



Kuva Harri Tarvainen

TOURCAST -

a Finnish tourism nowcasting and forecasting model

REPORT



TOURCAST -

a Finnish tourism nowcasting and forecasting model

REPORT

Ossi Nurmi, Henri Luomaranta, Paolo Fornaro



TABLE OF CONTENTS

TIIVISTELMÄ	4
ABSTRACT	6
INTRODUCTION	8
DATA SOURCES	9
Survey for accommodation establishments	10
Flight bookings data	12
METHODS AND MODELING	14
Imputation methods to deal with missing data	14
Models for forecasting short-term tourism volumes	14
PRELIMINARY RESULTS	15
Nowcasting based on accommodation survey	15
Monthly forecast based on flight bookings	18
Tourcast: a tourism forecasting model	20
CONCLUSIONS	24
REFERENCES	25



Kuva Mikko Niikinen

TIIVISTELMÄ

Matkailutoimialalla on selkeä tarve ajantasaiselle ja ennakoivalle tiedolle matkailun kehityksestä. Nykyiset tilastot ulkomaalaisten Suomeen suuntautuvasta matkailusta julkaistaan kuitenkin neljän viikon viiveellä.

”Matkailun ennakkointimalli” -hankkeen tavoitteena oli muodostaa viitekehys ulkomaalaisten matkailijoiden Suomen majoitusliikkeissä yöpymisten pikaestimoinnille (nowcast) ja lyhyen aikavälin ennustamiselle. Lähtömaista keskityttiin venäläisten, saksalaisten, yhdysvaltalaisien, kiinalaisten, japanilaisten ja korealaisten matkailuun.

Lähestymistapa perustuu ennusteisiin, jotka yhdistetään useista tilastollisista malleista ja koneoppimismenetelmistä. Tilastokeskus ja ETLA ovat aiemmin soveltaneet näitä menetelmiä Suomen bkt:n pikaestimaattien sekä päätoimialojen liikevaihdon ennakkotietojen tuottamiseen.

Matkailun ennakkointimalli tuottaa päättyneen kuukauden **pikaestimaatin** sekä enimmillään 6 kuukauden päähän ulottuvat **ennusteet**. Pikaestimaatti tuotetaan majoitustilaston liikekohtaisista tiedoista. Malli käyttää ajoissa vastanneiden liikkeiden tietoja ja puuttuvat tiedot estimoidaan tilastollisin menetelmin. Ennusteet perustuvat lentovarauksiin, joita ulkomaalaiset matkailijat tekivät kolmen suurimman globaalin jakelupalvelun (Amadeus, Sabre, Galileo) kautta. Lentovarausaineisto ei sisällä suoraan lentoyhtiöiltä tehtyjä varauksia. Rajapinnan aineistoon on Business Finlandille tuottanut espanjalainen yritys Forward Keys.



Mallien paremmuutta arvioitiin sen perusteella, kuinka paljon estimaattien prosenttivirheen itseisarvo poikkeaa tilaston lopullisista yöpymisistä. Tarkkuuden parantamiseksi on suositeltavaa ennustaa yöpymisten kertymää useammalle kuukaudelle yksittäisten kuukausien sijaan.

Mallin koetulokset tuotettiin ajanjaksolle 2015 tammikuu – 2019 heinäkuu. Pikaestimaatin virhe oli tarkastelujaksolla keskimäärin 1 - 2 prosenttia ulkomaalaisten kokonaisyöpymisille ja lähtömaittain 1 - 5 prosenttia. 1 - 6 kuukauden päähän ulottuvan ennusteen virhe ulkomaalaisten kokonaisyöpymisille oli keskimäärin 3 – 5 prosenttia. Yksittäisille maille virhe oli keskimäärin 4 – 11 prosenttia.

Matkailun ennakkointimallin avulla voidaan näin ollen tuottaa:

- Erittäin tarkka estimaatti päättyneen kuukauden yöpymisistä noin 10 päivän viiveellä
- Kohtalaisen tarkka ennuste ulkomaalaisten matkailijoiden kokonaisyöpymisistä lyhyellä (1 – 3 kuukautta) ja keskipitkällä (4 – 6 kuukautta) aikavälillä
- Riittävän tarkka ennuste useimpien lähtömaiden yöpymisistä lyhyellä aikavälillä

Mallit on toteutettu R- ohjelmointikielellä ja aineiston käsittelyä tehdään myös SAS-kielellä. Mallin jatkokehitystarpeiksi tunnistettiin mm. pitkän aikavälin- (6 – 12 kuukautta) sekä aluetason ennusteiden tuottaminen.

Hankkeen rahoitti Business Finland ja työn toteuttivat Ossi Nurmi ja Henri Luomaranta Tilastokeskuksesta sekä Paolo Fornaro ETLA:sta.



ABSTRACT

Despite the strong demand from the tourism industry for timely information, as well as for estimates of future developments, the current statistics on inbound tourism to Finland are published with a four weeks delay.

The “Tourism forecasting model (Tourcast)” -project was established to formulate a framework to estimate nowcasts and short-term forecasts of nights spent by foreign tourists at Finnish accommodation establishments. The project focused on the following countries of origin: Russia, Germany, USA, China, Japan and Korea.

The approach is based on combining predictions obtained from a range of statistical models and machine learning techniques. Prior to this project, Statistics Finland and Etila Economic Research (ETLA) have applied these techniques to produce flash estimates of the Finnish economy, as well as of turnovers of main industries in Finland.

The Tourcast-model produces **rapid estimates** for the most recent month and **forecasts** for up to 6 months ahead. The establishment-level dataset of accommodation statistics was used as main source for the rapid estimate. The model uses data from those establishments that responded early, and missing data were imputed using various statistical methods. Forecasts are based on flight bookings that foreign tourists made through 3 global distribution services (Amadeus, Sabre, Galileo). The flight bookings data do not include direct bookings from airlines. Interface to the data was provided for Business Finland by Forward Keys, a Spanish company.



The accuracy of the models was evaluated based on how much the estimates deviate from the final statistics on nights spent, in particular we compute the mean absolute percentage error. In order to improve forecasting accuracy, the recommended approach is to consider the cumulative nights predicted up to several months ahead, instead of evaluating the forecasts of individual months.

