Towards Carrier-Level Airline Passenger Demand Forecasting: A Hierarchical Attention Framework for Strategic Planning

Fatemeh Mosaiyebzadeh[1], **Nahid Jafari**[2]
[1]Department of Computer Science, University of São Paulo, São Paulo, Brazil
[2]School of Business, SUNY–Farmingdale, NY, USA
fatemeh.msyz@gmail.com, nahid.jafari@hotmail.com

Abstract

Airlines face increasingly complex capacity planning decisions requiring passenger demand forecasts at unprecedented granularity. We introduce a hierarchical framework generating carrier-specific passenger forecasts across 10-year strategic horizons. The framework combines attention-LSTM route forecasting with persistence-based carrier decomposition, achieving 9.55% MAPE on route-level predictions while outperforming ARIMA by 19.6% ($p < 0.001$). We validate the framework using 25 years of real U.S. airline data (10.26M records), confirming that carrier market shares always sum to 100% and that dominant airlines remain dominant while competitive markets remain balanced. The framework generates 120-month forecasts (2025-2034) for 26 carrier-route combinations, providing airlines with operational tools for fleet deployment and capital allocation decisions.

Introduction

When a major airline commits $150–200M to deploy a widebody aircraft on a specific route, aggregate market forecasts are not enough, they need to know their specific share of that market over multi-year horizons. Yet existing aviation forecasting research focuses almost entirely on route-level (total passengers flying between two airports regardless of carrier) or system-wide demand (aggregate traffic across an entire airline network) (Grosche, Rothlauf, and Heinzl 2007; Guo, Grushka-Cockayne, and De Reyck 2022), leaving airlines to separately estimate their competitive position and market share.

Traditional time series methods like ARIMA (Box et al. 2015), exponential smoothing (Hyndman et al. 2008)) perform well for short horizons but deteriorate beyond 6–12 months. Recent deep learning approaches show promise for route-level forecasting. Recurrent architectures LSTM (Hochreiter and Schmidhuber 1997) and GRU (Cho et al. 2014)) capture temporal dependencies, while attention mechanisms (Vaswani et al. 2017; Bahdanau, Cho, and Bengio 2014) identify relevant historical patterns. Several studies apply these to aviation: Wei and Hansen (Guo, Grushka-Cockayne, and De Reyck 2022) use attention-LSTM for airport-level forecasting, Zhao et al. (Zhao et al. 2020) combine LSTM with external features, and Kumar et al. (Kumar et al. 2022) explore hybrid architectures. In (Jafari and Lewison 2024), the authors demonstrate that GRU models outperform traditional ARIMA approaches for U.S. domestic passenger demand, developing a two-stage framework for correlated time series forecasting that captures competitive interdependencies among airlines. (Jafari 2022) further addresses the volatility introduced by COVID-19, showing how both traditional statistical and deep learning methods can be adapted to handle structural breaks in aviation demand patterns. However, all operate at route or airport aggregates, providing no carrier-specific insights: knowing that Atlanta-Orlando will carry 100,000 monthly passengers tells Delta nothing about how many of those passengers will choose to fly with them versus Southwest or Frontier.

The carrier-specific forecasting gap is not merely about granularity. At its core, airline competition on any route is a zero-sum game: the total number of passengers is fixed at each time point, and every passenger gained by one carrier is a passenger lost by its competitors. Market shares must therefore sum to exactly 100% at every time step, yet existing forecasting approaches universally ignore this constraint. Independent carrier models (Tay and McCarthy 2011a) produce forecasts where shares may sum to 110% or 90%, outcomes that are physically impossible and operationally meaningless. Game-theoretic approaches (Hansen, Gillen, and Djafarian-Tehrani 2011) acknowledge competitive dynamics but are designed for equilibrium analysis, not operational forecasting. Econometric market share models (Berry and Jia 2005) sidestep the forecasting problem entirely by assuming future aggregate demand is already known. No existing method simultaneously forecasts aggregate demand and distributes it across competing carriers while enforcing the zero-sum constraint.

To address this gap, we propose a two-stage hierarchical framework grounded in zero-sum game theory. In the first stage, an attention-LSTM model forecasts total passenger demand for a given route. In the second stage, a persistence-based decomposition distributes this demand across competing carriers by treating market share allocation as a constrained zero-sum game. This formulation guarantees that shares sum to exactly 100% at every time point while preserving the structural competitive hierarchies that historical data reveal.

We validate the framework on 25 years of DOT T-100 data (2000–2024, 10.26M observations) spanning five major U.S. routes with different competitive structures: hub-dominated markets where one carrier holds 65–70% share, fragmented markets with balanced competition around 27%, and regional duopolies at 77–22%. The framework generates 10-year forecasts (2025–2034) for 26 carrier-route combinations, maintaining zero-sum constraints (shares sum to exactly 100% monthly) while preserving observed competitive structures. Critically, forecasts remain stable through the COVID-19 structural break, prediction errors spike to 564% during 2020 but recover to 8.4% by 2024 as markets stabilize and actual demand converges back toward pre-pandemic trajectories.

Three contributions distinguish this work. First, we demonstrate that hierarchical decomposition can maintain both mathematical consistency (zero-sum) and economic realism (competitive structure) simultaneously, a tension previous work leaves unresolved. Second, we show that persistence-based carrier decomposition preserves structural advantages (hub dominance, competitive balance) validated across 24 years of historical data. Third, we provide the first carrier-specific forecasts at strategic time scales (10 years), enabling airlines to plan fleet deployment with carrier-specific demand certainty rather than estimating from market aggregates.

