Also, the R-squared for the estimated gravity is 0.9279, indicating a good fitting on the surveyed trip distribution.

Figure 20 - Gravity model fitting

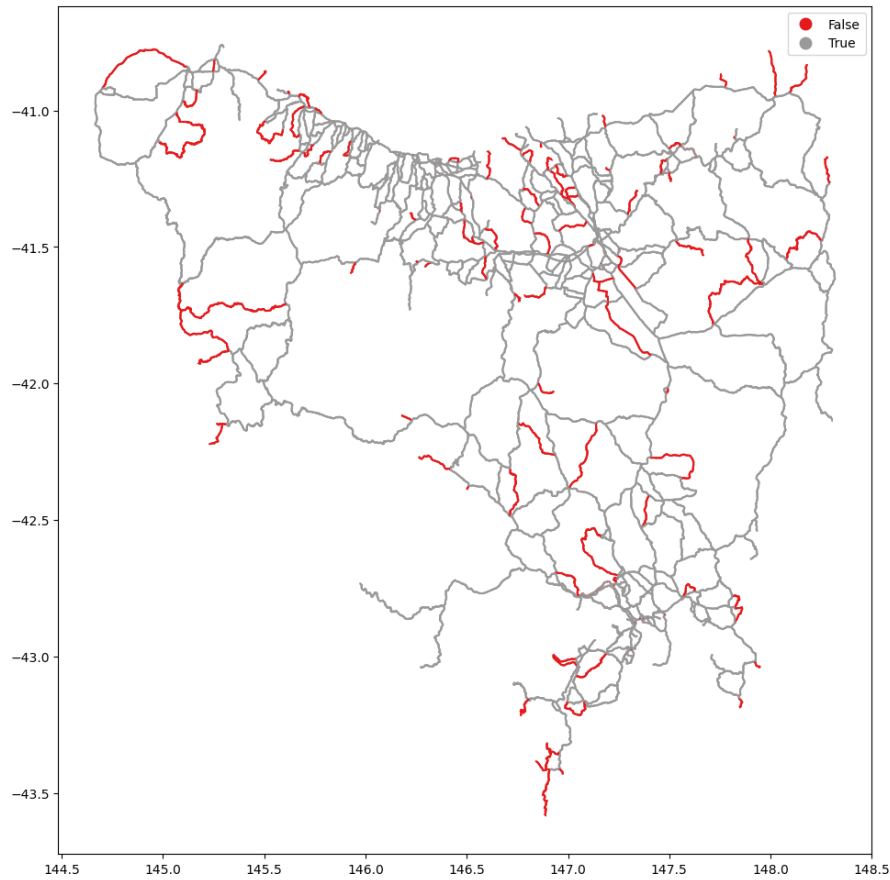


5.3. Route choice modelling

This section describes the modelled SA2 level route choice, based on the PSL method that was described in the methodology section.

In terms of the link coverage, 4776 out of 5213 links are used in the PSL modelled route choices, between SA2 O-D pairs, and most unused links are secondary roads which tend to have lower speed limits (so longer expected travel times), or some obvious detours. **Figure 21** presents the used and unused links in the Tasmania main road network.

Figure 21 - Links used in PSL-based route choice



Figures 22 to 25 present some comparisons between the modelled route choices and the corresponding Google Maps suggestions. In most cases, the modelled route choice is highly consistent with the Google Maps suggestions (as shown in plots) or is the same as the suggested route with the shortest estimated travel time. Also, we are aware that Google Maps assumes a small vehicle so it could take tertiary roads, which are not included in our main road network for freight vehicles, and this could result in some minor route selection differences. In general, without available empirical route choice data, we think the PSL-based modelled route choices are reasonable, and they satisfy the minimum data requirement in the integrated ODM inference framework.