The test results were produced for the period January 2015 up to July 2019. The error for the rapid estimates (nowcast) was on average 1 – 2 per cent for total foreign nights spent and 1 – 5 per cent for individual countries of origin. Error for the 1 to 6 months forecasts was on average 3 – 5 per cent for total foreign nights spent and 4 – 11 per cent for individual countries.

Using the Tourcast-model, it would therefore be possible to produce:

- A very accurate estimate for nights spent in the previous month with a publication lag of approximately 10 days
- A relatively accurate short- (1 to 3 months) to mid-term (4 to 6 months) forecast for total nights spent by foreign tourists
- A sufficiently accurate short-term forecast for most countries of origin

The models were implemented using R -programming language and data was also processed using SAS-language. Further development of the model may include introducing long-term (6 to 12 months) forecasting as well as providing estimates at a regional level.

The project was financed by Business Finland and the work was carried out by Ossi Nurmi and Henri Luomaranta from Statistics Finland and Paolo Fornaro from ETLA.

INTRODUCTION

International tourism is characterized by rapid changes and high seasonality. Tourism flows from different countries of origin are also quite unique and often unrelated to each other. The number of tourists from one country to another are affected and shaped by factors such as social media, international politics, currency exchange rates, flight ticket prices and holiday periods, just to name a few. Inbound tourism also has a major contribution to the Finnish economy. In 2017 the share of inbound tourism was 5.4 per cent of total Finnish exports of goods and services¹ and 17.8 per cent of Finnish service exports. Inbound and domestic tourism account for a 2.6% share of Finnish GDP.²

Well-established indicators for measuring tourism volume include nights spent and arrivals at hotels and other accommodation establishments, which are currently published by Statistics Finland with a publication lag of around 4 weeks. In addition to measuring past tourism flows, tourism policy makers and industries have a strong need to also forecast future tourism volumes in order to better anticipate policy and business needs.

Prior to this project, Statistics Finland has collaborated with the ETLA Economic Research (ETLA) to develop machine learning techniques and nowcasting models that provide rapid estimates of key indicators of Finnish economic activity³. These indicators include the Trend Indicator of Output (TIO) and quarterly GDP. The current project, Tourism forecasting model or “Tourcast”, was established in order to apply similar techniques within the context of tourism.

The objective of the Tourcast project is to develop a statistical model that provides nowcasts, as well as a short-term forecast spanning up to 6 months to the future, of tourism flows in Finland. At this stage, the focus was only on inbound tourism to Finland in total as well as for selected countries of origin, namely Russia, United Kingdom, Germany, China and Hong Kong, Japan, Korea and the United States. The two primary data sources used in the model are the survey of accommodation establishments and flight bookings to Finland by inbound tourists, provided by Forward Keys.

The project was financed by Business Finland and made in collaboration with Statistics Finland and the ETLA. The project task force consisted of Ossi Nurmi and Henri Luomaranta (Statistics Finland), Paolo Fornaro (ETLA) and Vesa-Pekka Juutilainen (Business Finland). The project steering group also included Katarina Wakonen, Susanne Heikkinen and Niko Marola from Business Finland.

¹ <http://pxnet2.stat.fi/PXWeb/sq/2c4f8d7b-0359-4b5b-b24f-cb2935d1ddad>

² <http://visitfinland.stat.fi/PXWeb/sq/bdf72025-650b-4713-b9e7-9c364a04a534>

³ https://www.stat.fi/static/media/uploads/tup/nowcasting_empecon.pdf

DATA SOURCES

This chapter describes the scope of the two primary data sources used by the Tourcast model. The nowcasting component of Tourcast uses the survey of accommodation establishments as source while the forecasting component takes advantage of the bookings of flight tickets to Finland by inbound visitors provided by Forward Keys.

SURVEY FOR ACCOMMODATION ESTABLISHMENTS

The accommodation statistics of Statistics Finland describe the supply and use of hotel- and other accommodation services. Demand for hotel services is measured by the number of overnight stays. Users of the services are divided into domestic and foreign tourists, where the latter are further broken down by country of residence.

Accommodation statistics cover tourists staying in accommodation establishments with at least 20 beds (or caravan pitches with electricity connection). In other words, it does not include excursionists (day trips) or tourists staying in own or unrented accommodation provided by relatives, friends or employers. Also rented accommodation in individual flats or houses, for example via Airbnb, are not included. Despite these coverage issues, the indicator based on nights spent by foreign tourists in hotels and other establishments is a very good proxy of the total inbound tourism volume in Finland.

Accommodation statistics are currently published approximately 4 weeks after the reference month, and establishments are requested to complete the online survey by the 6th day of each month. Roughly half of the establishments replied by this deadline in 2018. However, this half represents 67 per cent of total nights spent as the early responders are often the bigger hotels. The following graph presents the share of nights spent reported after a certain lag in 2018, for example by 5th day (t+5) or 25th day (t+25) after the end of reference month.




Figure 1 - Share of nights spent reported by a certain lag (days)

In the project, a data set was created on establishment level that includes the following variables:

- Id of the establishment
- Period (year & month)
- Response lead-time (number of days after the reference month)
- Number of nights spent by country of residence

With this data set, it's possible to realistically simulate the data availability at different points in time. The task of the nowcasting model is to impute the missing information for establishments that did not reply on time to the survey. For example, for $t+5$ the model would estimate the nights spent for those establishments that did not respond by the 5th day. The accuracy of these estimates can then be compared to the published accommodation statistics.⁴

