

# Neural Technologies for Predictive Analytics in Tourism Governance processes

Natalia, Pichierri<sup>1,†</sup>, Daniele, Margiotta<sup>1,2</sup>, Andriy, Shcherbakov<sup>2</sup>, Danilo, Croce<sup>1,\*,†</sup> and Roberto, Basili<sup>1,†</sup>

<sup>1</sup>Department of Enterprise Engineering, University of Rome, Tor Vergata

<sup>2</sup>Reveal S.r.l.

## Abstract

Tourism governance increasingly faces challenges from climate disruptions, overtourism, and global crises, demanding predictive and data-driven approaches. This paper presents a novel forecasting framework based on Transformer architectures, applied within Stendhal—an intelligent destination management platform including information about 231 Italian destinations and 2,553 municipalities. Integrating heterogeneous data sources such as official statistics, business registers, user-generated content, and media signals, Stendhal models complex tourism interdependencies. We evaluate three Transformer-based models (Informer, TimeXer, and a custom Informer-Ext) for provincial-level monthly tourism flow forecasting. Experiments on ten years of data (2014–2023) show significant gains over linear baselines, especially in predicting the direction of changes in tourist arrivals. Incorporating territorial metadata and relational features—including regional trends and COVID-19 indicators—further boosts predictive accuracy. Results demonstrate that advanced neural models provide accurate and interpretable economic forecasts for anticipatory, evidence-based tourism governance.

## Keywords

Predictive Analytics, Time Series Forecasting, Transformer Neural Networks, Tourism, Destination Management

## 1. AI for Predictive Tourism Governance

In 2024, tourism in Italy stands at the crossroads of three interlinked challenges: recovery from the COVID-19 pandemic, renewed pressures of overtourism, and growing impacts from the climate crisis. These dynamics compel Destination Management Organizations (DMOs), policymakers, and regional authorities to adopt strategies that are granular, proactive, and grounded in robust evidence. However, existing monitoring systems remain predominantly descriptive, ex-post, and fragmented, unable to fully represent tourism as an evolving relational system shaped by complex interactions among tourists, businesses, institutions, and communities [1, 2]. Recent literature advocates for a paradigm shift towards systemic approaches that integrate diverse data sources and explicitly account for socio-ecological interdependencies, supporting anticipatory and adaptive governance [3, 4, 5]. In this context, Artificial Intelligence (AI) and advanced analytics emerge as critical enablers for such transformative change [6].

This paper builds on **Stendhal**, a distributed AI system for destination governance developed between 2022 and 2024 via participatory action research, engaging experts in AI, tourism economics, and public policy alongside institutional stakeholders (available at <https://stendhal.isnart.it>). Central to Stendhal is an innovative, relational view of tourism destinations—not as static geographic or administrative entities, but as complex ecosystems of interdependent actors collaboratively shaping destination experiences and resilience. Currently operational across 231 destinations encompassing 2,553 Italian municipalities, Stendhal combines more than 80 structured and unstructured variables, including institutional statistics (ISTAT, Bank of Italy), business registers (Unioncamere), user-generated content (e.g., TripAdvisor reviews), and other media-derived signals. Through semantic enrichment and machine learning techniques, Stendhal enables the multidimensional analysis of tourism phenomena, overcoming traditional separations between demand and supply. The platform supports destination

*Ital-IA 2025: 5th National Conference on Artificial Intelligence, organized by CINI, June 23-24, 2025, Trieste, Italy*

\*Corresponding author.

✉ natalia.pichierri@alumni.uniroma2.eu (N. Pichierri); margiotta@revealsrl.it (D. Margiotta); shcherbakov@revealsrl.it (A. Shcherbakov); croce@info.uniroma2.it (D. Croce); basili@info.uniroma2.it (R. Basili)

ORCID 0000-0001-9111-1950 (D. Croce); 0000-0001-5140-0694 (R. Basili)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

governance by offering insights across four key dimensions: **tourism system integration**, measured via composite indicators spanning products, firms, infrastructure, governance, visitor profiles, and community engagement; **sentiment and experience analysis**, based on over four million tourist reviews categorized into thematic clusters; **destination reputation**, inferred from the sentiment of more than one million news articles; and **sustainability assessment**, through local indicators of economic impact, social cohesion, and environmental pressure.