Figure 22 - Route Choice (Smithton to Launceston) - PSL

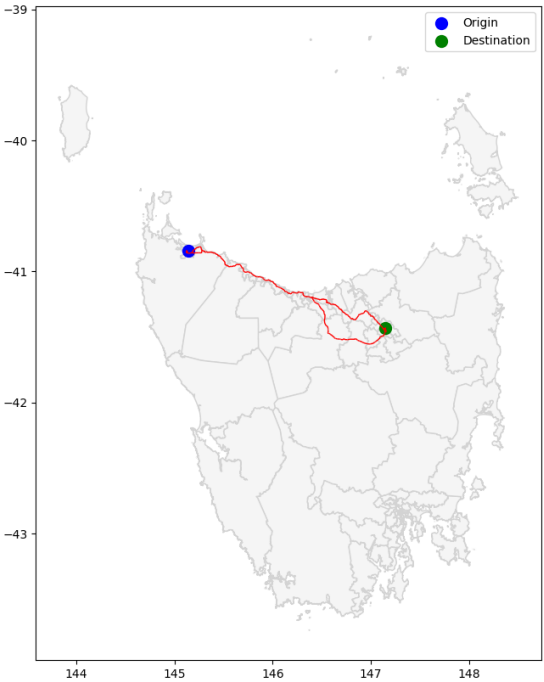


Figure 23 - Route Choice (Smithton to Launceston) - Google Maps

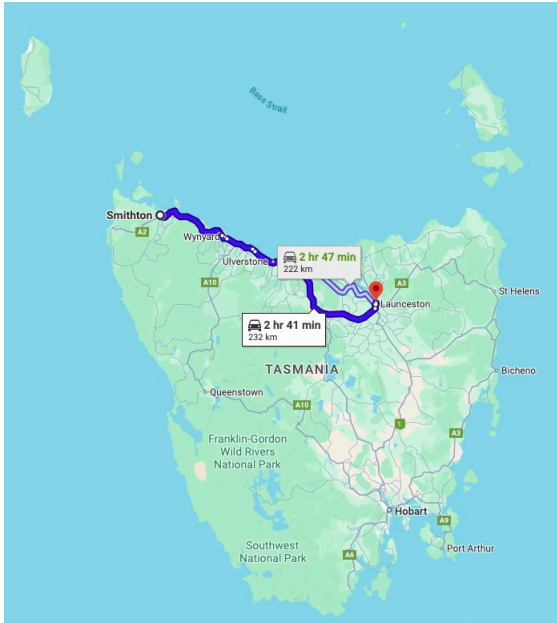


Figure 24 - Route Choice (Hobart to Scottsdale - Bridport) - PSL

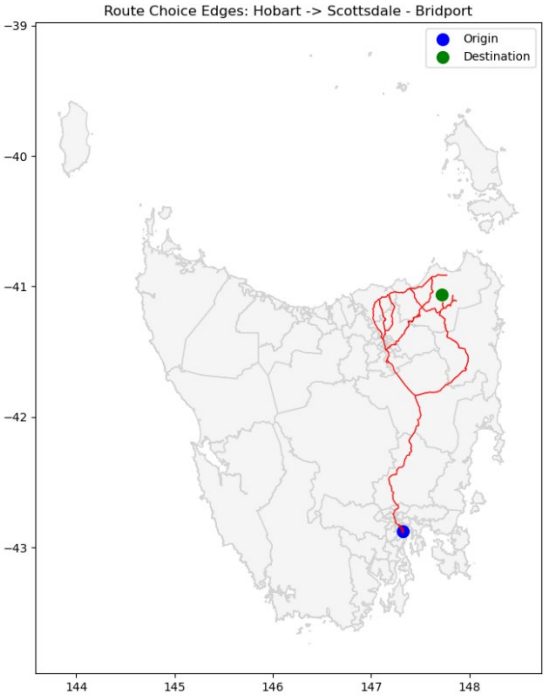
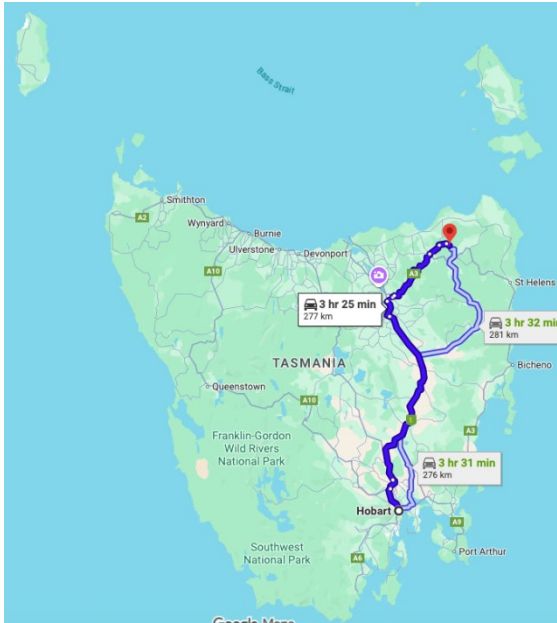


Figure 25 - Route Choice (Hobart to Scottsdale - Bridport) - Google Maps



5.4. Structural model effectiveness check

To validate the effectiveness of the structural model, which is crucial to allow it being used for baseline ODM projection in the future years, we firstly project the baseline ODM in 2014 (the RFMS year), then compare it to the RFMS ODM. The purpose of such comparison is to check if the difference between such two ODMs is statistically insignificant.

Figures 27 and 28 present the surveyed Tasmanian ODM in RFMS (left) and the projected baseline ODM for 2014 (right). The cosine similarity between two ODMs is about 0.96, suggesting the that two ODMs are structurally very similar (and their total trip volumes are the same), while we could still observe some noticeable differences between some OD pairs. This is the first evidence that the structural model is effective.

We know that any simplified model will not perfectly replicate the observed data, but do the differences in OD pair-level trips sabotage the effectiveness of the baseline projection? In this case we would think no. It is worth noting that the 2014 survey ODM represents a single, noisy realization of the freight system. On the other hand, the structural model (Land Use + Gravity) is designed to capture the systematic, reproducible relationships within that data, not the idiosyncratic noise or unobserved factors present in a single survey.

Therefore, conceptually, the structural model acts as an “information filter” that filters out the random, non-systematic noise from the 2014 survey and retains only the portion of the ODM that is explainable by stable, long-term drivers like employment and geography. This “filtered” baseline” is expected to be more robust for projection than the raw survey data precisely because it is less sensitive to the specific, non-repeatable conditions of 2014.

To further examine the quality of the projected baseline ODM, we want to know if two ODMs are largely statistically indistinguishable. Therefore, after the overall structural similarity comparison, we have conducted another comparison, considering the inherent randomness of the RFMS itself. For this comparison, we have incorporated the RSE associated with the surveyed ODM, and calculated the z-score for each OD pair using the formulas below:

$$SE_{ij} = RSE_{ij} \times T_{ij}^{survey}$$

$$z_{ij} = \frac{T_{ij}^{prediction} - T_{ij}^{survey}}{SE_{ij}}$$

Where T_{ij}^{survey} stands for the trips between origin i and destination j in the 2014 RFMS, $T_{ij}^{prediction}$ stands for the projected number of trips based on the structural model, and RSE_{ij} stands for the reported relative standard error for the trip number estimation in the survey.

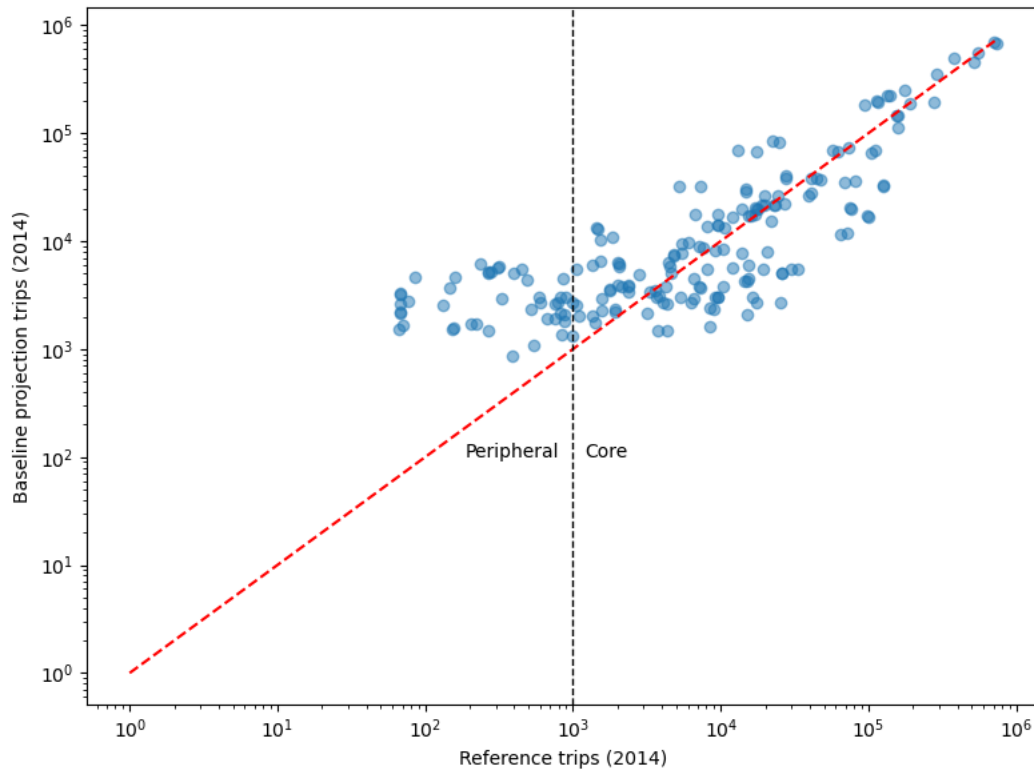
Across the full 225 by 225 ODM, 30% of OD pairs have $|z| < 1$, 50% have $|z| < 2$, and 57% have $|z| < 3$. However, 80 OD pairs ($\approx 36\%$ of the total) have less than 3 trips per day in the survey and collectively represent less than 0.3% of total trips. In these peripheral OD pairs, survey estimates are dominated by sampling variability and discreteness, and small absolute discrepancies yield large z-scores.