Global GDS Systems Market (Travel Networks), 2015

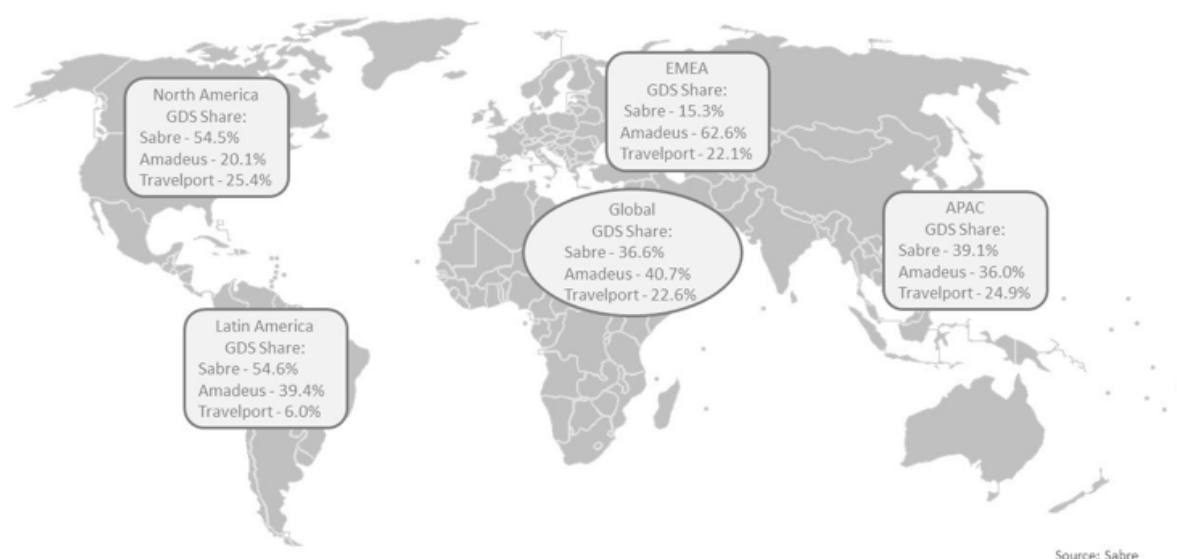


Figure 3 - The GDS market share by region, 2015

The combined data from all three GDS is made commercially available by Forward Keys,⁷ an enterprise based in Spain. The data used in this project was purchased from Forward Keys by Business Finland and made available to Business Finland via an API. The data covers the 7 countries of origin that are most relevant for Business Finland: Russia, United Kingdom, Germany, China and Hong Kong, Japan, Korea and the United States.

The most relevant information for the model concern trips and bookings. A **trip** takes place at a specific time when the passenger physically travels to Finland while **booking** adds another time dimension based on the **lead-time** when flight ticket was booked. For example, a flight ticket can be booked 1 week, 1 month or even 6 months prior to departure.

This is the main strength of the flight bookings data: it provides an outlook up to 6 months into the future on how many flight tickets were booked to Finland from each country of origin. In addition, it provides historical data on how many tickets were on the book 1 year ago, 2 years ago. The main weakness is related to completeness as direct bookings from airlines is unknown. In case flight bookings increase or decrease, it can indicate either a change in tourism volume **or** a change in the ratio of tickets sold through GDS versus directly by airlines.

⁷ <https://forwardkeys.com/>

METHODS AND MODELING

IMPUTATION METHODS TO DEAL WITH MISSING DATA

The adopted nowcasting strategy follows the idea of Fornaro, Luomaranta and Saarinen (2017) who use the time series structure of the variable of interest together with the additional information obtained from the microdata which is already available at the time of the nowcast is made.

As shown in Figure 1, a large share of data is already available at $t+10$ (around 80% of overnight stays). This constitutes a part of the final index that does not need to be estimated, denoted as x_t . We then also have an unknown part in the series, which corresponds to the replies that were not yet received. This is the part that needs to be imputed, denoted as \hat{x}_t . Now the final index can be formulated as

$$\hat{y}_t = x_t + \hat{x}_t \quad (1),$$

where \hat{y}_t is the final nowcast.

The two indices x_t and \hat{x}_t are obtained by summing over at each point in time the individual responses of the various establishments covered in the dataset, where the observation at time t is missing for \hat{x}_t . To obtain an estimate of the latter we use an automated ARIMA where we include the index x_t as external predictor (ARIMAX). For example, if we use a simple AR (1) process with an external predictor the model can be written as

$$\hat{x}_t = \mu + \phi_1 \hat{x}_{t-1} + \beta x_{t-1} + \epsilon_t \quad (2)$$

and the predicted value for \hat{x}_t would then be obtained by

$$\hat{x}_{t,pred} = \hat{\mu} + \hat{\phi}_1 \hat{x}_{t-1} + \hat{\beta} x_{t-1}$$

where the various parameter estimates are obtained by OLS (ordinary least squares). Once we have a predicted value for \hat{x}_t we can produce the nowcast of y_t by using (1).

We have explored various other options to impute the establishment data, such as the regularized iterative principal component (PCA) method described in Husson and Josse (2016) and implemented using the missMDA package for R, as well as using a simple imputation approach where we impute the value at time t for an establishment using the overnight stays at the same establishment one year before (this technique is similar to the one currently used by Statistics Finland). Finally, we tried taking the simple average of the three nowcast provided by the different imputation techniques.

MODELS FOR FORECASTING SHORT-TERM TOURISM VOLUMES

In order to form our set of predictions we use a combination approach similarly to Fornaro & Luomaranta (2019), where we combine the nowcasts obtained by a set of statistical models and machine learning techniques.

The set of predictors includes the lags of the target variable, and the latest observations of the flight bookings data. This data is the booking information of a certain nationality concerning the target month of interest, with different lead times included as separate variables. For example, for month t we have information on bookings made for flights to occur at $t+1$, $t+2$ up to $t+6$. Given that the techniques we rely on handle well a relatively large number of predictors, as well as the fact that working with monthly variables ensures us to have an adequate number of observations to estimate the models, we include all bookings variables, irrespective of the forecast horizon considered. Finally, we include a set of monthly dummies, D_{it} , which take value 1 if month t corresponds to period i .

For example, if we consider a simple linear regression framework, our model can be written as:

$$y_{t+h} = \beta_0 + \beta_1 y_t + \beta_2 \text{bookings}_{t+1} + \sum_{i=1}^{11} \gamma_i D_{it+h} + \epsilon_t. \quad (3)$$

Notice that bookings information tends to be available with a shorter lag compared to the overnight stays, which explains the discrepancy in the time index seen in model (3). Moreover, the formulation in (3) indicate that we are using a direct forecasting approach for predicting multiple periods ahead, instead of the iterative approach seen in other applications. Finally, the faster availability of flight bookings information means that we can use our models to obtain nowcasts of overnight stays.

It is important to underline that specification (3) is just an example, and that the forecasting formulation can change a lot depending on the technique we are using. As it will be seen in the results, we find that taking a simple average of the forecasts obtained by the different models tends to give a better forecasting performance compared to the individual models.