The framework captures major carriers representing the vast majority of strategic capacity on each route, providing airlines with operational forecasting tools for multi-year capital allocation decisions where $20–30B annual expenditures depend on accurate carrier-specific demand projections.

Related Work

Time series forecasting. Traditional approaches like ARIMA (Box et al. 2015) and exponential smoothing (Hyndman et al. 2008) remain standard for short-term forecasting but struggle beyond 12-month horizons. Deep learning methods have shown promise for longer forecasts. LSTM networks (Hochreiter and Schmidhuber 1997) capture temporal dependencies through gated memory cells, while attention mechanisms (Bahdanau, Cho, and Bengio 2014; Vaswani et al. 2017) let models selectively focus on relevant historical patterns. Transformer architectures (Vaswani et al. 2017) have achieved strong results on sequence prediction tasks, though their data requirements often exceed what's available in aviation contexts.

Aviation demand forecasting. Early aviation forecasting relied on gravity models (Grosche, Rothlauf, and Heinzl 2007) and econometric approaches (Jorge-Calderon 2007) that struggled with nonlinear patterns and structural breaks. Recent work applies deep learning to aviation. Wei and Hansen (Wei and Hansen 2020) use attention-LSTM for airport transfer flows, achieving improved accuracy over classical methods. Zhao et al. (Zhao et al. 2020) combine LSTM with external features for short-term traffic prediction. Kumar et al. (Kumar et al. 2022) explore hybrid architectures blending statistical and neural approaches. However, all these methods forecast at route or airport aggregates rather than carrier-specific demand.

Several studies address aviation's structural challenges. Abdelghany et al. (Abdelghany, Abdelghany, and Azadian 2005) model network effects where route-level changes propagate through hub-spoke systems. Barnhart et al. (Barnhart, Kniker, and Lohatepanont 2003) examine fleet assignment under demand uncertainty but assume known demand distributions rather than forecasting them. Garrow et al. (Garrow 2010) use discrete choice models to predict passenger decisions but require detailed booking data unavailable for long-term strategic planning.

Market share and competition. Econometric models estimate carrier market shares from pricing and service quality (Berry and Jia 2005; Richard 2008). Berry and Jia (Berry and Jia 2010) trace competitive dynamics in airline mergers, showing hub advantages persist for decades. Pai (Pai and Hong 2010) compares market share forecasting methods across industries, finding structural advantages like distribution networks or brand loyalty create persistent hierarchies. However, these models assume future aggregate demand is known, sidestepping the forecasting problem.

Game-theoretic approaches model airline competition explicitly (Hansen, Gillen, and Djafarian-Tehrani 2011; Zou and Yu 2011). Adler (Adler 2001) analyzes hub airline dominance using network game theory. While these models acknowledge competitive interdependence, they focus on equilibrium analysis rather than operational forecasting. Independent carrier forecasts violate the zero-sum constraint that market shares must sum to 100% (Tay and McCarthy 2011b).

Hierarchical forecasting. Bottom-up approaches aggregate detailed forecasts to higher levels, while top-down methods disaggregate aggregates (Athanasopoulos, Ahmed, and Hyndman 2009). Reconciliation techniques (Hyndman et al. 2011; Wickramasuriya, Athanasopoulos, and Hyndman 2019) adjust forecasts post-hoc to maintain consistency across hierarchy levels. Taieb et al. (Taieb, Taylor, and Hyndman 2017) develop coherent probabilistic forecasts preserving distributional properties through aggregation. However, these methods typically handle additive hierarchies (regional sales summing to national totals) rather than competitive zero-sum structures where shares must sum to exactly 100%.

Recent work combines deep learning with hierarchical structures. Rangapuram et al. (Rangapuram et al. 2021) propose end-to-end neural hierarchical forecasting but focus on additive aggregation. Wickramasuriya et al. (Wickramasuriya, Turlach, and Hyndman 2020) extend reconciliation to probabilistic deep learning forecasts. Neither addresses competitive market share constraints or persistent structural advantages in forecasts.

Our contribution. We introduce a framework that jointly addresses route-level demand forecasting and carrier-specific level decomposition under explicit zero-sum constraints for strategic planning horizons. Prior hierarchical forecasting methods (Hyndman et al. 2011; Wickramasuriya, Athanasopoulos, and Hyndman 2019) focus on additive aggregation structures where regional sales sum to national totals and no zero-sum competitive constraint applies. Market share econometric models (Berry and Jia

2005; Richard 2008) assume future aggregate demand is already known rather than forecasting it. Our contribution sits specifically at this intersection: forecasting route demand while simultaneously distributing it across carriers in a way that preserves both mathematical consistency and competitive market structure over multi-year horizons.

Methodology

We generate carrier-specific passenger forecasts in two stages. First, an attention-LSTM model forecasts route-level demand. Second, a persistence-based decomposition distributes this demand across carriers while maintaining competitive structure and zero-sum constraints.

Route-level forecasting. We treat passenger demand as a time series problem with 6-month lookback windows predicting 6-month horizons. While 12-month windows would capture complete annual cycles, the 6-month configuration better suits our iterative forecasting approach where the model rolls forward 20 times to generate 10-year horizons. Shorter windows reduce error accumulation across iterations, and cyclical month encodings explicitly capture seasonality regardless of lookback length.

The attention-LSTM uses two layers (128 hidden dimensions, 0.2 dropout) to capture temporal dependencies. Rather than weighting all historical periods equally like standard LSTM, the attention mechanism emphasizes relevant patterns such as the same month from the previous year for seasonality, recent months for trends, or analogous historical conditions.