Restricting attention to the remaining 145 “core” OD pairs, which account for over 99.7% of total trips, model fit is substantially higher: 45%, 77% and 83% of flows lie within ± 1 , ± 2 and ± 3 of survey standard errors, and the overall similarity further goes above 0.96. From a planning perspective, the gravity-based structural model therefore reproduces the material

freight pattern to within the statistical uncertainty of the survey, while deviations are concentrated in low-volume pairs of negligible policy relevance.

Figure 26 presents the trips in two ODMs in a scattered plot.

Figure 26 - Survey vs Baseline projected ODMs in 2014



From another angle, we further compared modelled flows against the survey's RSE-based confidence intervals for each OD pair. For each cell, we computed whether the predicted flow lay above the upper bound or below the lower bound implied by the survey estimate and its relative standard error. Aggregated across the entire matrix, only approximately 6% of total trips fall in OD pairs where the model prediction lies outside these confidence bounds, while 94% of trips occur in OD pairs where modelled flows are statistically indistinguishable from the survey.

The above examination results indicate that the structural model effectively reproduces the material freight pattern to within the precision of the survey itself, and that statistically significant discrepancies are confined to a small share of total flows.

5.5. ODM inference for 2024

This calibration exercise aims to infer the most likely underlying road freight ODM at the SA3 level for 2024. By fusing the projected baseline ODM with traffic counts from all 42 permanent counters, this use case seeks to maximize data utilization to provide the best possible estimate of the underlying freight movements.

To ensure a solvable and robust optimization, we imposed the following constraints:

- **Total System Travel:** Constrained to remain close to the baseline projection (respecting the macroeconomic drivers from the land use model), with a relaxation parameter of $\pm 1\%$ to facilitate convergence.
- **Traffic Counts:** Constrained to match observed counts (respecting the ground truth), with tolerance bands to account for potential measurement noise and day-to-day variability.

Notably, at this stage, we did not strictly constrain zonal trip production and attraction (P/A) to the land use model's specific zonal predictions. This decision addresses a known limitation of using PSL routing on an approximated network: without empirical route choice variance, strict zonal P/A constraints can create an over-determined system that fails to converge on a feasible solution. Future iterations incorporating empirical route choice data (e.g., GPS telematics) would allow for stricter zonal constraints, further enhancing inference quality.

Figures 29 and 30 present the baseline projection and the calibrated 2024 SA3 ODMs. The total freight task remains consistent, shifting marginally from 922,000 trips (baseline) to 931,000 trips (calibrated). This minor adjustment reflects the calibration process aligning the economic-based projection with actual observed traffic levels on the network.

Structurally, the cosine similarity between the calibrated ODM and the baseline projection is approximately 0.88. This indicates that while the fundamental spatial pattern derived from 2014 remains relevant, there is a visually perceivable structural drift (compared to the 2014 validation). This divergence highlights the value of the calibration step in capturing recent spatial shifts not predicted by land use changes alone. The interpretation of such structural difference will be discussed in detail in the next sub-section.

It is important to note that direct validation of inferred ODMs is inherently challenging in freight contexts, as no independent “ground truth” exists for comparison (noting that the RFMS was discontinued after 2014). While other estimation frameworks exist, such as CSIRO's TraNSIT model or periodic intercept surveys conducted by the Department of State Growth, these rely on bottom-up methodologies focused on specific supply chains or road segments. Due to fundamental differences in model scope, definitions, and aggregation levels, these sources are not suitable for direct “apples-to-apples” validation of a holistic, network-wide OD matrix. Consequently, OD estimation studies typically rely on internal consistency checks (such as network fit and structural plausibility) and convergence with complementary data sources rather than direct external validation. Our approach follows this established practice while incorporating multiple validation dimensions to maximize confidence.

A key validation metric is the symmetry of flows. As shown in Figure 31, the calibrated ODM exhibits near-perfect symmetry between origin-destination and destination-origin flows, with a correlation coefficient of 0.999. This symmetry matches the characteristics of the 2014 surveyed ODM (Figure 27) and confirms that the calibration process preserves the fundamental logic of a trip-based system, where annual outbound and inbound movements tend to balance.

Figure 31 - O-D and D-O flow comparison (Calibrated ODM)

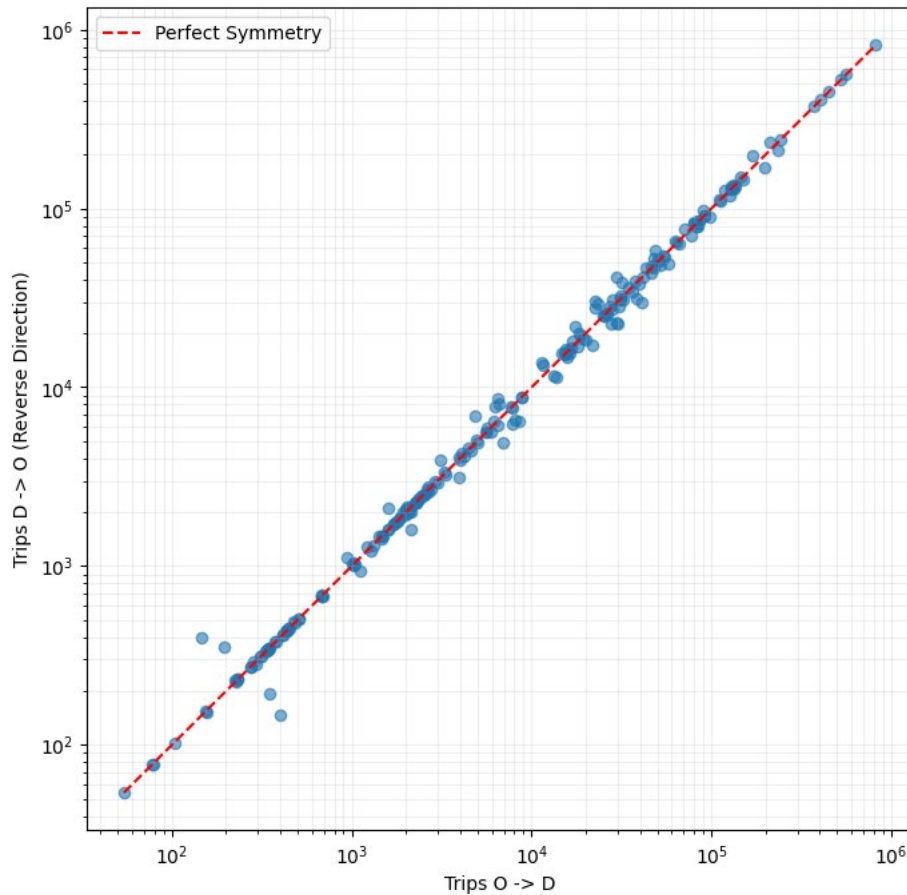



Figure 32 compares zonal trip production between the baseline projection and the calibrated ODM. While the majority of SA3 zones show strong alignment, noticeable deviations occur in specific zones (e.g., Burnie–Ulverstone and Meander Valley–West Tamar).