While Stendhal already supports comprehensive descriptive analytics and benchmarking, this study advances its capabilities into predictive modeling, exploring the feasibility of forecasting tourism flows—monthly arrivals and overnight stays—at the provincial level. This forecasting challenge involves complex temporal and spatial dependencies, heterogeneous and multidimensional datasets, and requires robust, interpretable predictive models.

A forecasting framework built on Stendhal’s dataset, which integrates over 40 tourism-related features, is here presented. We conduct a comparative evaluation of three advanced neural architectures—Informer [7], TimeXer [8], and a custom multi-source variant (Informer-Ext)—and assess their performance across multiple forecasting horizons (1-, 2-, and 12-months). Among the different variable classes, such as endogenous time series, exogenous signals, and time-independent territorial descriptors, this work examines which one contributes most to improve predictive accuracy with respect to traditional approaches, such as linear regressors and support vector machines. We also show that incorporating territorial and relational knowledge, in form of destination types and regional aggregations, enhances reliability and interpretability of the model in favor of a better anticipatory destination governance.

## 2. Neural Approaches to Forecasting over complex data streams

Forecasting aims to predict the future trajectory of one or more target variables  $\mathbf{Y}$ , based on an ordered sequence of past observations  $\mathbf{X}$  [9]. This predictive task is formally defined as the learning of a function:  $f : \mathbb{R}^{T \times d} \rightarrow \mathbb{R}^{H \times o}$  where  $T$  is the input sequence length,  $d$  the number of input features,  $H$  the forecasting horizon, and  $o$  the dimensionality of the output [8]. In high-complexity domains such as tourism, forecasting depends on an integrated use of heterogeneous sources, including temporal, spatial, and semantic data and it proceeds by the joint modeling of both endogenous and exogenous dependencies [10]. In the field of predictive modeling, two main paradigms are commonly adopted: classical approaches such as linear regressors (e.g., Support Vector Regression), and modern neural architectures based on Transformers (e.g., Informer [7] and TimeXer [8]).

Classical regressors such as Support Vector Regression (SVR) predict a single target variable per input instance, relying on supervised learning and manual feature engineering to transform time series into static vectors via statistical descriptors. While effective with limited data, these methods struggle to capture complex temporal dependencies [11].

Interestingly, Transformer-based neural models—such as Informer [7] and TimeXer [8]—can simultaneously forecast multiple targets, over various horizons by processing raw multivariate sequences and contextual features through learned embeddings. These models excel at capturing intricate temporal and cross-variable relationships, thanks to mechanisms like ProbSparse attention (Informer) and separated attention flows for endogenous and exogenous signals (TimeXer). Although these approaches improve flexibility and predictive accuracy, they require greater computational resources and more sophisticated training strategies. Tourism forecasting in particular demands models capable of handling heterogeneous inputs. These include temporal sequences but also geographic identifiers and descriptors for a variety of semantic dimensions. Moreover, they exhibit limited historical depth and usually many parallel time series, such as those regarding the many different individual regions or destinations. As a trade-off, linear models are computationally efficient but less expressive, whereas Transformers model more complex cross variable dependencies. While more demanding in training time and resources, the latter provide superior accuracy and adaptability in complex scenarios. Experimental results over realistic use cases in touristic destination management will be hereafter discussed.

### 3. Experimental Setting: Case Study and Objectives

Forecasting tourism flows over the Stendhal data is here studied by comparing linear models over manually engineered features with Transformer-based neural networks, trained on contextualized raw data. The task consists in predicting the monthly year-on-year percentage change in tourist arrivals for each Italian province. This formulation captures seasonal effects and external shocks relevant to tourism analytics [9, 10]. For example, if the forecasting task is performed in May and the target is to predict arrivals for June, the model is trained to estimate the percentage change between June 2025 and June 2024, that is:

$$\Delta_{\text{arrivals}} = \frac{\text{Arrivals}_{\text{June, 25}} - \text{Arrivals}_{\text{June, 24}}}{\text{Arrivals}_{\text{June, 24}}}$$

The same logic is applied for forecasting horizons of two, three, or more months ahead. For instance, a 3-month ahead forecast made in May 2025 will target the variation between August 2025 and August 2024, and so forth. This year-on-year delta formulation ensures that the predictive task remains robust to long-term trends and external shocks (such as pandemic effects), while maintaining a clear interpretability aligned with the expectations of tourism governance practitioners.