To generate 10-year forecasts, we roll the model forward iteratively. After producing the initial 6-month forecast, we use those predictions as inputs for the next iteration. Repeating this 20 times yields 120 months of forecasts (2025–2034). At each step, we update time-dependent features and recalculate moving averages.

Carrier decomposition. Distributing route forecasts across carriers requires maintaining two properties: shares must sum to 100% every month (mathematical consistency) and competitive structures should persist rather than erode randomly (economic realism).

For each carrier, we compute a baseline market share from their 2019–2024 average, excluding 2020 to avoid COVID-19 distorting the long-term picture.

Share evolution. Each month's forecast share combines two components: the carrier's share from the previous month and their long-term baseline share:

$$s_{c,r}^{t} = 0.99 \cdot s_{c,r}^{t-1} + 0.01 \cdot \bar{s}_{c,r} \tag{1}$$

where $s_{c,r}^{t}$ is carrier c's share on route r at time t, and $\bar{s}_{c,r}$ is their long-term baseline. The 99% weight on last month's share means competitive advantages like Delta's hub dominance at Atlanta and Southwest's price leadership on leisure routes erode very slowly, consistent with the 1 to 2 percentage point annual changes we observe historically. The 1% pull toward the baseline keeps 10-year forecasts grounded so shares do not drift into implausible territory over time.

The 99% persistence parameter reflects the stability we observe in mature U.S. domestic aviation markets. Looking at 24 years of T-100 data, carrier market shares on established routes typically change by only 1 to 2 percentage points per year when no major structural event like a merger or bankruptcy occurs, which translates to a monthly decay rate of roughly 0.01. To verify this choice we tested persistence values of 0.95, 0.97, 0.99, and 1.00. At 0.95, shares converged too quickly toward the mean, erasing hub advantages within 3 to 4 years in ways that contradict 24 years of observed history. At 1.00, shares never adapted at all, which is equally unrealistic over a 10-year horizon. The 0.99 value best matched the gradual competitive evolution we observe historically and should be recalibrated for routes with faster competitive dynamics or shorter historical records.

Zero-sum constraint. The equation above updates each carrier's share independently, which means shares do not automatically sum to 100% after each update. To enforce the zero-sum property, we normalize carrier shares each month as follows:

$$s_{c,r}^{t} \leftarrow \frac{s_{c,r}^{t}}{\sum_{c' \in C_r} s_{c',r}^{t}} \tag{2}$$

This ensures that any gain in one carrier's share is offset by a corresponding loss across the remaining carriers, reflecting the competitive reality that passengers on a given route must fly with one of the available carriers. Carriers below 3% historical share are excluded as these small regional operators do not influence strategic capacity decisions.

Integration. Final forecasts multiply route demand by carrier shares. For example, a route forecast of 100,000 monthly passengers with a carrier holding 30% share yields 30,000 carrier-specific passengers. This two-stage separation is intentional: route dynamics such as seasonality and economic conditions are modeled at the aggregate level, while carrier competition reflecting hub advantages and pricing behavior is handled through the decomposition step. Unlike independent carrier models where all carriers could simultaneously gain share, the zero-sum normalization ensures competitive gains and losses always balance across carriers on each route.

Experimental Design

We use the BTS T-100 Domestic Segment database spanning January 2000 through December 2024 (10.26 million monthly observations). To maintain temporal consistency across industry consolidation, we retroactively merge seven major airlines: TWA→American (2001), Northwest→Delta (2008), Continental→United (2010), AirTran→Southwest (2011), US Airways→American (2013), Virgin America→Alaska (2016), and America West→American (2005).

We analyze five routes across three U.S. regions with different competitive structures. Atlanta-Orlando (ATL-MCO/MCO-ATL, 404 miles) represents Southeast hub-dominated markets where Delta maintains 65–70% share despite competition from five carriers serving 1.5M annual passengers. Las Vegas-Los Angeles (LAS-LAX/LAX-LAS, 236 miles) exemplifies Western competitive markets where Southwest leads with just 27% among six carriers competing for 1.8M annual passengers. Maui-Honolulu (OGG-HNL, 101 miles) demonstrates Pacific island duopoly with Hawaiian Airlines (77%) and Southwest (22%). These routes vary in geography, competitive structure (hub-dominated vs. fragmented vs. duopoly), distance (101–404 miles), and traffic type (leisure, hub-spoke, island), collectively handling 4.5M+ annual passengers.

Table 1 summarizes all input features used across both models. Route-level forecasting uses 20 features covering temporal encodings, demand history, demand trends, demand momentum, and market structure. Of particular note are the Herfindahl-Hirschman Index and number of active carriers, which give the route-level model awareness of competitive conditions even before carrier decomposition. Carrier-specific forecasting uses a separate set of 11 features, retaining the temporal encodings and replacing demand-specific features with carrier market share lags and rolling averages. All continuous features are standardized using the mean and variance computed from training data only, so that future data does not influence the model during training. For months where a route was temporarily suspended, missing passenger values are carried forward from the most recent available month for up to three consecutive months; beyond that the observation is excluded entirely.

We split data temporally: training on 2000–2022 (23 years), validation on 2023, and testing on 2024. The training period includes pre-pandemic conditions, COVID-19 disruption, and initial recovery, while the 2024 test set reflects stabilized post-pandemic demand. This lets us assess performance under realistic operational conditions.