To assess whether these deviations represent model failure or expected variability, we benchmarked them against the RSE reported in the 2014 RFMS. Assuming the underlying relative uncertainty of zonal freight generation has remained comparable to 2014, we calculated Z-scores for the deviation between the baseline and calibrated productions. Results indicate that only two zones fall outside the ± 3 Z-score range. While this comparison is qualitative, acknowledging that uncertainty levels may have evolved, it provides indirect evidence that the calibrated ODM remains within a statistically reasonable range of the structural baseline, with outliers likely reflecting genuine localized shifts in freight intensity.

Figure 32 - Zonal trip production comparison (2024)



Lastly, recognizing the complexity of the non-linear optimization landscape and the potential for solutions to become trapped in local optima, we adopted the stability testing protocol described in our previous work (Sun et al., 2023). This involved introducing random perturbations to the baseline ODM projection and repeating the calibration process multiple times. The resulting ODMs demonstrated high consistency across all iterations, confirming the robustness of the solution presented above.

In conclusion, the 2024 calibrated ODM represents a synthesis of three information sources: the systematically reproducible structural patterns learned from 2014, the potential structural evolution over the subsequent decade, and the year-specific intensity captured by 2024 traffic counts. Multiple internal validation checks, including near-perfect trip symmetry ($r = 0.999$), zonal production consistency within survey RSE bands for most zones, and strong traffic count fit, indicate that the calibrated ODM is a statistically sound representation of 2024 freight movements. As is standard practice in OD estimation research, external validation using independent data sources (such as top-down commodity flow models or future survey benchmarks) would provide additional corroboration, and we recommend this as a priority for future research.

***Note:** The baseline ODM projections and the calibrated ODMs will be provided in the accompanying excel spreadsheets.

5.6. Freight movement survey power monitoring

This section presents experimental results testing the hypothesis that the structural knowledge derived from the 2014 Road Freight Movements Survey (RFMS) weakens over time. Determining the optimal timing for a new, capital-intensive travel survey remains a critical challenge for transport planning authorities. The integrated inference framework developed here provides a quantitative evidence base to support such decisions.

Before interpreting the results, it is necessary to clarify the data configuration choices made to ensure a rigorous longitudinal comparison.

- **Consistent Observational Density:** As detailed in the Data Input section, 42 permanent traffic counters were available in Tasmania by the end of 2024. However, only 26 of these have been consistently operational from 2015 to 2024. While utilizing all available counters in any given year maximizes data utilization for specific point-in-time estimates (as done in the previous subsection), variable observational density introduces artifacts into longitudinal analysis. A calibrated ODM is shaped by the density of its constraints; thus, a year with more counters is more tightly constrained than a year with fewer. To isolate genuine structural drift from artifacts of changing sensor density, we restricted this longitudinal analysis to the consistent set of 26 counters. This ensures that temporal variations in structural deviation metrics reflect changes in underlying flow patterns, not changes in observability.
- **Static Network Topology:** We utilized a static snapshot of the Tasmanian main road network (as of August 2025) for all calibration years. Given the mature nature of the network, where capacity expansions over the last decade have been minimal and safety-focused rather than connectivity-altering, this assumption is expected to introduce negligible assignment error. Methodologically, holding the network topology constant acts as an experimental control, ensuring that any observed divergence is attributable to the evolution of demand patterns rather than changes in network supply or routing logic.

With the above data configurations, we have a consistent “ruler” to measure the change in the structural deviation across years.

As described in the methodology section, given the hierarchical regularization, the optimization target is to find an optimal SA2 ODM that minimises the weighted sum of:

- The structural distance between the inferred SA3 ODM and the baseline SA3 ODM projection, and
- The structural distance between the inferred SA2 ODM and the inferred SA2 ODM alternative (which represents how flows should be distributed if the 2014 internal spatial structure were perfectly preserved given the calibrated SA3 totals).

In the absence of empirical knowledge of the relative importance of the SA3 and SA2 levels’ structural similarity, in this modelling exercise, we gave them equal weights. As a result, the two levels of structural distance are reported and analysed separately, since they are associated with distinct interpretation.

Figure 33 presents the evolution of SA3-level structural distance (cosine distance) from 2015 to 2024. The series of blue dots (representing the consistent 26-counter calibration) exhibits a clear, monotonic increasing trend, rising from 0.04 in 2015 to over 0.08 in 2024. This trend confirms that the explanatory power of the 2014-based structural model is progressively degrading that, over time, the calibration process is forced to apply increasingly large adjustments to the baseline projection to satisfy observed traffic counts. A linear regression (green dashed line) indicates an average annual increase in cosine distance of approximately 0.005.

Critically, the red triangle represents the structural distance for 2024 when all 42 counters are used. This value (0.12) is significantly higher than the 26-counter estimate (0.08). This discrepancy yields two vital insights for planning authorities:

- **The "Conservative Baseline".** The green trend line represents a conservative estimate of structural decay. It serves as a stable monitoring signal to forecast when a pre-determined deviation threshold will be breached, triggering the need for a new survey.
- **The Value of Sensor Density.** The higher deviation revealed by the full counter set suggests that with sparse observation (26 counters), the model may still be "over-fitting" to the baseline due to a lack of constraints. Increasing counter density forces the model to confront more traffic reality, revealing non-structural anomalies earlier. This suggests that investment in denser counting networks not only improves estimation accuracy but also accelerates the detection of divergence from land-use baselines.