**Dataset Description.** The dataset used in this study covers the period from 2014 to 2023, with a monthly temporal granularity. The data are organized at the provincial<sup>1</sup> level, with additional information available for over 2,500 municipalities across Italy. For each province, the series typically includes approximately 120 data points per variable (i.e., one for each month in the ten-year span). When modeling at the monthly level, the effective number of data points per province and month is thus inherently limited by the historical depth of the available records. This relative “shortness” is a key challenge in time series forecasting, as it constrains the ability of both classical and neural models to learn complex temporal dependencies [12]. The main variables included in the dataset are: *i*) Tourist arrivals (*arrivi*) and overnight stays (*presenze*), reported monthly for each province, and further disaggregated by visitor origin (residents vs. non-residents) and accommodation type (hotels vs. non-hotel facilities); *ii*) Segmented indicators: time series split by residency status and accommodation category, enabling a granular analysis of tourism flows; Exogenous features: contextual and territorial indicators, including composite indices describing tourism vocations, COVID-19 impact variables, and regional/provincial identifiers; *iii*) Other metadata: month identifiers (one-hot encoded), municipality codes, and higher-level aggregation keys (province, region).

**Methodological Framework: Linear vs Neural Models.** This study contrasts two fundamentally different modeling paradigms for forecasting tourism flow variations: a linear approach with explicit manual feature engineering, and neural models based on Transformer architectures. Both approaches are evaluated on the same forecasting task and dataset, under identical data splits, evaluation metrics, and validation setups. The dataset, provided by ISNART as part of the Stendhal initiative, consists of monthly tourism data from 2014 to 2023 at the provincial level and includes more than 2,500 municipalities. Alongside arrivals and overnight stays, the data provide a wide set of segmented indicators (by residency status and accommodation type), exogenous features such as composite territorial vocation indices, COVID-19 impact variables, as well as province and region identifiers and other metadata (month identifiers, municipality codes, higher-level keys).

In the linear approach, each raw time series is transformed into an extensive set of aggregated features designed to summarize its recent temporal dynamics. For each of the eight core series—covering all combinations of arrivals/presences, residents/non-residents, and hotel/non-hotel accommodations—statistical descriptors such as mean, standard deviation, last value, and percentage change are computed over multiple rolling windows (2, 4, 6, and 12 months), both for the recent past and for equivalent periods in the previous year. This yields a comprehensive feature vector encoding both short- and long-term temporal patterns, local volatility, and recent momentum for each indicator. These quantitative features are complemented by categorical variables encoding the temporal and spatial context—calendar month (as one-hot), province, and region. The Support Vector Regressor (SVR)[11],

<sup>1</sup>Only provinces with a comparable number of available data were considered.

implemented via the KeLP framework [13], is chosen as the main algorithm due to its robustness in high-dimensional and heterogeneous spaces.

**Table 1**

Output selection strategies for 1-month and 2-month forecasts. The table summarizes the different approaches adopted to select the final prediction when neural architectures (e.g., Informer, TimeXer) are used for multi-step forecasting. For each target horizon (1 or 2 months ahead), multiple model configurations are possible depending on the prediction length (*pred\_len*) and the selection or aggregation strategy applied to the model’s output sequence: *first* refers to the first available output, *second* to the second position in the prediction window, and *avg* denotes the arithmetic mean across all overlapping predictions for a target month. Models trained with longer output windows (e.g., *pred\_len*=12) can contribute predictions for both short- and medium-term targets through appropriate selection or averaging. These choices are important for ensuring comparability of results across model variants and forecasting horizons, especially when input sequences are short and overlapping predictions are generated. The same evaluation setup, with harmonized output selection, is applied to both linear (SVR) and neural models, as well as to the naive baseline, to enable a fair pointwise comparison for each province and target month.