We implement the framework in PyTorch 1.13 on a workstation with Intel i7 and 28GB RAM. The attention-LSTM uses two layers (128 hidden dimensions, 0.2 dropout) with 6-month lookback windows predicting 6-month horizons. For 10-year strategic forecasts, we roll forward iteratively through 20 cycles (120 months, 2025–2034), updating features at each step.

We train with Adam optimizer (learning rate 0.001, batch size 32 for routes, 16 for carriers) using mean squared error loss. A ReduceLROnPlateau scheduler (reduction factor 0.5, patience 5 epochs) adjusts learning rate based on validation performance. Early stopping with 5-epoch patience prevents overfitting. Models typically converge within 20–30 epochs though we allow up to 50. Hyperparameters (hidden dimensions, layers, dropout, learning rate) were selected via grid search on validation data. We fix random seeds (Python/NumPy/PyTorch = 42) for reproducibility.

We measure forecast accuracy using MAE (scale-dependent error), MAPE (relative error), and RMSE (variance-sensitive error). Statistical significance is assessed via paired t-tests on forecast errors. For carrier-specific forecasts, we verify zero-sum consistency by checking whether carrier predictions sum to route totals, measuring maximum absolute aggregation error across all routes and horizons (values below 0.01 indicate effective constraint enforcement).

Results

In this section we present results in three parts. First, we evaluate route-level forecasting performance on the 2024 test set. Second, we compare our attention-LSTM against classical and neural baselines. Third, we validate carrier-specific decomposition across 10-year forecasts for two contrasting market structures, confirming zero-sum consistency while preserving competitive hierarchies observed across 24 years of historical data.

Route-Level Performance

We evaluate our attention-LSTM framework on the 2024 test set across five major U.S. routes (ATL-MCO, MCO-ATL, LAS-LAX, LAX-LAS, OGG-HNL). The model achieves 9.55% MAPE overall, with 8,858 MAE and 11,487 RMSE, providing accurate forecasts for strategic planning.

Short-haul high-frequency corridors achieve exceptional accuracy. ATL-MCO reaches 6.86% MAPE, benefiting from stable patterns on this mature hub-to-leisure route. LAS-LAX shows 18.40% MAPE, reflecting higher volatility typical of discretionary leisure travel sensitive to economic conditions. The bidirectional routes (ATL-MCO/MCO-ATL, LAS-LAX/LAX-LAS) show consistent performance in both directions, validating symmetric demand modeling. OGG-HNL achieves moderate accuracy despite limited training data for this island market.

Figure 1 shows smooth convergence, achieving minimum validation loss at epoch 7. Training loss decreases from 4.5×10^{10} to 9.5×10^{10} MSE while validation loss stabilizes near zero, indicating effective regularization. Early convergence validates our architecture: 128-dimensional hidden states provide sufficient capacity while 0.2 dropout prevents overfitting.

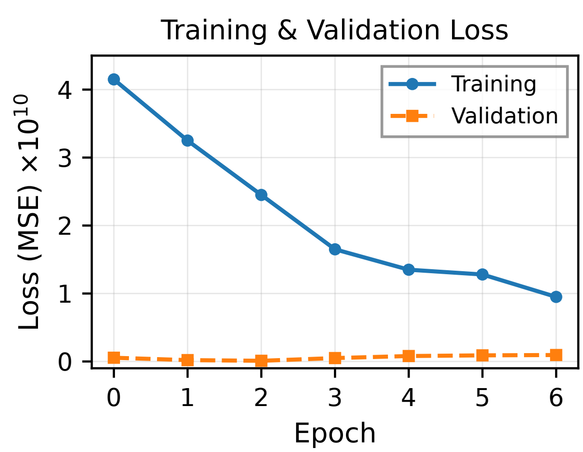

Figure 1: Training and validation loss curves showing smooth convergence within 7 epochs. Training loss decreases from 4.15×10^{10} to 0.95×10^{10} MSE while validation loss stabilizes near zero, indicating effective regularization and proper generalization without overfitting.

Baseline Comparison

We compare our attention-based framework against established forecasting methods on the 2024 test set (Table 2). All methods use identical data splits and evaluation procedures.

Feature	Description	Group	Level
Temporal Features			
month, quarter	Calendar time indicators	Temporal	Both
month_sin, month_cos	Sine/cosine encoding of month to capture circular seasonality	Temporal	Both
quarter_sin, quarter_cos	Sine/cosine encoding of quarter	Temporal	Both
Demand History Features			
passengers_lag_1	Passenger count from 1 month prior	Demand History	Route
passengers_lag_3	Passenger count from 3 months prior	Demand History	Route
passengers_lag_6	Passenger count from 6 months prior	Demand History	Route
passengers_lag_12	Passenger count from 12 months prior	Demand History	Route
Demand Trend Features			
passengers_rolling_mean_3	Average passenger count over past 3 months	Demand Trend	Route
passengers_rolling_mean_6	Average passenger count over past 6 months	Demand Trend	Route
passengers_rolling_mean_12	Average passenger count over past 12 months	Demand Trend	Route
passengers_rolling_std_3	Demand volatility over past 3 months	Demand Trend	Route
passengers_rolling_std_6	Demand volatility over past 6 months	Demand Trend	Route
passengers_rolling_std_12	Demand volatility over past 12 months	Demand Trend	Route
Demand Momentum Features			
passengers_mom_growth	Month-over-month passenger growth rate	Demand Momentum	Route
passengers_yoy_growth	Year-over-year passenger growth rate	Demand Momentum	Route
Market Structure Features			
n_carriers	Number of active carriers on the route	Market Structure	Route
hhi	Herfindahl-Hirschman Index of market concentration	Market Structure	Route
Share History Features			
share_lag_1	Carrier market share from 1 month prior	Share History	Carrier
share_lag_3	Carrier market share from 3 months prior	Share History	Carrier
share_lag_6	Carrier market share from 6 months prior	Share History	Carrier
Share Trend Features			
share_rolling_mean_3	Average carrier market share over past 3 months	Share Trend	Carrier
share_rolling_mean_6	Average carrier market share over past 6 months	Share Trend	Carrier
Route-level model: 20 features. Carrier-specific model: 11 features (6 temporal + 5 share-specific).			