Figure 33 - SA3 level structural distance over years

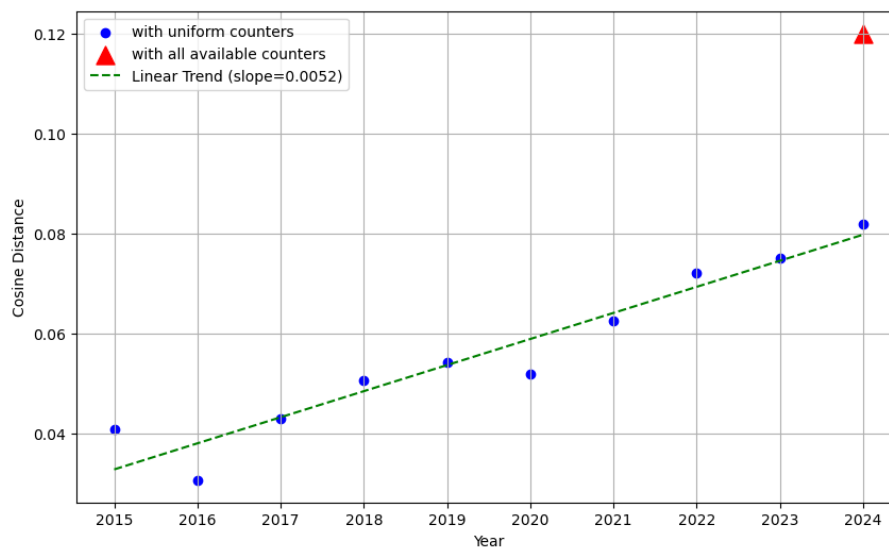



Figure 34 presents the SA2-level structural distance. Unlike the direct deviation shown at the SA3 level, this metric captures micro-spatial drift. It quantifies the tension between how flows must be distributed to satisfy local traffic counts versus how they should be distributed if the 2014 intra-regional spatial logic were intact. The SA2-level distance largely mirrors the increasing trend of the SA3 level, confirming that structural degradation is occurring simultaneously at both the strategic (inter-regional volume) and tactical (intra-regional distribution) scales.

Figure 34 - SA2 level structural distance over years



In conclusion, both SA3 and SA2 metrics demonstrate a consistent, measurable decay in the validity of the 2014 structural prior. This confirms the hypothesis that survey-derived structural knowledge has a finite lifespan, and provides BITRE with a quantifiable method to determine the optimal timing for future data collection cycles.

6. DISCUSSION

6.1. Bridging the data deficit with hybrid inference

The primary challenge addressed in this study is the "data deficit" characterizing the Australian freight landscape: the disconnect between comprehensive but outdated structural data (2014 RFMS) and current but fragmented observational data (traffic counts). Our results demonstrate that a pure "structural" approach (Gravity model alone) ignores the temporal drift of the last decade, while a pure "data-driven" approach (counts alone) is mathematically under-determined at the route level.

Such insight pronounces the value of the hybrid approach. In the 2024 ODM calibration case, calibrated ODM represents a posterior estimate that fuses these sources. The high internal consistency (O-D symmetry $r \sim 0.999$) and the preservation of macro-structural logic (SA3 level similarity > 0.8) confirm that it is possible to generate robust, updated freight matrices without a new expensive survey, provided the structural prior is explicitly modelled.

6.2. Monitoring structural decay

An equally significant contribution of this framework is the ability to quantify the "shelf life" of survey data. By tracking the structural distance between the projected baseline and the calibrated ODM using a consistent set of traffic counters, we identified a clear structural decay from 2015 to 2024.

From a policy implication perspective, this metric transforms the survey scheduling problem from a political decision into an evidence-based one. Instead of arbitrary fixed intervals (e.g., "every 5 years"), authorities can define a "structural deviation threshold" (e.g., cosine distance > 0.1). When the framework's monitoring signal breaches this threshold, it triggers the investment case for a new comprehensive survey.

This framework could also serve as an investment guidance. The comparison between the "26-counter" and "42-counter" deviation (Figure 33) highlights the value of sensor density. The conservative baseline (26 counters) provides the minimum trend, while the full sensor set reveals the full extent of divergence, given the current data availability. This suggests that investment in permanent counter networks is not just for operations, but is a strategic asset for maintaining the validity of long-term planning models.

6.3. The role of hierarchical regularization

Our methodology introduced a hierarchical regularization term which anchors SA2 flows to land-use patterns while allows SA3 flows to scale with traffic counts.

From a methodological perspective, this effectively resolves the limitations of PSL-based route choice in data-sparse environments. Without empirical GPS trajectories, route choice models struggle to discriminate between neighbouring zones. Hierarchical regularization acts as a "soft constraint," ensuring that the calibrated micro-structure at the SA2 level remains economically plausible even when the traffic count data is insufficient to uniquely identify them.

Importantly, this architecture provides future scalability and is "GPS-ready." As telematics data becomes available, the "Alternative SA2 ODM" can be updated from a static land-use derivative to an empirically observed distribution, seamlessly improving the model's fidelity without architectural redesign.

6.4. Limitations and future directions

While the pilot case study using Tasmanian data demonstrates the framework's efficacy, several limitations must be acknowledged to guide future scaling to the mainland.

- **Route Choice Proxy:** The reliance on PSL with a static road network is a necessity of the current data landscape. It captures topology but misses the operational realities such as congestion, vehicle accessibility, and turn bans. Future iterations should prioritize integrating freight vehicle specific empirical route choices with the actual freight vehicle-allowed road network.
- **External Validation:** As noted previously, the lack of a ground-truth ODM precludes direct error measurement. While internal checks (symmetry, RSE benchmarking) are positive, cross-validation against emerging top-down datasets (e.g., commodity flow models) or partial intercept surveys remains a priority.
- **The scaling work:** Tasmania is a closed system with no external road gateways. Scaling this to mainland Australia will require modelling inter-state "external" trips. The framework can accommodate this by treating state border crossings as "special generators" with high traffic count observability, effectively constraining the model at the state boundaries.
- **Potential improvement on the structural model:** The form of structural model in the pilot case study has been constrained by the number of data points. More sophisticated methods for the structural learning could be tested when scaling up to mainland Australia.

7. CONCLUSIONS

Australia's freight planning capabilities have long been constrained by a fundamental "data deficit": the reliance on the discontinued 2014 Road Freight Movements Survey (RFMS) as the sole comprehensive structural snapshot, contrasted with a growing but fragmented ecosystem of observational data (e.g., traffic counts). This report has presented an integrated inference framework designed not merely to fill this gap, but to bridge it by systematically fusing the structural legacy of the 2014 survey with the observational reality of modern traffic data.

7.1. Methodological contribution

We have developed a hierarchical inference framework that operates on two principles: **structural preservation** and **empirical calibration**. By estimating a land-use and gravity-based structural model from the 2014 RFMS, we successfully extracted the stable economic logic driving freight generation. We then demonstrated that this logic can be projected forward and calibrated against contemporary traffic counts using a single-level optimization approach. This hybrid architecture solves the "under-determination" problem inherent in count-based estimation by using the structural model to stabilize the solution, particularly at the fine spatial resolution (SA2) where count data is sparse.

