

# Linked Data for Cross-Domain Decision-making in Tourism

Marta Sabou, Adrian M.P. Braşoveanu, and Irem Önder

MODUL University Vienna, Austria

{marta.sabou, Adrian.brasoveanu, irem.onder}@modul.ac.at

## Abstract

In today's global economy, tourism managers need to consider a range of factors when making important decisions. Besides traditional tourism indicators (such as arrivals or bednights) they also need to take into account indicators from other domains, for example, economy and sustainability. From a technology perspective, building decision support systems that would allow inspecting indicators from different domains in order to understand their (potential) correlations, is a challenging task. Indeed, tourism (and other indicators), while mostly available as open data, are stored using database centric technologies that require tedious manual efforts for combining the data sets. In this paper we describe a Linked Data based solution to building an integrated dataset as a basis for a decision support system capable of enabling cross-domain decision-making. Concretely, we have exposed tourism statistics from TourMIS, a core source of European tourism statistics, as linked data and used it subsequently to connect to other sources of indicators. A visual dashboard explores this integrated data to offer cross-domain decision support to tourism managers.

**Keywords:** statistical data; decision support; linked data; TourMIS.

## 1 Introduction

Does the economic crisis influence the travelling behaviour of German tourists within Europe? Do (mass-)tourism arrivals to Spain impact the environment of the host country? In today's interconnected world, benchmarking, forecasting and, in general, decision-based activities of tourism managers cannot solely rely on tourism statistics (such as arrivals, bednights and capacity) but must also consider data from other domains. Understanding correlations between tourism and economic indicators, for example, can prepare tourism managers to plan their activities according to future financial predictions. Or, contrasting tourism and environmental indicators, as another example, can shed light on environmental impacts (e.g., increased CO<sub>2</sub> emissions in heavily touristic areas) potentially caused by (mass)tourism in host countries thus helping to avoid them in the future. Therefore, such cross-domain queries are an important decision support instrument for tourism decision makers.

Most tourism decision support systems usually only cater for investigating tourism indicators in isolation from economic or sustainability indicators, as we describe in Section 2.2. The main reason for this is that technologically, merging data from various sources of indicators is challenging and relies on much manual effort. Indeed, as we point out in Section 2.1, until recently, statistical data was made available using database centric technologies that require tedious manual efforts when combined. For example, a system that would allow correlating arrivals from Germany to Prague with the GDP of Germany, must be “told” through explicitly programmed mappings that the country “Germany” in one data source (say, TourMIS) is actually the same as the country “DE” in another data source (say, World Bank). Large-scale data integration,

therefore, remains challenging because (1) it requires significant manual effort to hardcode correspondences of the entities used by two data sources and (2) these manual mappings must be maintained to account for changes in the naming conventions used by the two datasets.

Semantic Web and Linked Data (LD) technologies have been developed to support the intelligent integration of data on the Web (Berners-Lee et al., 2001). Core to the idea of the Semantic Web is to represent information not as mere text or numbers, but rather in terms of structures that have a machine understandable meaning. Linked Data technologies provide a mechanism to publish this “intelligent” data on the Web and creating links between elements of different datasets, thus facilitating their integration. While originally developed for integrating facts derived from textual web pages, a new area of application for Semantic Web/Linked Data technologies, is the semantic representation and integration of statistical data. However, to date no significant amount of tourism statistics have been published as linked data.

To advance the state of the art we published TourMIS (Wöber, 2003), a core source of European tourism statistics, as Linked Data. TourMIS includes data from over 100 European cities, 33 European nations and 19 source markets. The main purpose of TourMIS is to present information for tourism decision makers. The data in TourMIS is provided by either the city tourism organizations or national tourism organizations. The data set consists of bednights spent at a destination (including all types of accommodations), arrivals at accommodations, capacity and occupancy rates. TourMIS has over 18000 registered users ranging from representatives of national, provincial, regional and city tourism organizations, which are involved in long-term, strategic planning of the tourism development of a region to tourism suppliers such as suppliers of accommodation, food, travel, as well as travel agencies and tour operators, which are mostly interested in local tourism demand forecasts. On average, they query TourMIS once every three minutes, thus leading to 200K queries per year.