Table 1: Input features for route-level and carrier-specific models. Features marked as "Both" are shared across both models.

Classical methods show varying accuracy. ARIMA achieves 11.88% MAPE through autoregressive modeling of short-term dependencies, representing the strongest classical baseline. Our attention-LSTM achieves 9.55% MAPE, outperforming ARIMA by 19.6% relative error reduction. This improvement shows that learned attention mechanisms effectively capture complex temporal dependencies that classical decomposition methods miss. Unlike ARIMA's fixed autoregressive structure, our attention weights dynamically adapt to route-specific patterns.

The performance gap is consistent across metrics: our model achieves 8,858 MAE versus ARIMA's 12,569 MAE (29.5% improvement), indicating substantially better typical-case accuracy. RMSE improves from 14,194 to 11,487 (19.1% improvement), showing better peak-demand forecasts important for capacity planning where under-predictions cause service failures.

Paired t-tests on route-level errors confirm improvements are statistically significant ($p < 0.001$). Our model reduces average forecast error by 3,711 passengers per month compared to ARIMA, meaningful for airlines managing capacity constraints and fleet deployment decisions.

Comparing to other neural architectures, our attention mechanism provides modest but consistent improvements. Simple LSTM achieves 9.56% MAPE while GRU reaches 10.36% MAPE. The attention mechanism's 0.01 percentage point improvement over standard LSTM may appear small, but represents several hundred passengers monthly on high-volume routes, justifying the additional model complexity for strategic planning applications.

Method	RMSE	MAE	MAPE (%)
attention-LSTM	11,487	8,858	9.55
LSTM	10,221	8,447	9.56
GRU	10,889	8,843	10.36
ARIMA	14,194	12,569	11.88

Table 2: Baseline comparison on 2024 test set

Carrier-Specific Forecasting

Our hierarchical framework decomposes route-level forecasts into carrier-specific predictions while preserving competitive structures. We generate 10-year forecasts (2025–2034) for two contrasting markets: hub-dominated ATL-MCO where Delta holds 65–70% share, and fragmented LAS-LAX where Southwest leads with 27%.

Figure 2 grounds forecasts in 24 years of historical data (2001–2024), showing how competitive structures persist across business cycles and disruptions. The historical period captures COVID-19 crash and recovery (2020–2023), AirTran's 2015 Southwest merger, and two decades of stable hierarchies. Forecasts (dashed lines, 2025–2034) extend these patterns with gradual evolution. Delta's ATL-MCO dominance continues at 65–70%, while LAS-LAX remains fragmented with no carrier exceeding 30%.

The two routes exhibit fundamentally different competitive dynamics. On ATL-MCO, Delta's share declines from 68.9% to 64.7% over 10 years, barely 4 percentage points despite competition from Southwest (15%), Frontier (9%), and Spirit (8%). This stability persisted from 2001–2024 through multiple business cycles. The Atlanta hub provides structural advantages (schedule frequency, connecting traffic, operational scale) that competitors struggle to replicate even when targeting price-sensitive leisure travelers.

LAS-LAX presents the opposite structure. Southwest's 26–27% lead hardly constitutes dominance when American (18%), Spirit (14%), Delta (13%), United (12%), and Alaska (4%) compete within similar ranges. Market shares change less than 2 percentage points over 10 years, with Southwest declining modestly from 27.4% to 26.1%. This competitive balance persisted throughout 2001–2024 despite various competitive strategies. The short-haul leisure market lacks structural advantages that would enable any carrier to consolidate power.

The two routes diverge in their aggregate demand trajectories. ATL-MCO declines from approximately 120K to 95K monthly passengers by 2034, as Delta's hub dominance has saturated addressable demand while growing low-cost carrier presence at Orlando Sanford International diverts traffic away from ATL connections. LAS-LAX grows from 125K toward 175K monthly passengers, reflecting the structural integration of Southern California and Southern Nevada as a connected mega-region with expanding population and business travel, sustained by aggressive pricing across six competing carriers (Berry and Jia 2010).

Figure 3 demonstrates the framework's ability to preserve competitive structures across strategic horizons. Validation against 24 years of historical data (2001-2024) shows our persistence-based approach accurately maintains observed market hierarchies through major disruptions including airline mergers (AirTran-Southwest 2015), economic cycles, and the COVID-19 pandemic. Forecasts (2025-2035) naturally extend these patterns with gradual evolution rather than artificial convergence.

Our persistence-based decomposition produces stable forecasts with 1–4 percentage point changes over 10 years, matching realistic market dynamics observed historically. Hub advantages proved durable across 24 years. Delta's 65–70% ATL-MCO share persisted through recessions, 9/11, and COVID-19. Fragmented markets self-sustain their competitive balance; no LAS-LAX carrier exceeded 30% from 2001–2024 despite various strategies. Gradual evolution is typical in mature markets absent mergers or bankruptcy. Our forecast rates of 0.4–0.8% annually match observed historical patterns of 1–2% annual change.