PRELIMINARY RESULTS

In this chapter we present the results separately for the nowcasting and forecasting components of the tourism forecasting model. Finally, we combine these components into an integrated tourism forecasting model that provides rolling forecasts for inbound tourism up to 6 months ahead.

The accuracy of the tourism forecasting models are evaluated using mean absolute percentage error (MAPE), which is defined as:

$$MAPE = \frac{100\%}{T} \sum_{t=1}^T \left| \frac{A_t - F_t}{A_t} \right|$$

where A_t is the actual nights spent and F_t is the forecasted nights spent for each period t . MAPE measures how much the forecast deviates from the final published nights spent in percentage terms, on average.

NOWCASTING BASED ON ACCOMMODATION SURVEY

The nowcast was performed for each combination of month and time lag ($t+5$, $t+10$, $t+15$) for the period between January 2015 up to October 2019. The nowcasting model uses the real nights spent data that was available, for example by 5th day in case of $t+5$. The model imputes missing nights spent for those accommodation establishments that did not respond by this date.

Four different imputation methods, described in the previous sections, were tested:

1. Simple method: use data from previous year and same month (for example use January 2018 data for January 2019)
2. MissMDA
3. ARIMA with external predictors (ARIMAX)
4. Combination: average of all 3 above

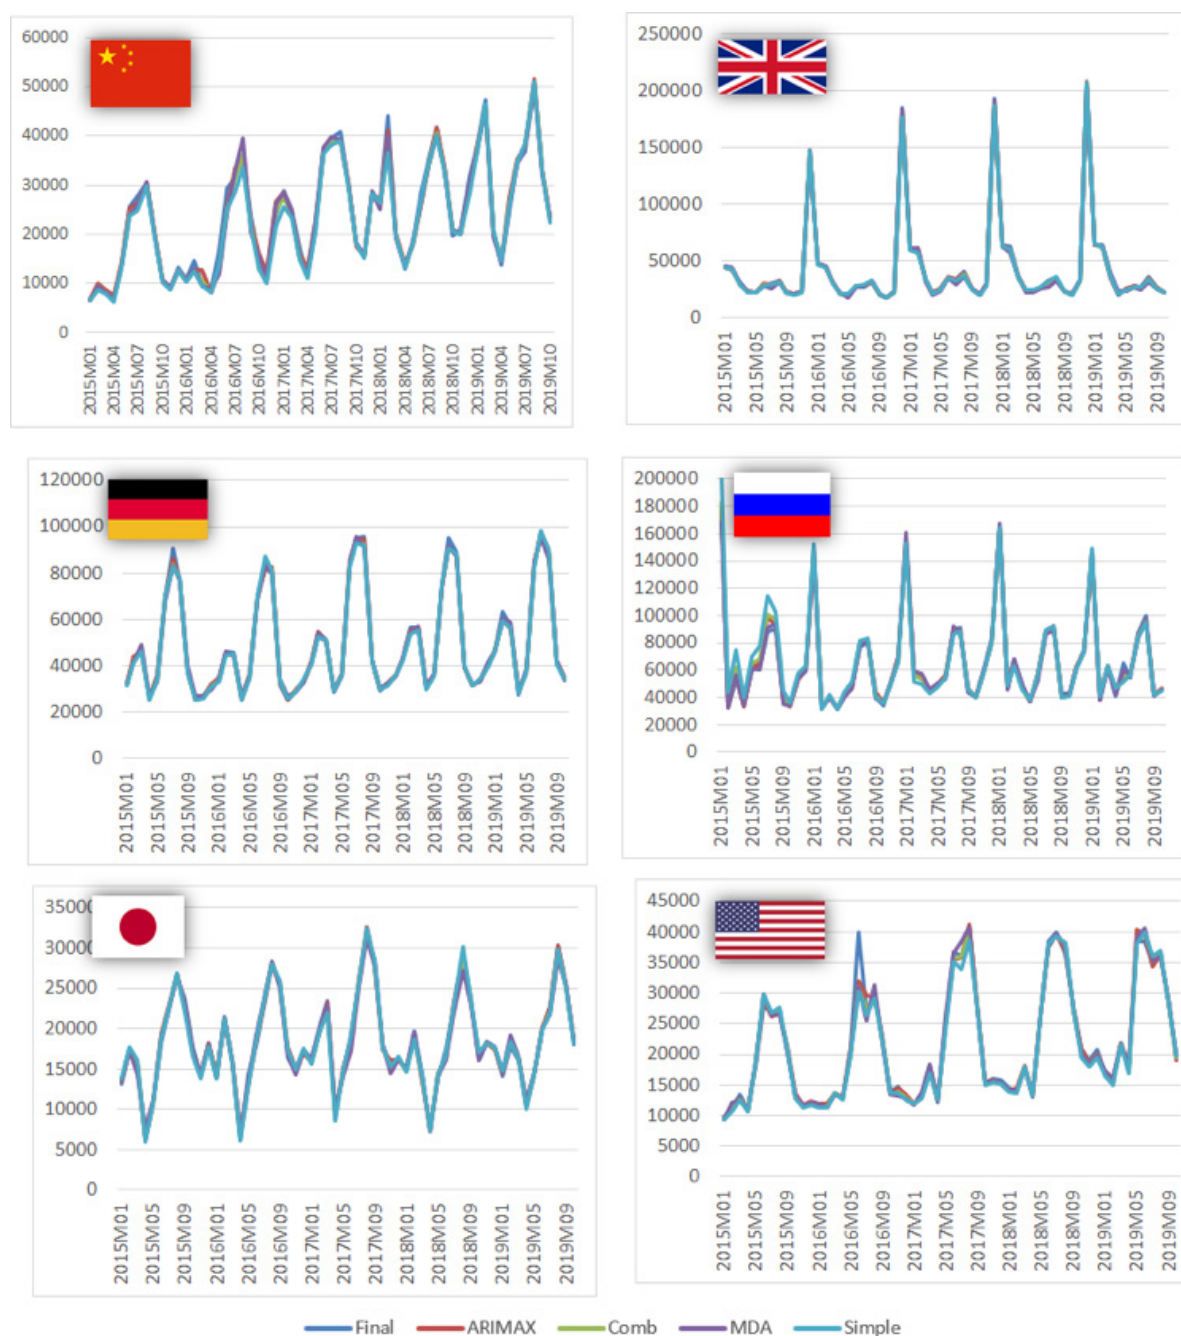
The first published preliminary statistics were also included as it also deviates slightly from the final statistics. This gives the possibility to benchmark the models also against the preliminary statistics which are typically published with a time lag of 20 – 25 days. We report the nowcasting performance of the various models in Table 1 and look at the historical performance for individual countries in Figure 4.