Based on the created dataset we are creating a decision support system that allows investigating the correlation of indicators across several domains, in particular, tourism, economy and sustainability. The resulting LD dataset as well as the decision support dashboard have been built during the ETIHQ ([www.etihq.eu/](http://www.etihq.eu/)) project and therefore are named accordingly as the ETIHQ dataset and the ETIHQ dashboard. To the best of our knowledge ETIHQ is the first online tourism decision support system that connects linked data and tourism statistics from multiple data sources (e.g. Worldbank, Eurostat). The goal of this paper is to describe the technical solution for providing TourMIS as LD, to introduce the cross-domain decision dashboard and to reflect on the main lessons learned while using this novel technology for tourism.

## **2 Related Work**

We review existing tourism statistics sources and tourism decision support systems.

### **2.1 Sources of Tourism Indicators**

The ETIHQ repository is unique in terms of its content when compared to other sources of tourism statistics. Given their importance, tourism indicators are published by a variety of international organisations along with indicators in other areas such as

economy and sustainability. The UN provides a multitude of datasets, including also various tourism indicators from UN's World Tourism Organization (UNWTO), such as, arrivals, departures and tourist expenditures, measured per country and year. Much of this data is open, however, only provided in the proprietary Excel format. Additional UNWTO data is only available as PDF downloads. Eurostat provides a wealth of European level statistics, including various tourism indicators (capacity, arrivals, bednights, expenditures etc) for all European countries. Measurements are provided on a monthly basis. Data can be downloaded in a variety of formats including Excel, CSV, HTML, SPSS and PDF. The World Bank provides open access to over 8000 indicators ([data.worldbank.org](http://data.worldbank.org)), including also a variety of tourism indicators measured annually and at a country level. The available datasets can be downloaded in XML and Excel formats, as well as accessed through an API.

ETIHQ is comparable to the sources of tourism statistics described above, but it is more detailed and this makes its publishing beneficial to European tourism stakeholders. Indeed, while other sources provide annual measurements (except Eurostat), at country level, ETIHQ contains both annual and monthly measurements and it focuses on individual cities. Additionally, ETIHQ also identifies key markets based on tourists' origin, a feature not offered by any of the data sources we surveyed, although market information is essential for tourism promotion organizations in developing their international advertising campaigns.

We conclude that there are various sources that offer a multitude of indicators in the area of tourism and beyond. A general trend is offering this data as open data, primarily through downloading it in popular (non-semantic) encodings. For now, the publication of these datasets as LD has been primarily undertaken by third-parties rather than data owners, resulting in low-quality datasets that simply expose database tables into an RDF format without enriching them with domain-specific metadata or linking them to other datasets. Upon inspection, many of these datasets were slow or returned erroneous (or no) data.

## **2.2 Tourism Decision Support Systems**

Tourism decision support systems have been implemented in various areas such as destination recommendation systems. Some examples of previous research in this area include studies about case-based travel recommendations (Ricci, Cavada, Mirzadeh, & Venturini, 2006), travel decision styles and destination recommendations (Zins & Grabler, 2006), and creating adaptive recommender systems using neural net approaches (Mazanec, 2006). There are some destination recommendation systems that have been developed as well such as DieToRecs (Ricci, Fesenmaier, et al., 2006), TourBO (Franke, 2006) and MobyRek (Ricci & Nguyen, 2006). However, the studies related to travel recommendation systems are directed towards the consumer (i.e., the individual tourists) rather than the tourism managers therefore focusing on different content (i.e., destinations, touristic offers, events) than the ETIHQ decision dashboard. Other types of decision support systems such as yield management systems are available for tourism managers, which have been used primarily within the hospitality industry.