7.2. Key findings from the Tasmania pilot

The application of this framework to Tasmania yielded three critical insights:

- **Robust estimation:** The calibrated 2024 ODM provides a statistically sound estimate of current freight movements. Internal validation metrics, including near-perfect O-D symmetry ($r=0.999$) and zonal production consistencies within historical error bands, confirm that the model produces behaviourally realistic flows that respect both economic fundamentals and observed traffic.
- **Quantifiable structural decay:** Through longitudinal analysis using a consistent set of 26 traffic counters, we identified a clear increase in the structural deviation between the 2014-based baseline and the calibrated reality. This confirms that the explanatory power of the 2014 survey is progressively weakening, with an average annual structural drift measurable at both the strategic (SA3) and tactical (SA2) levels.
- **The value of observational density:** We demonstrated that increasing the number of traffic counters significantly alters the estimated ODM structure, revealing deviations that sparse networks miss. This highlights that dense, permanent monitoring networks are essential not just for operations, but for maintaining the strategic accuracy of national freight models.

7.3. Strategic implications for BITRE

This framework equips BITRE with a comprehensive, three-pillar strategic toolset:

- **Operational estimation:** It provides an immediate mechanism to generate validated "best-estimate" freight matrices for non-survey years. This maximizes the return on investment of existing data assets, turning fragmented traffic counts into coherent network-wide intelligence.

- **Long-term forecasting:** By anchoring freight generation to fundamental economic drivers (employment and land use), the framework serves as a robust projection engine. It enables BITRE to model future freight scenarios (e.g., 2030, 2040) based on socio-demographic forecasts, providing the structural demand layer required for long-term infrastructure investment analysis.
- **Survey cycle optimization:** It transforms survey planning from a static schedule into an evidence-based decision. By continuously monitoring the structural degradation trend, BITRE can scientifically determine when the behavioural validity of the 2014 baseline has been exhausted, ensuring that future costly surveys are commissioned exactly when — and only when — they are needed.

7.4. Strategic outlook and future development

While the pilot confirms the framework's efficacy, scaling to a national level will require addressing the "Island Effect" by modelling inter-state external trips and integrating richer route-choice data (such as GPS telematics) to replace the static network proxies used here. Nevertheless, this study establishes a replicable, scientifically rigorous pathway to modernize Australia's freight modelling capacity, moving from reliance on discrete, historical snapshots to a continuous, data-driven understanding of the national freight task.

While this report focuses on the methodological development and validation of a road freight OD inference framework, it forms part of a broader strategic vision for a nationally integrated freight data system. The longer-term objective is to establish a sustainable data architecture in which survey data, roadside counters, and telematics sources are harmonised within a consistent inference framework to produce regular, transparent, and quality-assured estimates of freight flows across Australia.

In this vision, roadside counters provide the structural anchor for annual updates; periodic survey collections refresh network-wide OD structure; and telematics data enhance route-choice calibration and temporal resolution. Over time, this integrated framework could be operationalised through a national visualisation and analytics platform to support infrastructure planning, freight precinct development, intermodal strategy, emissions analysis, and resilience planning.

Further work is currently underway to explore the feasibility of scaling this approach nationally and embedding it within an operational data platform. The findings of the present study provide the methodological foundation for that next phase of development.

REFERENCES

- Abbas, S., Sampedro, G.A., Almadhor, A., Abisado, M., Marzougui, M., Kim, T., Alasiry, A., 2024. PODE: privacy-enhanced distributed federated learning approach for origin-destination estimation. *PeerJ Comput. Sci.* 10, e2050. <https://doi.org/10.7717/peerj-cs.2050>
- Australian aggregate freight forecasts – 2022 update, 2022. , BUREAU OF INFRASTRUCTURE AND TRANSPORT RESEARCH ECONOMICS RESEARCH REPORT. Bureau of Infrastructure and Transport Research Economics.
- Barna, T. (Ed.), 1963. *Structural Interdependence and Economic Development*. Palgrave Macmillan UK, London. <https://doi.org/10.1007/978-1-349-81634-7>
- Basar, G., Cetin, M., Nichols, A.P., 2018. Comparison of vehicle re-identification models for trucks based on axle spacing measurements. *J. Intell. Transp. Syst.* 22, 517–529. <https://doi.org/10.1080/15472450.2018.1441027>
- Basso, F., Pezoa, R., Tapia, N., Varas, M., 2022. Estimation of the Origin-Destination Matrix for Trucks That Use Highways: A Case Study in Chile. *Sustainability* 14, 2645. <https://doi.org/10.3390/su14052645>
- Behara, K.N.S., Bhaskar, A., Chung, E., 2022. Single-level approach to estimate origin-destination matrix: exploiting turning proportions and partial OD flows. *Transp. Lett.* 14, 721–732. <https://doi.org/10.1080/19427867.2021.1932182>
- Bekhor, S., Ben-Akiva, M.E., Ramming, M.S., 2006. Evaluation of choice set generation algorithms for route choice models. *Ann. Oper. Res.* 144, 235–247. <https://doi.org/10.1007/s10479-006-0009-8>
- Bell, M.G.H., 1991. The estimation of origin-destination matrices by constrained generalised least squares. *Transp. Res. Part B Methodol.* 25, 13–22. [https://doi.org/10.1016/0191-2615\(91\)90010-G](https://doi.org/10.1016/0191-2615(91)90010-G)
- Ben-Akiva, M., Bierlaire, M., 1999. Discrete Choice Methods and their Applications to Short Term Travel Decisions, in: Hall, R.W. (Ed.), *Handbook of Transportation Science, International Series in Operations Research & Management Science*. Springer US, Boston, MA, pp. 5–33. https://doi.org/10.1007/978-1-4615-5203-1_2
- BITRE, 2022. Australian aggregate freight forecasts – 2022 update (No. Research Report 154). Bureau of Infrastructure and Transport Research Economics.
- Blogg, M., Conor Semler, Hingorani, M., Troutbeck, R., 2010. Travel Time and Origin-Destination Data Collection using Bluetooth MAC Address Readers, in: *Australasian Transport Research Forum 2010 Proceedings*. Presented at the Australasian Transport Research Forum, Canberra, Australia.
- Boeing, G., 2025. Modeling and Analyzing Urban Networks and Amenities With OSMnx. *Geogr. Anal.* 57, 567–577. <https://doi.org/10.1111/gean.70009>
- Bonnel, P., Fekih, M., Smoreda, Z., 2018. Origin-Destination estimation using mobile network probe data. *Transp. Res. Procedia* 32, 69–81. <https://doi.org/10.1016/j.trpro.2018.10.013>
- Bonnel, P., Hombourger, E., Olteanu-Raimond, A.-M., Smoreda, Z., 2015. Passive Mobile Phone Dataset to Construct Origin-destination Matrix: Potentials and Limitations. *Transp. Res. Procedia* 11, 381–398. <https://doi.org/10.1016/j.trpro.2015.12.032>
- Bovy, P.H.L., 2009. On Modelling Route Choice Sets in Transportation Networks: A Synthesis. *Transp. Rev.* 29, 43–68. <https://doi.org/10.1080/01441640802078673>
- Bovy, P.H.L., Fiorenzo-Catalano, S., 2007. STOCHASTIC ROUTE CHOICE SET GENERATION: BEHAVIORAL AND PROBABILISTIC FOUNDATIONS. *Transportmetrica* 3, 173–189. <https://doi.org/10.1080/18128600708685672>
- Carvalho, L., 2011. A Bayesian Statistical Approach for Inference on Static Origin-Destination Matrices. <https://doi.org/10.48550/arXiv.1012.1047>