	Average	Domestic	Foreign	Russia	Germany	UK	USA	Japan	China	Korea	Hong Kong
Preliminary	1,3 %	0,9 %	1,2 %	1,6 %	1,1 %	1,1 %	0,8 %	1,0 %	1,2 %	1,2 %	2,3 %
t+5 Simple	6,5 %	1,7 %	3,4 %	11,1 %	4,6 %	5,1 %	5,5 %	4,9 %	9,3 %	8,0 %	11,7 %
t+5 MDA	6,4 %	1,9 %	2,9 %	4,6 %	4,4 %	9,4 %	4,1 %	6,6 %	7,0 %	9,5 %	13,4 %
t+5 ARIMAX	4,6 %	1,0 %	1,7 %	4,5 %	3,5 %	3,6 %	4,9 %	4,3 %	5,8 %	7,3 %	9,6 %
t+5 Comb	4,7 %	1,1 %	2,2 %	4,9 %	3,2 %	4,6 %	3,6 %	3,9 %	5,9 %	6,8 %	10,5 %
t+10 Simple	4,2 %	1,1 %	2,0 %	7,3 %	2,6 %	2,9 %	3,7 %	2,7 %	5,5 %	6,1 %	7,5 %
t+10 MDA	3,8 %	1,1 %	1,5 %	2,9 %	2,5 %	4,9 %	3,0 %	3,2 %	4,6 %	6,8 %	7,4 %
t+10 ARIMAX	3,2 %	0,8 %	1,2 %	3,2 %	1,9 %	1,9 %	2,5 %	2,6 %	5,4 %	4,7 %	7,5 %
t+10 Comb	3,0 %	0,8 %	1,3 %	3,6 %	1,9 %	2,7 %	2,4 %	2,4 %	4,0 %	4,6 %	6,8 %
t+15 Simple	3,5 %	1,0 %	1,7 %	5,7 %	2,2 %	2,3 %	2,8 %	2,0 %	4,5 %	5,3 %	7,6 %
t+15 MDA	3,2 %	0,9 %	1,3 %	2,4 %	2,1 %	4,3 %	2,6 %	2,6 %	4,0 %	5,6 %	6,2 %
t+15 ARIMAX	2,7 %	0,7 %	1,0 %	2,7 %	1,7 %	1,7 %	2,7 %	1,8 %	4,2 %	3,9 %	7,1 %
t+15 Comb	2,6 %	0,7 %	1,1 %	2,8 %	1,6 %	2,3 %	2,2 %	1,6 %	3,4 %	3,8 %	6,4 %

Table 1 - MAPE for each model and time lag

The colors in the table indicate the relative accuracy of each model for a certain country. For example, the first published preliminary foreign nights spent deviate 1,2 % from the final foreign nights. If the foreign nights were estimated after 5 days (t+5) using ARIMAX method, the deviation would be 1,7%. Interestingly, the foreign ARIMAX nowcast is slightly more accurate than the preliminary statistics.

In general, the most accurate results are obtained by using either ARIMAX method or by using a combination of all methods. In terms of time lag, there is a significant improvement in accuracy going from t+5 to t+10 nowcast. The accuracy improves only slightly by t+15. This is expected, as nearly 80% of nights spent have been collected by t+10 and 85% by t+15.

Table 1 - MAPE for each model and time lag

The visual inspection of the $t+10$ monthly nowcasts for the main countries also confirms the accuracy of the models. The nights spent by tourists from each country have a specific monthly seasonality which the models accurately follow. There are only a few exceptional periods, where the nowcasts significantly deviate from the final figures. These exceptions can take place if certain significant accommodation establishments have not responded by $t+10$ and they have a great contribution to the nights spent for that country.

To conclude these results, the recommended nowcasting model should be ran at roughly $t+10$ days and the model should be based either on ARIMAX only or a combination of simple, MissMDA and ARIMAX methods.

MONTHLY FORECAST BASED ON FLIGHT BOOKINGS

The forecasting component estimates the nights spent for upcoming months based on flight bookings. A basic ARIMA model was benchmarked against the combination of nowcasts produced by 15 different models. The results of the forecasting exercise are shown by first taking a case study of a 1 month forecast for Russian tourism, Figure 5, and then by providing the full results for all countries and all forecast periods up to 6 months.




Figure 5 - Final data and 1 month ahead forecast for nights spent by Russian tourists

The graph presents the final data for Russian tourists and a 1 month ahead forecast using ARIMA or a combination of 15 other models. The results from the models are shown as lower decile (10%), mean and upper decile (90%). The final statistics on Russian nights are typically between the upper and lower decile of the forecasts. This makes the mean often close to the final figure. The ARIMA model often produces too low estimates as the ARIMA forecast is in some cases even below the lower decile of the other models. The average error (MAPE) for ARIMA is 28 % while for the mean of 15 models the average error is 17 %.

Table 2 - Forecast errors using flight bookings data

Country	Model	Mean Absolute Percentage Error (MAPE)					
		1 month	2 months	3 months	4 months	5 months	6 months
All countries	15 models	7 %	8 %	8 %	9 %	10 %	10 %
	ARIMA	9 %	10 %	11 %	12 %	12 %	13 %
	Improvement	2 %	2 %	3 %	2 %	2 %	3 %
China	15 models	17 %	16 %	16 %	17 %	19 %	21 %
	ARIMA	36 %	39 %	34 %	27 %	24 %	30 %
	Improvement	19 %	23 %	18 %	10 %	5 %	10 %
Germany	15 models	7 %	8 %	9 %	9 %	10 %	10 %
	ARIMA	11 %	11 %	11 %	12 %	12 %	12 %
	Improvement	3 %	3 %	2 %	3 %	2 %	2 %
Japan	15 models	12 %	12 %	10 %	14 %	22 %	21 %
	ARIMA	14 %	14 %	14 %	14 %	15 %	14 %
	Improvement	1 %	2 %	4 %	0 %	-7 %	-7 %
Korea	15 models	18 %	18 %	18 %	19 %	21 %	23 %
	ARIMA	23 %	23 %	22 %	22 %	24 %	24 %
	Improvement	5 %	5 %	4 %	3 %	3 %	1 %
Russia	15 models	17 %	24 %	21 %	23 %	25 %	28 %
	ARIMA	28 %	31 %	30 %	30 %	31 %	34 %
	Improvement	12 %	7 %	9 %	7 %	6 %	7 %
United Kingdom	15 models	12 %	11 %	12 %	12 %	11 %	12 %
	ARIMA	19 %	23 %	23 %	21 %	19 %	17 %
	Improvement	8 %	12 %	11 %	9 %	7 %	5 %
United States	15 models	10 %	10 %	12 %	11 %	12 %	13 %
	ARIMA	11 %	11 %	11 %	11 %	12 %	13 %
	Improvement	1 %	1 %	-1 %	0 %	0 %	0 %