One example of a tourism decision support system is BASTIS ([bastis-tourism.info](http://bastis-tourism.info)), which combines different types of data under one online system. It has three parts, a

wiki page, a blog and a forum where individuals can get information about heritage tourism in Baltic Sea region. The data comes from EUROSTAT, TourMIS, Statistisches Bundesamt Deutschland (German Statistical Office) and Deutsche Zentrale für Tourismus (National Tourism Organization of Germany). The heritage sites include churches, castles, museums, natural sites and attractions, and national parks from 10 different Baltic countries, 13 source markets and more than 40 heritage sites. Each heritage site has a profile that has basic information regarding the location of the site, the number of visitors, source markets, satisfaction of visitors and the duration of their visit at the heritage site. Furthermore, the heritage sites can be benchmarked with each other and the information can be used by decision makers involved in heritage tourism in the Baltic Sea region. However, this system is limited since it covers the Baltic Sea Region alone, the information is only about tourism domain although it comes from different sources and the data is integrated by hard-coded data source connectors.

On the other hand, the TourMIS interface already supports a variety of decisions such as detecting trends in specific source markets, monthly development of European city tourism and best performing European cities according to bednights in comparison to the previous quarter, which would help Destination Marketing Organizations (DMO) to take appropriate actions. For instance, a DMO can use TourMIS and can decide to invest more in a source market (country of origin of tourists) that shows an increasing trend such as the Russian market. However, a shortcoming of TourMIS is that the data is limited to that contained in the system's database. Although TourMIS is a good tourism information system, the data does not include external factors such as economy of a source market or sustainability indicators that are linked to tourism industry, which can explain the trends as well. TourMIS is free to use and open to anyone; access to the website requires only a web-form based registration.

PATA ([mpower.pata.org/](http://mpower.pata.org/)(Pacific Asian Travel Association) mpower is a tourism decision support system regarding the Asia Pacific countries. This system is a closed one; the detailed information is only accessible to members, which requires a fee to be paid. Its members include 80 national and regional tourism offices, more than 30 major airline companies, major hospitality groups and travel agencies. The indicators that are available are visitor information (e.g. number of international arrivals, gender, mode of transport, etc.), expenditure, source markets, aviation information including flight frequency, seat availability and air passenger traffic, accommodation (e.g. occupancy, number of available rooms, RevPAR (revenue per available room), etc.) digital indicators (e.g. top 10 travel websites, time spent on websites, % of internet users, etc.) and forecasts of arrivals and expenditure.

From the above we conclude that (1) the majority of tourism DSS are aimed at tourists rather than tourism managers; (2) while some tourism DSS systems aimed at managers include data from other data sources, these are from the tourism domain and are integrated in a hard-coded manner. All reviewed systems focus on inspecting statistical indicators from the tourism domain alone thus leaving an important market for cross-domain decision support systems.

### 3 A Brief Introduction to Semantic Web and Linked Data

Semantic Web and Linked Data technologies aim (1) to enrich raw data with semantic (that is, machine processible) information and (2) to publish it using Web based languages on the Web thus allowing links between datasets to be explicitly specified.

Core to the idea of the Semantic Web are ontologies (Gruber, 1993), formal domain models describing concepts in a domain and their relationships in ways that machines can process and reason with these descriptions. For example, a tourism ontology would describe concepts such as “City” or “Market”. Data items (e.g., Prague) are then described in terms of ontology concepts (e.g., by associating Prague to the concept “City”). Since the goal of the Semantic Web technologies is to make data public on the Web, ontology elements as well as each raw data element to be described in the Semantic Web are assigned a unique web URL, for example, <http://data.etihq.eu/etihq/page/dataset/Prague>, for the data element Prague or <http://data.etihq.eu/etihq/page/ontology/City> for the “City” concept of the ontology. The structure of the URLs usually indicates the name of the dataset (in our case “etihq”) as well as the type of the entity, which can be either part of the abstract data model (i.e., the ontology) or of the dataset.