- Cascetta, E., 1984. Estimation of trip matrices from traffic counts and survey data: A generalized least squares estimator. *Transp. Res. Part B Methodol.* 18, 289–299. [https://doi.org/10.1016/0191-2615\(84\)90012-2](https://doi.org/10.1016/0191-2615(84)90012-2)
- Castillo, E., Menéndez, J.M., Sánchez-Cambronero, S., 2008. Predicting traffic flow using Bayesian networks. *Transp. Res. Part B Methodol.* 42, 482–509. <https://doi.org/10.1016/j.trb.2007.10.003>
- Cetin, M., Monsere, C.M., Nichols, A.P., 2011. Bayesian Models for Reidentification of Trucks Over Long Distances on the Basis of Axle Measurement Data. *J. Intell. Transp. Syst.* 15, 1–12. <https://doi.org/10.1080/15472450.2011.544571>
- Cipriani, E., Gemma, A., Mannini, L., Carrese, S., Crisalli, U., 2021. Traffic demand estimation using path information from Bluetooth data. *Transp. Res. Part C Emerg. Technol.* 133, 103443. <https://doi.org/10.1016/j.trc.2021.103443>
- Cohen, H., Horowitz, A., Pendyala, R., 2008. Forecasting Statewide Freight Toolkit (No. NCHRP Report 606). Transportation Research Board (TRB), Washington, D.C., USA.
- De Rolland, C., Bayart, C., 2025. From zonal aggregated mobile phone data to origin-destination matrices: Revisiting methodologies developed for traffic counts data. *Transp. Res. Part Policy Pract.* 202, 104715. <https://doi.org/10.1016/j.tra.2025.104715>
- Demissie, M.G., Kattan, L., 2022. Estimation of truck origin-destination flows using GPS data. *Transp. Res. Part E Logist. Transp. Rev.* 159, 102621. <https://doi.org/10.1016/j.tre.2022.102621>
- Dey, S., Winter, S., Tomko, M., 2020. Origin–Destination Flow Estimation from Link Count Data Only. *Sensors* 20, 5226. <https://doi.org/10.3390/s20185226>
- Di, X., Shi, R., Mo, Z., Fu, Y., 2023. Physics-Informed Deep Learning for Traffic State Estimation: A Survey and the Outlook. *Algorithms* 16, 305. <https://doi.org/10.3390/a16060305>
- Duanmu, J., Foytik, P., Khattak, A., Robinson, R.M., 2012. Distribution Analysis of Freight Transportation with Gravity Model and Genetic Algorithm. *Transp. Res. Rec. J. Transp. Res. Board* 2269, 1–10. <https://doi.org/10.3141/2269-01>
- Fiorenzo-Catalano, M.S., 2007. Choice set generation in multi-modal transportation networks. Delft, Netherlands.
- Fisk, C.S., 1989. Trip matrix estimation from link traffic counts: The congested network case. *Transp. Res. Part B Methodol.* 23, 331–336. [https://doi.org/10.1016/0191-2615\(89\)90009-X](https://doi.org/10.1016/0191-2615(89)90009-X)
- Freight Demand Modeling and Data Improvement Handbook, 2017. . U.S. Department of Transportation Federal Highway Administration.
- Gajda, J., Sroka, R., Burnos, P., 2020. Sensor Data Fusion in Multi-Sensor Weigh-In-Motion Systems. *Sensors* 20, 3357. <https://doi.org/10.3390/s20123357>
- Graells-Garrido, E., Opitz, D., Rowe, F., Arriagada, J., 2023. A data fusion approach with mobile phone data for updating travel survey-based mode split estimates. *Transp. Res. Part C Emerg. Technol.* 155, 104285. <https://doi.org/10.1016/j.trc.2023.104285>
- Hess, S., Quddus, M., Rieser-Schüssler, N., Daly, A., 2015. Developing advanced route choice models for heavy goods vehicles using GPS data. *Transp. Res. Part E Logist. Transp. Rev.* 77, 29–44. <https://doi.org/10.1016/j.tre.2015.01.010>
- Holguín-Veras, J., Jaller, M., Destro, L., Ban, X. (Jeff), Lawson, C., Levinson, H.S., 2011. Freight Generation, Freight Trip Generation, and Perils of Using Constant Trip Rates. *Transp. Res. Rec. J. Transp. Res. Board* 2224, 68–81. <https://doi.org/10.3141/2224-09>
- Holguín-Veras, J., Jaller, M., Sánchez-Díaz, I., Campbell, S., Lawson, C.T., 2014. Freight Generation and Freight Trip Generation Models, in: *Modelling Freight Transport*. Elsevier, pp. 43–63. <https://doi.org/10.1016/B978-0-12-410400-6.00003-3>