Using the mean of 15 models improves the forecast accuracy over a standard ARIMA model for most countries and forecasting horizons. The improvement is typically 1 to 10 percentage points but can be as big as 20 as in the case of China. The short-term (1 - 3 months) forecasts are in general more accurate than the longer term (4 - 6 months). Different weighting schemes for the 15 models were also tested but this did not yield any significant improvements over a standard arithmetic mean.

The overall accuracy of the monthly forecasts is average at best. A forecast error of less than 10 per cent is achieved only for Germany and all countries in total. The forecast accuracy is rather poor for Russia, China and Korea. Tourism from these countries has been very volatile especially during 2015 - 2017. This volatility affects also the forecast accuracy for these countries.

As a conclusion, the monthly forecasts based on flight bookings are significantly less accurate than the nowcasting component. While the nowcasting accuracy is often within 2 - 3 per cent of the final data, the forecasting is within 7 - 10 per cent at best. For some countries the forecasting error can be up to 25 per cent off. Further fine tuning of the model is needed in order to reach a more acceptable range of accuracy as presented in the next chapter. However, it is important to underline that bookings data bring a clear advantage, in terms of forecasting performance, compared to a standard ARIMA setting.

TOURCAST: A TOURISM FORECASTING MODEL

The tourism forecasting model, “Tourcast”, proposed here does not forecast the nights spent for each individual month but instead it provides estimates for seven cumulative periods each month. The cumulative estimates presented in this chapter were calculated by summing up the forecasts for individual months.

The cumulative periods are:

- nowcast for the previous month (t+10 days based on accommodation survey)
- nowcast + 1 month forecast
- nowcast + 2 months forecast
- ...
- nowcast + 6 months forecast

In other words, the users will get a sense of how tourism flows changed last month and will develop up to 6 months into the future. There are at least three strong arguments in favor of using cumulative forecasts as opposed to forecasting individual months:

1. The changes in nights spent and forecast errors are greater for individual months than for an aggregated period of more than one month.
2. The nowcast is very accurate and the forecast accuracy is improved by including the nowcast as the first month of the period.
3. The tourists arrive at a certain month but some of the nights spent may take place in the following months. Using a rolling period mitigates this issue considerably.

We report, in Figure 6 and 7, the 12-month moving average of the cumulative forecasts obtained by the forecast combination approach.




Figure 6 - Rolling 12 months forecast error across all 7 countries of origin

The graph highlights the benefits of using a cumulative forecast instead of forecasting individual months for all 7 countries of origin. Even the longest period, nowcast + 6 months, has a maximum error of 8 per cent, while the forecast for the individual month 6-step ahead has an error of up to 16 per cent.