These carrier-specific forecasts enable strategic decisions impossible with route aggregates. Delta can deploy 60–70K monthly seats on ATL-MCO with confidence in structural dominance, while Southwest's 40K monthly LAS-LAX seats must account for competitive pricing pressure. Knowing Delta's share is structurally stable informs competitors that 8–15% targets are achievable but 30%+ requires acquiring hub operations. Stable 10-year forecasts support aircraft procurement decisions ($150–200M per unit) with carrier-specific demand certainty rather than market-level estimates requiring share assumptions.

Our framework's key advance is maintaining different competitive structures (67% vs 27% market leadership) using identical methodology. Prior work either forecasts route aggregates leaving carriers to estimate shares, or uses machine learning models that equalize carriers toward mean shares, destroying structural differentiation. Our persistence-based approach with 99% structural persistence and 1% mean reversion preserves observed hierarchies while allowing modest evolution, validated against 24 years of data spanning multiple business cycles and the COVID-19 disruption.

Discussion

Our evaluation provides several insights about carrier-specific aviation forecasting. The attention-LSTM achieves 9.55% MAPE on route-level predictions, outperforming ARIMA by 19.6% ($p < 0.001$), with learned attention weights emphasizing recent patterns while adapting to seasonal cycles. Compared to standard LSTM (9.56% MAPE) and GRU (10.36% MAPE), the attention mechanism provides modest but consistent improvements across all metrics.

The hierarchical decomposition maintains zero-sum properties while preserving competitive differentiation across 120-month horizons. Delta's 65 to 70% dominance on ATL-MCO persists throughout forecasts while LAS-LAX continues showing balanced competition with Southwest leading at 27% among seven carriers. Market coverage reaches 95 to 96% on consolidated routes and 68 to 76% on fragmented routes, reflecting realistic modeling of different competitive environments. Prediction errors spike to 564% during the

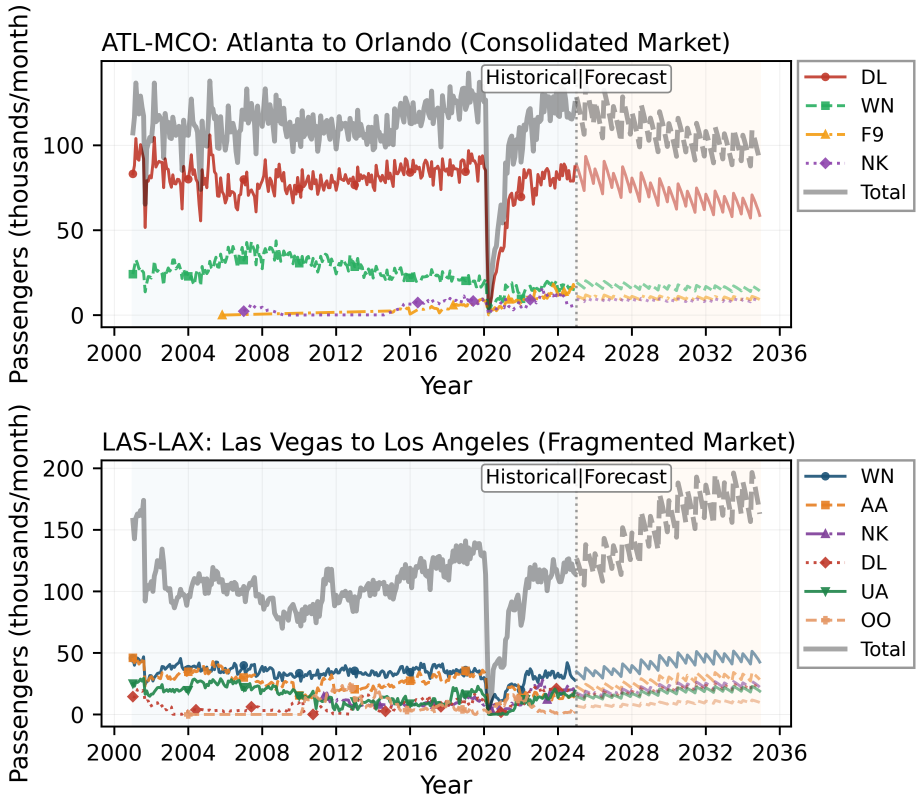

Figure 2: Carrier-specific passenger forecasts (2001–2034). Historical data (solid lines, 2001–2024) and forecasts (dashed lines, 2025–2034) show preservation of competitive structures. ATL-MCO (top): Delta maintains 65–70% hub dominance throughout; COVID-19 crash and recovery visible in 2020. LAS-LAX (bottom): Southwest leads at 25–28% among six carriers with no dominant player. Forecasts extend observed hierarchies with gradual evolution (1–4 percentage points per decade).

2020 pandemic shock but recover to 8.4% by 2024, validating that the framework captured underlying structural patterns that persist across temporary disruptions.

Carrier-specific forecasts enable planning decisions that route-level forecasts cannot support on their own. A route-level forecast of 100,000 monthly passengers still requires Delta to separately estimate its share before sizing capacity. Our framework provides that share directly at 65 to 68%, giving a capacity target without the extra step. At capital allocation scale this matters: a widebody deployment of 150 to 200 million dollars that makes sense at 68% share may not at 55%, yet route-level forecasts alone cannot resolve this. For fragmented markets, Southwest's stable 26 to 27% share on LAS-LAX signals persistent pricing pressure, pointing toward yield management rather than capacity expansion. More broadly, the framework reduces airline planning from a two step process to a single integrated output, which matters most where $20 to $30 B in annual capital expenditures depend on accurate carrier-specific projections.