1-month forecast			2-month forecast		
Model	pred_len	Strategy	Model	pred_len	Strategy
1-Month	1	first	1-Month	1	N/A
2-Month (Early Selection)	2	first	2-Month (Early Selection)	2	second
2-Month (Averaging)	2	avg	2-Month (Averaging)	2	N/A
12-Month (Early Selection)	12	first	12-Month (Early Selection)	12	second
12-Month (Averaging)	12	avg	12-Month (Averaging)	12	avg from second

The neural approach employs Transformer-based architectures—Informer<sup>2</sup>, TimeXer<sup>3</sup>, and custom extensions—which process raw, multivariate time series directly, without explicit feature engineering. Input sequences include arrivals, presences, and their segmentations, while static and categorical variables (e.g., province, vocation, month) are integrated via learnable embeddings or concatenation. Preprocessing is minimal: quantitative features are normalized both globally and locally or encoded as year-on-year changes; territorial variables are aggregated from municipal to provincial level using either arithmetic or population-weighted means, labeled as *v1* and *v2* inputs. Model inputs are further enriched by testing various combinations of endogenous and exogenous features, such as regional aggregates, COVID-19 indicators, and segmentations by visitor or accommodation type. Some variants provide each province with its own or overall regional summary to capture cross-region dependencies. Categorical features are encoded either as embeddings or one-hot codes, as appropriate. Both approaches face the challenge of limited data per province, as a decade of monthly records yields relatively few samples for supervised learning. Additionally, Transformer-based models can produce multiple predictions for the same target month due to their sliding window, multi-step forecasting approach. To ensure robust and comparable evaluation across horizons, we test different output selection strategies—such as taking the first or second output, or averaging overlapping predictions—as detailed in Table 1. A simple “naive” baseline is included for reference: for any given month and province, the predicted delta is simply the year-on-year percentage change observed in the previous year for the same month (e.g., the forecast for June 2025 is set equal to the delta observed for June 2024). This baseline offers a lower bound and highlights the added value of more complex modeling. All predictions are generated for each province and target month in the test set, ensuring fair, point-wise evaluation over the full spatiotemporal grid, and assessing the ability of each method to generalize under identical data conditions.

**Evaluation Metrics.** All models are trained and validated using a temporal cross-validation scheme, with the period 2015–2022 for training, the first semester of 2023 for validation, and the second semester for testing, following [12]. Two key metrics are used: Mean Absolute Error (MAE), which measures the average magnitude of prediction errors and is scale-invariant, and the Sign Agreement Ratio (SIGN), which quantifies the proportion of times the model correctly predicts the direction (increase or decrease) of arrivals, reflecting operational needs in tourism management. Where relevant, results are also disaggregated by region, forecasting horizon, and model variant.

**Result discussion.** We evaluate forecasting performance using two complementary metrics: the Mean

<sup>2</sup><https://github.com/zhouhaoyi/Informer2020>

<sup>3</sup><https://github.com/thuml/TimeXer>

Absolute Error (MAE) and the Sign Agreement Ratio (SIGN), the latter indicating the frequency with which the predicted and observed year-on-year variation in arrivals have the same sign. The analysis focuses on the last six months of 2023 across 103 Italian provinces. Results in Table 2 (SIGN) reveal that TimeXer achieves the best overall directional accuracy. For the 1-month horizon, the most effective configuration combines local and regional arrival trends, disaggregated tourist flows (arrivals and presences), COVID-19 signals, and standardized province-level values; the final forecast is obtained as the arithmetic mean of outputs from a two-month prediction window. This multi-source input, alongside output averaging, maximizes the likelihood of correctly anticipating the direction of monthly changes. At a 2-month horizon, TimeXer again leads in SIGN, especially when the model is trained with a 12-month prediction window and predictions are averaged. These results suggest that including relational and territorial information, as well as ensemble-like averaging across overlapping model outputs, is critical for capturing complex interdependencies in tourism flows.