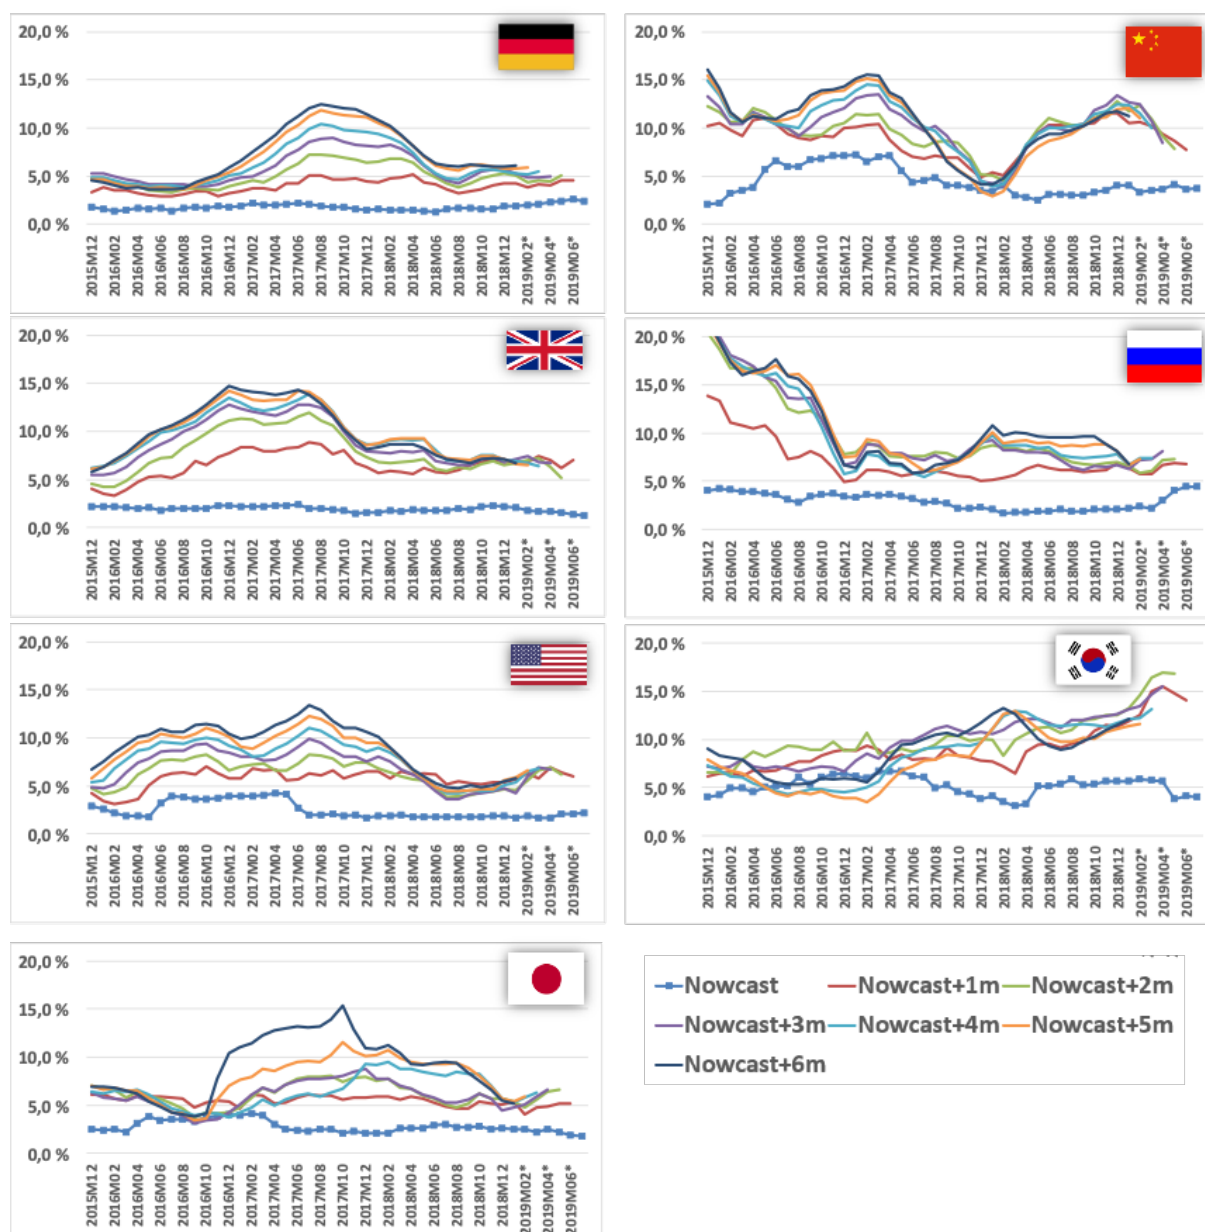


Figure 7 - Rolling 12 months forecast error for each country of origin

In most cases, the forecast is less accurate for the longer time horizons. This is clearly seen in the case of Germany, where the 1-month forecast is always more accurate than the 2-month forecast, which is in turn more accurate than 3-month forecast etc. Some countries are more difficult to forecast than others, for example China and Russia, where even the 1-month error ranges between 5 to 15 per cent. The tourism from these two countries have been highly volatile in the past. At the other end, countries with a more stable and predictable tourism seasonality, such as Germany, can be forecasted with less than 5 per cent error for the 1-month ahead period. To support this fact, we plot the relation between forecast error and the variation in tourism flows of different countries, in Figure 8.

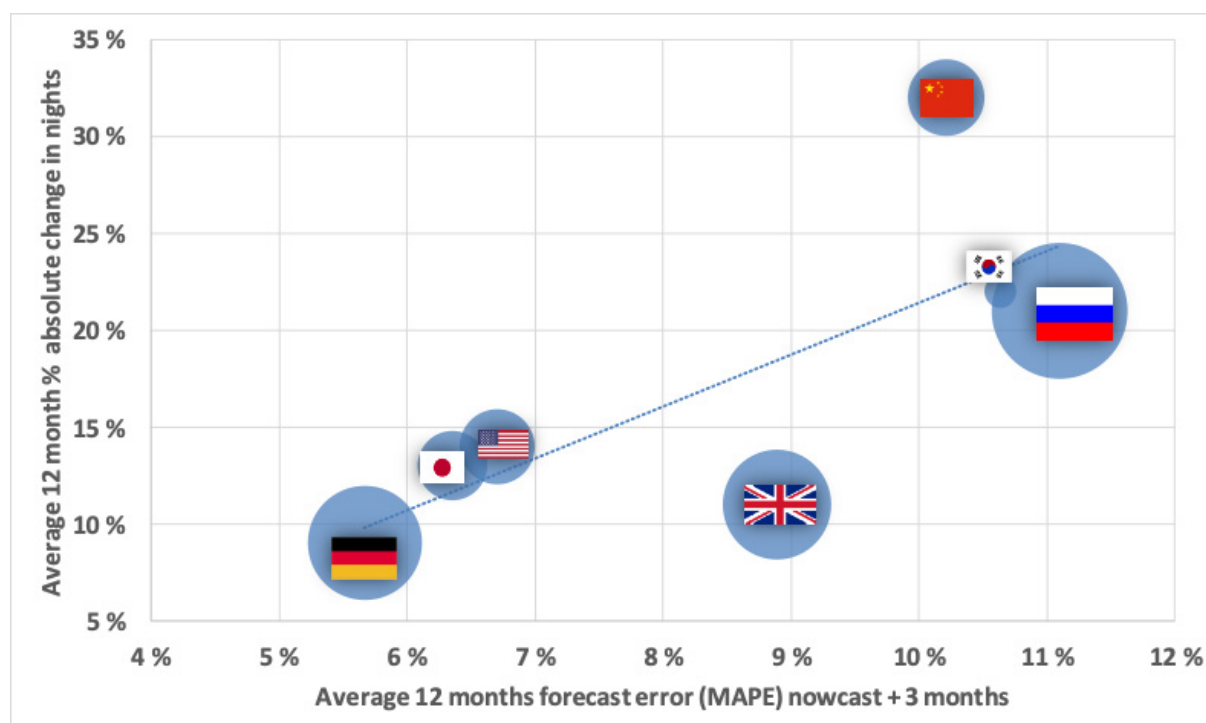


Figure 8 - Correlation between nights spent volatility and forecast error (nowcast + 3 months)

The graph illustrates the correlation between the change of nights spent and the forecast error (nowcast + 3 months) for the period between January 2015 – July 2019. The dotted line is the expected result in case of perfect correlation. The circle size indicates the number of nights spent.

The basic rule is that countries with bigger changes in nights spent are also more difficult to forecast. There are also some interesting exceptions to this rule, as three countries are clearly outside the correlation line: China, Russia and United Kingdom.