While relational databases rely on a relation (i.e., table like) data model, the Semantic Web adopts a triple (or graph based) model with all data being represented as triples. For example, to declare that Prague is a City, a triple is created stating that < Prague, isA, City>. To encode triple-based Semantic Web data, a set of languages have been developed, most notably RDF (Resource Description Format) and OWL (Web Ontology Language). The RDF representation of the triple above is:

```
<ebo:City rdf:ID="eds:Prague" />
```

where the ebo and eds prefixes point to the “<http://data.etihq.eu/etihq/page/ontology>” and “<http://data.etihq.eu/etihq/page/dataset/>” namespaces respectively.

One advantage of exposing datasets online using Web based standards is that links can be created between the elements of these datasets (similarly to hyperlinks in HTML). Most often, an “owl:sameAs” link is created between URL’s from different datasets that represent the same real-life entity. For example, a single triple is sufficient to declare that two URLs in the ETIHQ and WorldBank data respectively refer to the same entity - Germany:

```
eds:Germany owl:sameAs  
<http://worldbank.270a.info/classification/country/DE.html>
```

Thanks to this link it is now possible for a computer program to understand that the term “Germany” as used in TourMIS is the same as the term “DE” used in the WorldBank dataset and that data about these terms can be combined.

### 4 Creating the ETIHQ Linked Data Dataset

To create the ETIHQ dataset from TourMIS we executed the following steps:

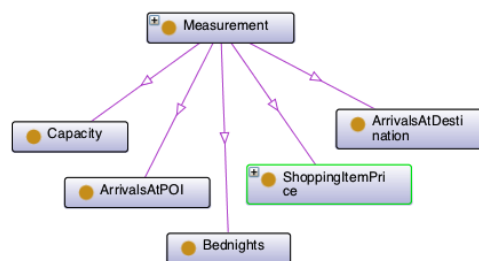
- *Data Cleaning* – This step prepares the raw data for being published as Linked Data and it is especially needed for legacy databases.

- *Semantic Modelling* – Selecting vocabularies, creating ontologies and a structure for the dataset.
- *RDB2RDF conversion* – Databases are transformed into triples using the process defined in the Semantic Modelling step.
- *Interlinking* – Linked Data offers simple mechanisms, such as links, to extend a dataset with connections towards other linked datasets.
- *Linked Data Interface Publishing* – exposes the data of the Web both in a way that it can be browsed by people and queried programmatically.

#### 4.1 Semantic Modelling

The semantic structure of the proposed solution includes a variety of ontologies:

- RDF Data Cube (QB - [www.w3.org/TR/vocab-data-cube/](http://www.w3.org/TR/vocab-data-cube/)) - is the emerging semantic standard for modeling data of statistical nature and therefore it was used as a main vocabulary to publish the TourMIS data;
- PROV-O ([www.w3.org/TR/prov-o/](http://www.w3.org/TR/prov-o/)) - is the standard vocabulary for specifying provenance information and its concepts were reused to allow representing information about who and when contributed data to TourMIS;
- ETIHQ Domain Ontologies - we also developed a set of ontologies specific to the tourism domain that complement the generic models above with concepts describing the content of TourMIS.



**Figure 1:** Main concepts in the Base Ontology.

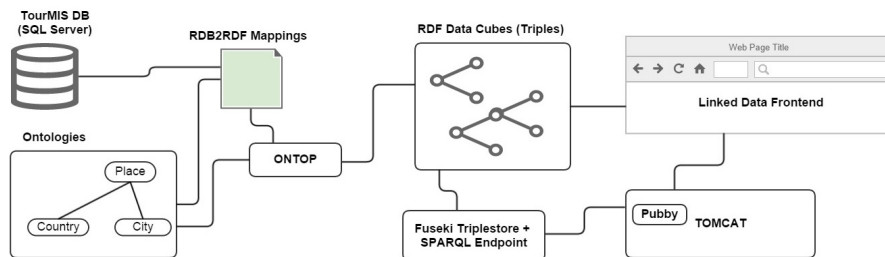
The ETIHQ domain ontologies describe three major areas of TourMIS. The base ontology has the “<http://www.etihq.eu/ontology/>” namespace (prefix ebo that stands for "etihq base ontology") and has been developed as an extension of the TourMIS ontology described in (Sabou et al., 2013). Core to this base ontology is the concept of *Measurement* which has five more specific concepts corresponding to all the statistical indicators in TourMIS and depicted in Figure 1. Additionally, the ontology models classes such as *Destination*, *Market* and *Currency*. Two further class hierarchies, one corresponding to *PointOfInterest* and one to *ShoppingItem*, are imported from two corresponding ontologies. The second ontology is the POI ontology declared under the <http://www.etihq.eu/ontology/poi/> namespace which models concepts corresponding to the POI categories (e.g., museums, churches, theatres) in TourMIS as specialisations of the *PointOfInterest* concept. Finally, the