- Holguín-Veras, J., Thorson, E., 2000. Trip Length Distributions in Commodity-Based and Trip-Based Freight Demand Modeling: Investigation of Relationships. *Transp. Res. Rec. J. Transp. Res. Board* 1707, 37–48. <https://doi.org/10.3141/1707-05>
- Horowitz, A.J., 2010. Origin–destination table disaggregation using biproportional least squares estimation. *Transportation* 37, 689–703. <https://doi.org/10.1007/s11116-010-9273-1>
- Leontief, W.W., 1936. Quantitative Input and Output Relations in the Economic Systems of the United States. *Rev. Econ. Stat.* 18, 105. <https://doi.org/10.2307/1927837>
- Li, W., Luo, C., He, Y., Wan, Y., Du, H., 2023. Estimating Inter-Regional Freight Demand in China Based on the Input–Output Model. *Sustainability* 15, 9808. <https://doi.org/10.3390/su15129808>
- Ma, W., Qian, Z. (Sean), 2018. A Generalized Single-Level Formulation for Origin–Destination Estimation under Stochastic User Equilibrium. *Transp. Res. Rec. J. Transp. Res. Board* 2672, 58–68. <https://doi.org/10.1177/0361198118782041>
- Maher, M.J., 1983. Inferences on trip matrices from observations on link volumes: A Bayesian statistical approach. *Transp. Res. Part B Methodol.* 17, 435–447. [https://doi.org/10.1016/0191-2615\(83\)90030-9](https://doi.org/10.1016/0191-2615(83)90030-9)
- McHugh, C., Camargo, P., van Vuren, T., 2021. Freight modelling in Australia in 2021—A data availability perspective, in: *Australasian Transport Research Forum 2021 Proceedings*. Presented at the 42nd Australasian Transport Research Forum (ATRF), Brisbane, Australia.
- Michau, G., Pustelnik, N., Borgnat, P., Abry, P., Bhaskar, A., Chung, E., 2019. Combining traffic counts and Bluetooth data for link-origin-destination matrix estimation in large urban networks: The Brisbane case study. <https://doi.org/10.48550/ARXIV.1907.07495>
- Mitchell, D., Kurniawan, T., 2015. Estimating Australian Commodity Freight Movements: A Linear Programming Approach, in: *Australasian Transport Research Forum 2015 Proceedings*. Presented at the Australasian Transport Research Forum, Sydney, Australia.
- National Freight and Supply Chain Strategy 2025, 2025. . Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts.
- Ortúzar S., J. de D., Willumsen, L.G., 2024. *Modelling transport*, Fifth edition. ed. Wiley, Hoboken, NJ.
- Peng, T., Gan, M., Ou, Q., Yang, X., Wei, L., Rødal Ler, H., Yu, H., 2024. Railway cold chain freight demand forecasting with graph neural networks: A novel GraphARMA-GRU model. *Expert Syst. Appl.* 255, 124693. <https://doi.org/10.1016/j.eswa.2024.124693>
- Pitfield, D.E., 1978. Freight Distribution Model Predictions Compared: A Test of Hypotheses. *Environ. Plan. Econ. Space* 10, 813–836. <https://doi.org/10.1068/a100813>
- Pompigna, A., Mauro, R., 2020. Input/Output models for freight transport demand: a macro approach to traffic analysis for a freight corridor. *Arch. Transp.* 54, 21–42. <https://doi.org/10.5604/01.3001.0014.2729>
- Quick Response Freight Manual II, 2007. . Cambridge Systematics.
- Ramming, M.S., 2002. *Network knowledge and route choice*.
- Shen, W., Wynter, L., 2012. A new one-level convex optimization approach for estimating origin–destination demand. *Transp. Res. Part B Methodol.* 46, 1535–1555. <https://doi.org/10.1016/j.trb.2012.07.005>
- Spieß, H., 1990. A gradient approach for the OD matrix adjustment problem.
- Sun, W., Vij, A., Kaliszewski, N., 2023. A flexible and scalable single-level framework for OD matrix inference using IoT data. *Transp. Res. Part Policy Pract.* 175, 103775. <https://doi.org/10.1016/j.tra.2023.103775>
- T9 Urban freight demand modelling, 2021. . Australian Transport Assessment and Planning (ATAP) Guidelines. The Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts, Canberra, Australia.

- TAG UNIT M1.1 - Principles of Modelling and Forecasting, 2024. , Transport Analysis Guidance (TAG). Department for Transport, UK.
- TAG Unit M4 - Forecasting and Uncertainty, 2025. , Transport Analysis Guidance (TAG). Department for Transport, UK.
- Tamin, O.Z., Willumsen, L.G., 1989. Transport demand model estimation from traffic counts. *Transportation* 16. <https://doi.org/10.1007/BF00223044>
- Uddin, M., Liu, Y., Lim, H., 2023. Exploring the Effects of Population and Employment Characteristics on Truck Flows: An Analysis of NextGen NHTS Origin-Destination Data, in: *International Conference on Transportation and Development 2023*. Presented at the International Conference on Transportation and Development 2023, American Society of Civil Engineers, Austin, Texas, pp. 503–513. <https://doi.org/10.1061/9780784484883.044>
- Van Zuylen, H.J., Willumsen, L.G., 1980. The most likely trip matrix estimated from traffic counts. *Transp. Res. Part B Methodol.* 14, 281–293. [https://doi.org/10.1016/0191-2615\(80\)90008-9](https://doi.org/10.1016/0191-2615(80)90008-9)
- Viswanathan, K., Beagan, D., Mysore, V., Srinivasan, N., 2008. Disaggregating Freight Analysis Framework Version 2 Data for Florida: Methodology and Results. *Transp. Res. Rec. J. Transp. Res. Board* 2049, 167–175. <https://doi.org/10.3141/2049-20>
- Wang, F., Chen, C., 2018. On data processing required to derive mobility patterns from passively-generated mobile phone data. *Transp. Res. Part C Emerg. Technol.* 87, 58–74. <https://doi.org/10.1016/j.trc.2017.12.003>
- Yang, H., 1995. Heuristic algorithms for the bilevel origin-destination matrix estimation problem. *Transp. Res. Part B Methodol.* 29, 231–242. [https://doi.org/10.1016/0191-2615\(95\)00003-V](https://doi.org/10.1016/0191-2615(95)00003-V)
- Yang, H., Sasaki, T., Iida, Y., Asakura, Y., 1992. Estimation of origin-destination matrices from link traffic counts on congested networks. *Transp. Res. Part B Methodol.* 26, 417–434. [https://doi.org/10.1016/0191-2615\(92\)90008-K](https://doi.org/10.1016/0191-2615(92)90008-K)
- Yang, Y., Jia, B., Yan, X.-Y., Chen, Y., Song, D., Zhi, D., Wang, Y., Gao, Z., 2023. Estimating intercity heavy truck mobility flows using the deep gravity framework. *Transp. Res. Part E Logist. Transp. Rev.* 179, 103320. <https://doi.org/10.1016/j.tre.2023.103320>
- Yang, Y., Jia, B., Yan, X.-Y., Jiang, R., Ji, H., Gao, Z., 2021. Identifying intracity freight trip ends from heavy truck GPS trajectories. <https://doi.org/10.48550/arXiv.2106.09881>
- Yao, Y., Ogden, R.T., Zeng, C., Chen, Q., 2021. Bivariate Hierarchical Bayesian Model for Combining Summary Measures and their Uncertainties from Multiple Sources. <https://doi.org/10.48550/arXiv.2109.07560>
- Zanjani, A.B., Pinjari, A.R., Kamali, M., Thakur, A., Short, J., Mysore, V., Tabatabaee, S.F., 2015. Estimation of Statewide Origin–Destination Truck Flows from Large Streams of GPS Data: Application for Florida Statewide Model. *Transp. Res. Rec. J. Transp. Res. Board* 2494, 87–96. <https://doi.org/10.3141/2494-10>