The forecast performs better than expected for China and to some degree also for Japan and United States, as these countries are above the correlation line. The main hypothesis is that most tourists from these countries arrive to Finland by air making the flight bookings a more representative sample than for more nearby countries.

Most of the Russians arrive using other modes of transport than air. There was also a rapid downturn in Russian tourism in early 2015 which can be considered an external shock that the forecasting model cannot predict.

The tourism from United Kingdom is unique in the sense that more than one third of UK nights are spent in December, compared to only 10% for other countries. Due to this extremely high seasonality of UK tourism, it is also more difficult to forecast.

We finally report the MAPE for cumulative forecasts of tourism flows from different countries of origin, with the out-of-sample period going from January 2015 up until July 2019, in Table 3.

Table 3 - Summary of mean and maximum monthly MAPE per country

Country	Indicator	Mean Absolute Percentage Error (MAPE)					
		Nowcast+1m	Nowcast+2m	Nowcast+3m	Nowcast+4m	Nowcast+5m	Nowcast+6m
All countries	Mean	3,3 %	4,1 %	4,2 %	4,4 %	4,3 %	4,4 %
China		9,0 %	9,7 %	10,2 %	10,8 %	10,9 %	11,4 %
United Kingdom		6,5 %	7,9 %	8,9 %	9,6 %	9,9 %	10,0 %
Germany		3,9 %	5,0 %	5,7 %	6,1 %	6,7 %	7,0 %
Japan		5,6 %	6,5 %	6,4 %	6,6 %	7,7 %	8,4 %
Korea		9,5 %	11,0 %	10,6 %	9,3 %	8,8 %	9,9 %
Russia		7,7 %	10,7 %	11,1 %	10,9 %	11,4 %	11,2 %
United States		5,7 %	6,1 %	6,7 %	7,4 %	7,8 %	8,2 %
All countries	Max	11,4 %	10,4 %	11,2 %	11,4 %	12,0 %	12,6 %
China		25,2 %	28,8 %	30,0 %	27,0 %	28,6 %	30,0 %
United Kingdom		8,9 %	12,0 %	12,8 %	13,9 %	14,2 %	14,8 %
Germany		10,2 %	12,4 %	13,8 %	14,8 %	15,3 %	15,5 %
Japan		18,8 %	22,5 %	20,5 %	21,7 %	31,9 %	48,8 %
Korea		33,6 %	42,8 %	32,1 %	29,0 %	26,8 %	28,9 %
Russia		31,8 %	35,2 %	35,5 %	37,2 %	37,8 %	42,0 %
United States		18,5 %	19,2 %	20,5 %	20,0 %	20,4 %	21,0 %

The table summarizes the accuracy of Tourcast in terms of mean and maximum forecasting error. The color scale ranges from green (good accuracy) to red (poor accuracy). Even for the countries where the forecast accuracy is good on average, there have been some periods with poor accuracy as indicated by the maximum error.

As a conclusion, for the period from January 2015 to July 2019, the tourism forecasting model was able to provide accurate estimates for the short- (1 - 3 months) to mid-term (4 – 6 months) forecast for the aggregate of all countries.

An acceptable level of accuracy was achieved for

- Short-term forecast for Germany, UK, Japan and USA
- 1 month forecast for UK and Russia

The following estimates are not accurate enough, additional data sources may be needed to further improve the model for:

- Short- to mid-term forecast for China and Korea
- Mid-term forecast for Germany, UK, Russia, Japan and USA

It should also be pointed out, that a sufficient level of forecast accuracy is highly subjective. For some purposes, a forecast error of more than 5 per cent may be acceptable but for others it could be far too inaccurate. The conclusions presented above reflect the subjective levels of performance that we are happy with considering the project objectives.

CONCLUSIONS

The project has provided evidence that it is possible to combine existing and new data sources with machine learning and statistical techniques in order to gain a grasp of the developments of tourism flows in the immediate past and into the future. The recommended approach is to compile cumulative forecasts that combine nowcasting based on accommodation survey with up to 6 months forecast based on flight bookings to Finland. This approach provides significantly better accuracy than forecasting individual months.

The Tourcast model can provide a very accurate nowcast and relatively accurate short-term (1 to 3 months) forecasts for nights spent of most countries of origin. For the aggregate of all countries, the model can even fairly accurately predict the mid-term (4 to 6 months) nights spent.

By the nature of the models, external shocks that do not appear in flight bookings will result in poor forecast accuracy, until the model has adjusted to the new level of nights spent, if the shocks are persistent. Such downward shock occurred for example in Russian tourism to Finland in 2015. The shocks may be mitigated by introducing other data sources at a later stage such as currency exchange rates or visas granted that might predate the change in nights spent of Russians.

These are the requirements for establishing Tourcast as a regular monthly production process:

- Data sources
 - Accommodation survey data: nights spent at establishment level by country by month
 - Published accommodation statistics; preliminary and final data by country by month
 - API access to flight bookings data
- System requirements
 - SAS for preparing the accommodation survey data and compiling the results
 - R-programming language for running the nowcasting and forecasting models

The main objective for 2020 is to set up a production process that provides regular monthly updates of the current Tourcast model. In addition, several paths for future development of Tourcast have been identified. These may include:

- Regional level forecasts:
 - 4 main tourism regions of Finland (Helsinki metropolitan area, Lakeland, Coast and Archipelago, Lapland)
 - 19 regions of Finland
 - Long term forecasts (6 to 12 months) by using flight capacity data that is also available via API but not used in the current project
 - Improving forecast accuracy by using additional data sources in modeling: hotel bookings data, granted visas, traffic sensor data, mobile positioning data etc.
-

REFERENCES

Fornaro, Paolo, Luomaranta, Henri & Saarinen, Lauri (10.1.2017). “Nowcasting Finnish Turnover Indexes Using Firm-Level Data”.

ETLA Working Papers No 46. <http://pub.etla.fi/ETLA-Working-Papers-46.p>

Fornaro, Paolo, Luomaranta, Henri (30.11.2019) “Nowcasting Finnish real economic activity: a machine learning approach” Empirical Economics, pp 1-17

Josse J, Husson F (2016) missMDA: a package for handling missing values in multivariate data analysis. J Stat Softw 70(1):1–31

Trevor Hastie, Robert Tibshirani, and Jerome Friedman. The elements of statistical learning: data mining, inference and prediction. Springer, 2 edition, 2009