Shopping ontology (<http://www.etihq.eu/ontology/shopping/> namespace) defines the *ShoppingItem* concept and its main subclasses as available in TourMIS and categorised in four main categories, namely, *Accommodation*, *Entertainment*, *FoodAndBeverage* and *Transport*.

As part of the semantic modelling step we also specified the datasets that should be created from TourMIS, in terms of the structure required by the QB Vocabulary. Five distinct datasets were created corresponding to arrivals, bednights, capacity, arrivals at points of interest and shopping statistics respectively. The dimensions of these datasets as well as the type of their measurements have been specified using concepts from the ETIHQ ontologies.

## 4.2 RDB2RDF Conversion and Linking

The TourMIS database is powered by a Microsoft SQL Server. In order to translate it into Linked Data, we used an RDB2RDF mapping (Relational to RDF mapping). The mappings file contains all the SQL queries that are used to create the RDF triples together with the rules used to create the right types and URIs for the RDF resources. These files are used by the Ontop Protégé plug-in to generate RDF files (Figure 2).



**Figure 2:** Overview of the main tools used for extracting TourMIS data and publishing it as Linked Data.

We wrote 60 mappings to generate 5 different datasets from TourMIS amounting to around 16 million triples. The RDF files generated by Ontop are in the RDF Data Cube (QB) format, therefore the values for one month or one year are mapped to the values of observations. The resulting triples are uploaded to the Fuseki triplestore (which a dataset specialized for storing RDF style data).

As a final step of the dataset creation process we specified links between the entities of ETIHQ and other datasets. We used the Silk link discovery tool (Isele et al, 2010) to create links to the DBpedia ([dbpedia.org](http://dbpedia.org)) and Geonames ([www.geonames.org](http://www.geonames.org)) datasets. We successfully linked all entities referring to countries and cities.

## 4.3 Linked Data Publishing

Once the data from TourMIS has been extracted, represented as RDF and stored in a specialized dataset, we published it on the Web. To that end, we provided a SPARQL frontend using YASGUI so that users can query the data. SPARQL is a query language similar to SQL in databases but adapted to query the graph like data model of RDF. Besides allowing query-based access, the ETIHQ data (both ontologies and

data triples) is exposed in a Web-page like format using the Pubby application (deployed in Tomcat). Thanks to this, users can navigate the online linked data repository using any web browser. Several online access points to our dataset can be accessed at <http://data.etihq.eu>.

#### 4.4 Linked Data Access Approaches

We experimented with several different approaches for collecting, querying and aligning the data from various sources. Firstly, we used federated SPARQL, where each data source is queried via a SERVICE request, and a federated SPARQL engine (e.g., Jena ARQ) aggregates the results. This approach often leads to queryTimeouts and the data alignment has to be encoded in the query itself. Secondly, we stored all data dumps in local repositories and queried them using SPARQL. Thirdly, we indexed the data from different domains using a search server (ElasticSearch, Sindice, Lucene, etc) and created a Search API that gets the data from all the sources. This can be done via custom scripts or plugins for the respective server. We preferred to create a custom Python script, as this allowed to also align the various data sources and bring the data into a common format that can be easily visualised. Using this approach, data is queried using server specific JSON-based query. We are considering opening this search API to external users under different open and commercial licenses, if the number of interested users is large. The fourth approach for querying the data involves the use of the visual dashboard (described in Section 5) that we developed on top of the Search API, as it comes with both visualizations and export functionalities (csv or Excel for the data; SVG or PNG for the charts). We are also considering offering RDF export functionalities in the future.