**Table 2**

SIGN – Forecasting performance across models and horizons for the predicted variable: year-on-year percentage change in tourist arrivals ( $\Delta_{arrivals}$ ). Model abbreviations: *Inf* = Informer, *TX* = TimeXer, *prov* = province ID, *v1/v2* = territorial vocation (arithmetic or population-weighted mean), *+Reg* = regional aggregate features. Best-performing configurations typically combine endogenous time series with exogenous territorial descriptors. Input composition and output selection strategies (see Table 1) significantly impact results.

Model	1M SIGN	Arch.	Input	2M SIGN	Arch.	Input
1m	0.777 0.785	Inf <sub>v1</sub> TX	$\Delta arr + cov$ $\Delta arr + arr + cov$			
2m es	0.777 0.786	Inf TX + Reg	$\Delta arr + cov$ $\Delta arr + arr + cov + \Delta pres (std)$	0.782 0.783	Inf <sub>v2</sub> TX <sub>prov</sub>	$\Delta arr + cov$ $\Delta arr + cov + pres$
2m avg	0.787 <b>0.796</b>	Inf <b>TX + Reg</b>	$\Delta arr$ <b><math>\Delta arr + arr + cov + pres + \Delta pres (std)</math></b>			
12m es	0.774 0.779	Inf <sub>v1</sub> TX <sub>prov</sub>	$\Delta arr + cov$ $\Delta arr + arr + cov + pres + split$	0.772 0.786	Inf <sub>v1</sub> TX <sub>prov</sub>	$\Delta arr + cov$ $\Delta arr + cov + split + \Delta split (std)$
12m avg	0.778 0.791	Inf <sub>v1</sub> TX <sub>prov</sub>	$\Delta arr + cov$ $\Delta arr + cov$	0.817 <b>0.836</b>	Inf <b>TX + Reg</b>	$\Delta arr$ <b><math>\Delta arr + arr + pres + \Delta pres + split (std)</math></b>
Baseline	<b>0.722</b>	–	–	<b>0.722</b>	–	–
SVR	<b>0.751</b>	–	<b>256 handcrafted-features</b>	<b>0.754</b>	–	<b>256 handcrafted-features</b>

**Table 3**

MAE – Forecasting performance across models and horizons for the predicted variable: year-on-year percentage change in tourist arrivals ( $\Delta_{arrivals}$ ).

Model	1M MAE	Arch.	Input	2M MAE	Arch.	Input
1m	0.085 <b>0.090</b>	Inf <b>TX</b>	$\Delta arr + arr + cov (std)$ <b><math>\Delta arr + split</math></b>			
2m es	0.083 0.093	Inf <sub>v1</sub> TX	$\Delta arr + cov$ $\Delta arr + pres + \Delta pres + split$	0.089 0.108	Inf <sub>e1</sub> TX + Reg	$\Delta arr + cov$ $\Delta arr + cov + pres + \Delta pres (std)$
2m avg	<b>0.083</b> 0.097	Inf <sub>v1</sub> TX	<b><math>\Delta arr + cov</math></b> $\Delta arr + pres + \Delta pres + split$			
12m es	0.093 0.109	Inf <sub>v1</sub> TX	$\Delta arr + cov$ $\Delta arr + pres + \Delta pres + split$	0.096 0.126	Inf <sub>prov</sub> TX <sub>v2</sub>	$\Delta arr$ $\Delta arr + cov$
12m avg	0.090 0.117	Inf <sub>v1</sub> TX	$\Delta arr + cov$ $\Delta arr + arr + pres + split$	<b>0.086</b> 0.117	Inf TX	<b><math>\Delta arr</math></b> $\Delta arr + arr + pres + split$
Baseline	<b>0.176</b>	–	–	<b>0.176</b>	–	–
SVR	<b>0.074</b>	–	<b>256 handcrafted-features</b>	<b>0.074</b>	–	<b>256 handcrafted-features</b>