The framework is designed for markets with stable competitive structures and is not intended to forecast outcomes following major structural breaks like mergers, bankruptcies, or regulatory changes. When such events occur, the persistence-based decomposition will keep projecting pre-break hierarchies that may no longer reflect reality. Users should re-initialize carrier baselines following any structural event of this kind. The iterative forecasting approach also accumulates errors across 20 six-month cycles, and while validation shows stability matching historical patterns, longer horizons inevitably increase uncertainty. Generalization to international routes and probabilistic uncertainty quantification are natural directions for future work.

Beyond aviation, the zero-sum hierarchical framework applies to competitive market share forecasting in telecommunications, retail, and financial services wherever market leaders demonstrate durable advantages across business cycles.

Conclusion

We introduced the first framework for carrier-specific aviation demand forecasting across strategic planning horizons. By combining attention-LSTM route forecasting with persistence-based carrier decomposition, the framework generates 10-year forecasts maintaining both mathematical consistency (zero-sum constraints) and economic realism (competitive structure preservation).

Validation on 25 years of T-100 data demonstrates strong performance: 9.55% MAPE on route-level predictions, outperforming ARIMA by 19.6%, with carrier-specific decom-

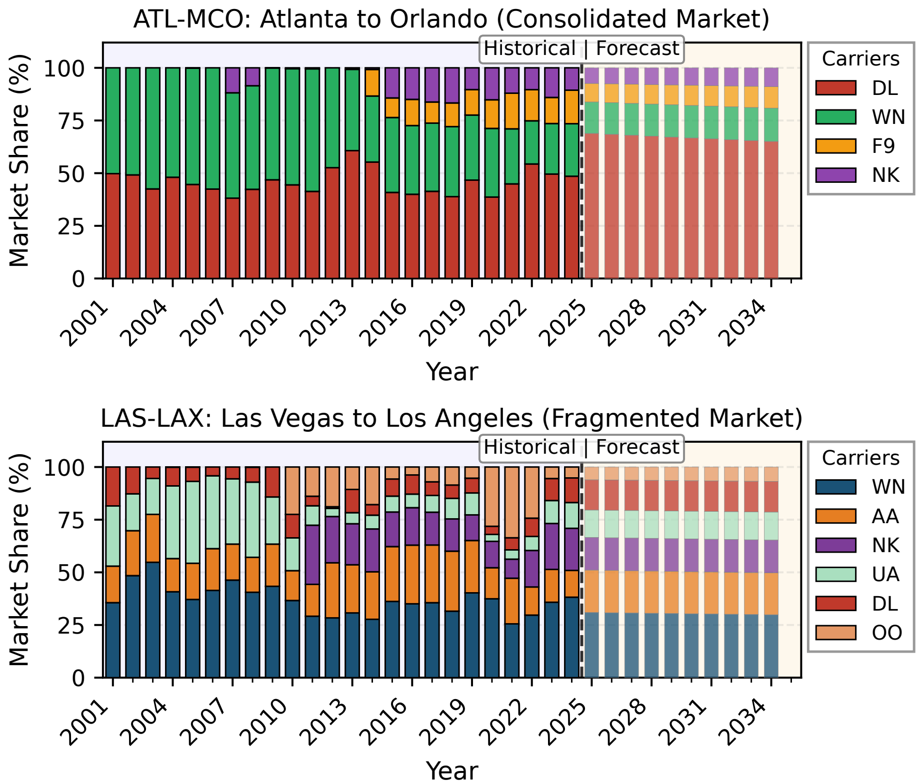

Figure 3: Market share evolution (2001–2035). Solid bars show actual historical shares (2001–2024); semi-transparent bars show forecasts (2025–2035). ATL-MCO (top): Delta maintains hub dominance at 65–68% throughout, with Southwest, Frontier, and Spirit competing for the remainder. LAS-LAX (bottom): Fragmented competition across six carriers persists across the full period, with Southwest leading at 27–30%. All bars sum to 100% each year. Historical data from BTS T-100 (2001–2024); forecasts from hierarchical attention-LSTM with persistence-based decomposition (2025–2035).

position preserving observed competitive hierarchies across 24 years of historical data. The framework handles diverse market structures from hub-dominated routes (Delta 65-70% on ATL-MCO) to fragmented markets (Southwest 27% on LAS-LAX), maintaining zero-sum constraints while capturing structural advantages that persist across business cycles and disruptions including COVID-19.

This capability addresses a critical gap in aviation forecasting. Airlines currently forecast route aggregates and separately estimate their competitive position, introducing uncertainty in fleet deployment and capacity planning decisions. Our framework provides direct carrier-specific demand forecasts supporting multi-year capital allocation where \$20–\$30 B annual expenditures depend on accurate projections.

The persistence-based decomposition offers a methodological contribution beyond aviation. Where hierarchical forecasting typically handles additive aggregation (regional sales summing to national totals), our zero-sum structure addresses competitive markets where gains for one participant necessarily imply losses for others. This approach generalizes to telecommunications market share, retail competitive positioning, and other domains exhibiting persistent competitive hierarchies.

Future work could extend the framework in several directions: incorporating pricing and service quality features into carrier decomposition, modeling network effects where route-level changes propagate through hub-spoke systems, and adapting the zero-sum constraint to handle market entry and exit dynamics. The framework could also integrate real-time demand signals and booking data to update forecasts as new information becomes available, enabling dynamic capacity adjustment beyond static strategic planning.