## 5 Visual Dashboard

The visual dashboard allows selecting and concurrently visualising tourism, economic and sustainability indicators, thus being, to the best of our knowledge, the first multi-domain decision support system in tourism. The current dashboard combines information from TourMIS, WorldBank and EuroStat. While TourMIS provides European level tourism indicators, from the other two sources we select economics and sustainability indicators (see Table 1).

Economic Indicators	Sustainability Indicators
GDP growth (annual %);	CO2 emissions (kt);
Inflation, consumer prices (annual %);	Forest area (% of land area);
Consumer price index (2005 = 100);	Roads, paved (% of total roads);
Official exchange rate (LCU per US\$, period avg);	Agricultural land (% of land area).
Unemployment rate, total (%of total labor force).	

**Table 1: Economic and Sustainability Indicators from the World Bank.**

Figure 3 displays a screenshot of the dashboard, which includes the following main panels: (i) the top horizontal interface bar allows selecting the time-frame of the analysis; (ii) the left-most frame can be used to define new queries over statistical data, displays these queries and allows their management (e.g., (de)selection); (iii) the top part of the middle pane contains a chart based visualisation element that visualises the currently selected queries; (iv) the lower part of the middle frame allows zooming



into data details for the selected search; (v) the right-most pane contains geographic data visualisations.

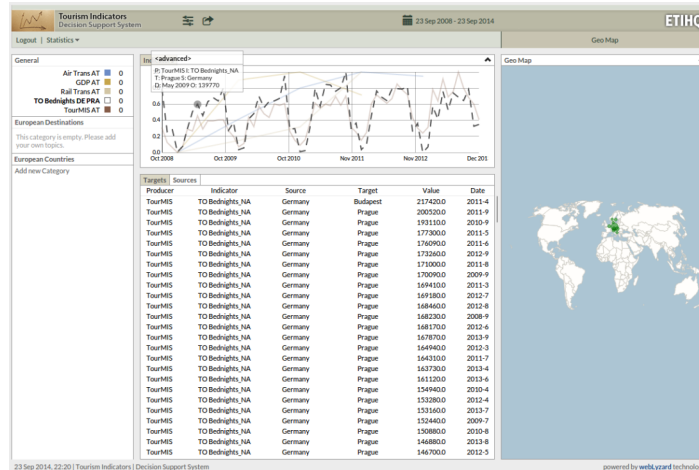


Figure 3: The ETIHQ Visual Dashboard (etihq.weblyzard.com).

The ETIHQ dashboard can be used by DMO managers to enhance their decision making and understand the big picture including economical factors that affect tourism demand. Consider the example of a DMO manager who wants to understand *whether and how the economic crisis influences the traveling behavior of German tourists in Europe*. He could follow a set of steps, which we now overview and exemplify how can be concretely performed with the Dashboard.





Figure 4: Search (topic) addition (left) and naming (right).

1) *Adding a new search*: the manager could start by searching for the number of beds reserved by German tourists in any European city, for example Prague. The definition of a search in the visual interface corresponds to the notion of a topic. A new topic is defined by pushing the wheel button and selecting the “Add Topic” menu item (Figure 4, left). A topic corresponds to an indicator, that is a slice of the data in the selected time interval with market (source) and destination (target) as fixed dimensions. Searches should be identified with meaningful names consisting of the data source of the indicator (TO stands for TourMIS), the name of the indicator (i.e., Bednights) and the dimension values that are chosen (in our example, these would be DE and PRA). Accordingly, we name the example search as “TO Bednights DE PRA” (Figure 4, right).

After adding and naming the search, we proceed to defining its content. For this, we hover over the new topic and press the wheel button that appears to its right. This action will replace the chart view in the top-middle pane of the interface with a dialog field that allows defining the topic (see Figure 5). It enables selecting the data source (currently, World Bank, Eurostat, TourMIS), indicators (the indicators from the menu), markets and destinations (both can be cities or countries). Once the relevant selections have been made, we choose “Save” which will close the dialog box.