Turning to MAE (Table 3), the best 1-month accuracy (0.083) is achieved by an extended Informer variant incorporating territorial descriptors and COVID-19 indicators, with predictions averaged over a 2-month window. For the 2-month horizon, standard Informer—using only arrival variation as input and averaging predictions from a 12-month model—yields the lowest MAE (0.086). These findings highlight the Informer’s ability to efficiently leverage minimal contextual inputs when tasked with precise value estimation. Support Vector Regression (SVR), despite achieving competitive MAE values, especially when equipped with a comprehensive set of 256 handcrafted features, consistently underperforms neural models in terms of SIGN. This points to a limitation in capturing the directionality and dynamics of seasonal shocks, a key aspect for operational tourism forecasting. SVR’s reliance on manual feature engineering and costly hyperparameter tuning further limits its scalability relative to neural models that learn latent structure directly from raw data. Across all configurations, the naive baseline—repeating the



previous year's delta—provides a consistent lower bound, underlining the value added by both advanced linear and neural strategies. The most effective approaches across all horizons consistently integrate both endogenous time series and exogenous territorial descriptors, while the precise input composition and output selection strategy—summarized in Table 1—substantially affect final performance.

## 4. Conclusions

This study demonstrates that Transformer-based models, particularly custom variants of Informer and TimeXer, provide robust and reliable forecasts of provincial tourism flows even with limited historical data. Performance improves notably when integrating territorial vocations, COVID-19 indicators, and granular segmentations of arrivals and presences. Transformer architectures consistently surpass classical SVR in capturing the directionality of year-on-year changes, as measured by SIGN, highlighting their advantage for operational forecasting. While no single contextual feature dominates across all scenarios, combining multiple territorial and relational descriptors, such as territorial vocations or COVID-19 indicators, generally enhances accuracy, although their impact varies across models and prediction horizons. Future work will explore alternative Transformer variants, task generalization to other tourism indicators, and adaptive strategies for contextual feature selection.

## Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

## References

- [1] C. Cooper, Managing tourism knowledge: a review, *Tourism Review* 73 (2018) 507–520.
- [2] S. M. Allameh, Antecedents and consequences of intellectual capital: The role of social capital, knowledge sharing and innovation, *Journal of Intellectual Capital* 19 (2018) 858–874.
- [3] G. I. Crouch, J. R. B. Ritchie, Tourism, competitiveness, and societal prosperity, *Journal of Business Research* 44 (1999) 137–152.
- [4] T. Mihalič, Conceptualising overtourism: A sustainability approach, *Annals of Tourism Research* 84 (2020) 103025.
- [5] M. Volgger, G. Erschbamer, H. Pechlaner, Destination design: New perspectives for tourism destination development, *Journal of Destination Marketing & Management* 19 (2021) 100561.
- [6] Z. Doborjeh, N. Hemmington, M. Doborjeh, N. K. Kasabov, Artificial intelligence: a systematic review of methods and applications in hospitality and tourism, *International Journal of Contemporary Hospitality Management* 34 (2022) 1154–1176.
- [7] H. Zhou, S. Zhang, J. Peng, S. Zhang, J. Li, H. Xiong, W. Zhang, Informer: Beyond efficient transformer for long sequence time-series forecasting, *CoRR abs/2012.07436* (2020).
- [8] Y. Wang, H. Wu, J. Dong, G. Qin, H. Zhang, Y. Liu, Y. Qiu, J. Wang, M. Long, Timexer: Empowering transformers for time series forecasting with exogenous variables, in: *NeurIPS 2024*, Vancouver, BC, Canada, December 10 - 15, 2024, 2024.
- [9] B. Zhang, Y. Pu, Y. Wang, J. Li, Forecasting hotel accommodation demand based on lstm model incorporating internet search index, *Sustainability* 11 (2019) 4708.
- [10] H. Song, R. T. Qiu, J. Park, A review of research on tourism demand forecasting: Launching the annals of tourism research curated collection on tourism demand forecasting, *Annals of tourism research* 75 (2019) 338–362.
- [11] A. J. Smola, B. Schölkopf, A tutorial on support vector regression, *Statistics and computing* 14 (2004) 199–222.
- [12] N. D. Bokde, Z. M. Yaseen, G. B. Andersen, Forecasttb—an r package as a test-bench for time series forecasting—application of wind speed and solar radiation modeling, *Energies* 13 (2020) 2578.
- [13] S. Filice, G. Castellucci, G. D. S. Martino, A. Moschitti, D. Croce, R. Basili, Kelp: a kernel-based learning platform, *Journal of Machine Learning Research* 18 (2018) 1–5.