References

Abdelghany, A.; Abdelghany, K.; and Azadian, F. 2005. Modeling Airline Passengers' Choice of Flight Schedule in a Hub-and-Spoke System. *Transportation Research Part E*, 41(3): 201–216.

Adler, N. 2001. Competition and Network Effects in Airline Markets. *Transportation Research Part A*, 35(4): 329–346.

Athanasopoulos, G.; Ahmed, R. A.; and Hyndman, R. J. 2009. Hierarchical Forecasts for Australian Domestic Tourism. *International Journal of Forecasting*, 25(1): 146–166.

Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

Barnhart, C.; Kniker, T. S.; and Lohatepanont, M. 2003. Airline Fleet Assignment with Time Windows. *Transportation Science*, 37(1): 1–20.

Berry, S.; and Jia, P. 2005. Airline hubs: Costs, markups and the implications of customer heterogeneity. *NBER Working Paper*, (11347).

Berry, S.; and Jia, P. 2010. Tracing the Woes: An Empirical Analysis of the Airline Industry. *American Economic Journal: Microeconomics*, 2(3): 1–43.

Box, G. E.; Jenkins, G. M.; Reinsel, G. C.; and Ljung, G. M. 2015. *Time Series Analysis: Forecasting and Control*. John Wiley & Sons.

Cho, K.; Van Merriënboer, B.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.

Garrow, L. A. 2010. Discrete Choice Modeling for Aviation. *Journal of Air Transport Management*, 16(4): 205–210.

Grosche, T.; Rothlauf, F.; and Heinzl, A. 2007. Gravity models for airline passenger volume estimation. *Journal of Air Transport Management*, 13(4): 175–183.

Guo, X.; Grushka-Cockayne, Y.; and De Reyck, B. 2022. Forecasting airport transfer passenger flow using real-time data and machine learning. *Manufacturing & Service Operations Management*, 24(6): 3193–3214.

Hansen, M. M.; Gillen, D.; and Djafarian-Tehrani, R. 2011. Competition and network effects in airline markets. *Transportation Research Part B: Methodological*, 45(4): 592–602.

Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural Computation*, 9(8): 1735–1780.

Hyndman, R. J.; Ahmed, R. A.; Athanasopoulos, G.; and Shang, H. L. 2011. Optimal Combination Forecasts for Hierarchical Time Series. *Computational Statistics & Data Analysis*, 55(9): 2579–2589.

Hyndman, R. J.; Koehler, A. B.; Ord, J. K.; and Snyder, R. D. 2008. *Forecasting with Exponential Smoothing: The State Space Approach*. Springer.

Jafari, N. 2022. The chaos on US domestic airline passenger demand forecasting caused by COVID-19. *International Journal of Business Forecasting and Marketing Intelligence*, 7(3): 241–258.

Jafari, N.; and Lewison, M. 2024. Forecasting air passenger traffic and market share using deep neural networks with multiple inputs and outputs. *Frontiers in Artificial Intelligence*, 7: 1429341.

Jorge-Calderon, J. D. 2007. Analysis of Air Travel Demand in an International Context. *Journal of Air Transport Management*, 13(5): 264–271.

Kumar, S.; Hussain, L.; Banarjee, S.; and Reza, M. 2022. A hybrid deep learning approach for multi-step ahead prediction. *Expert Systems with Applications*, 186: 115746.

Pai, P.-F.; and Hong, W.-C. 2010. Comparative Study of Market Share Forecasting Methods. *International Journal of Production Economics*, 127(1): 165–178.

Rangapuram, S. S.; Seeger, M. W.; Gasthaus, J.; Stella, L.; Wang, Y.; and Januschowski, T. 2021. End-to-End Learning of Coherent Probabilistic Forecasts for Hierarchical Time Series. *International Conference on Machine Learning*, 8832–8843.

Richard, O. 2008. Airport Market Power and Pricing in US Airline Markets. *Transportation Research Part E*, 44(1): 89–111.

Taieb, S. B.; Taylor, J. W.; and Hyndman, R. J. 2017. Coherent Probabilistic Forecasts for Hierarchical Time Series. In *International Conference on Machine Learning*, 3348–3357.

Tay, R.; and McCarthy, P. S. 2011a. Market share analysis. *Handbook of Transport Modelling*, 140–158.

Tay, R.; and McCarthy, P. S. 2011b. Market Share Models. *Handbook of Transport Economics*, 140–158.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30.

Wei, W.; and Hansen, M. 2020. Forecasting Airport Transfer Passenger Flow Using Real-time Data and Machine Learning. *Transportation Research Part E*, 139: 101965.

Wickramasuriya, S. L.; Athanasopoulos, G.; and Hyndman, R. J. 2019. Optimal Forecast Reconciliation for Hierarchical and Grouped Time Series Through Trace Minimization. *Journal of the American Statistical Association*, 114(526): 804–819.

Wickramasuriya, S. L.; Turlach, B. A.; and Hyndman, R. J. 2020. Forecasting Hierarchical and Grouped Time Series Through Probabilistic Coherence. *arXiv preprint arXiv:2006.02043*.

Zhao, Z.; Chen, W.; Wu, X.; Chen, P. C.; and Liu, J. 2020. LSTM network: A deep learning approach for short-term traffic forecast. *IET Intelligent Transport Systems*, 14(2): 68–75.

Zou, L.; and Yu, C. 2011. Empirical Analysis of Route Entry and Exit in US Airline Markets. *Transportation Research Part E*, 47(3): 317–332.