**Figure 5: Defining the content of a search (topic).**

2) *Investigating the search results.* As soon as the topic definition dialog box is closed, the data related to this search is retrieved and visualized in the charts view (entitled “Indicators”). The first time a topic’s data is visualized, the corresponding trend line is a dashed line, as shown in Figure 3. The newly added topic also triggers various changes in the rest of the interface.

The data displayed in the tables (middle pane) changes. This pane will create as many sub-panes as the number of dimensions for the visualized indicators. For our example, the TourMIS Bednights indicator has two dimensions, namely source and target, so two corresponding panes will be created. The “Targets” table, keeps the source value fixed (in our case, “Germany”) and varies the values for the Target cities, thus displaying the number of German tourists going to all European destinations. The table can be sorted based on the value field, thus allowing to quickly identify the most/least popular destination for Germans – it appears for example that Budapest is a very popular destination for German tourists. Similarly, the “Source” table keeps the source fixed to “Prague”, but varies the source markets, thus allowing detecting those tourist groups that visit Prague the most/the least.

A click on the pane name will trigger a change in the Geo Map (right pane of the interface), which displays the tabular data visually. The data for a particular market is summed up (from months to yearly data), and a visual representation of the connections between markets and destinations in the form of arrows is created (bigger arrows mean more tourists in the selected interval). Following the steps above, the manager can investigate the behavior of German tourists to other European destinations as well, for example, Budapest.

3) *Adding a search over an indicator from another domain.* The previous steps allow exploring the visiting behavior of German tourists to Prague and other European cities. To understand whether this behavior correlates with the German economic situation, an economic indicator can be added as a new topic, such as the GDP Growth from World Bank (displayed as an interrupted line in Figure 6). Figure 6 super-imposes German GDP (from WorldBank) as well as Bednights to Prague and

Budapest, as these indicators have been selected for visualization in the “General” pane (the color on the right side of a topic corresponds to the graph color on the chart – e.g., light blue for German Bednights to Prague).

The resulting chart shows that there is a certain seasonality of the German visits in Prague. The peak for each year is October (Are Germans escaping from Oktoberfest?). By inspecting German arrivals to both Prague and Budapest, it appears that German tourists seem to be influenced more by the seasonality of the business year (more visits during summer) than the crisis, as the patterns seem consistent from the end of 2008 to the end of 2014 and unaffected therefore by the slight GDP drop from 2009. Searches for further destinations (e.g., Copenhagen, Dubrovnik, Venice) can be added to confirm this hypothesis.

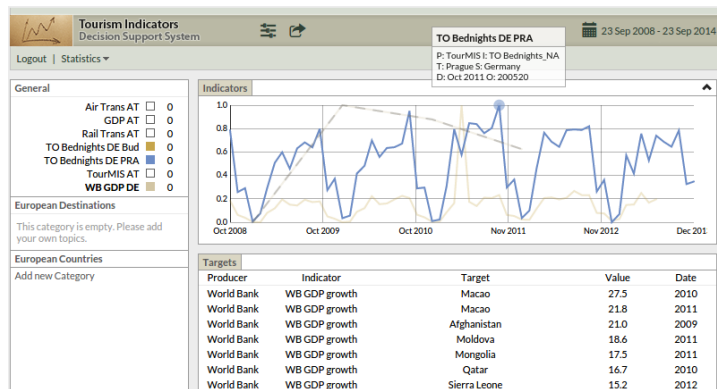


Figure 6: Superimposing economic and tourism indicators.

## 6 Summary and Lessons Learned

Although the tourism domain heavily relies on complex decision-making, it is currently difficult to build decision support systems for this domain that would be capable of seamlessly integrating and visualising data from multiple data sources of tourism (and other) indicators. LD technologies, on the other hand, when adopted at large scale, greatly facilitate data integration at the syntactic and semantic level alike by providing a uniform data encoding format, as well as the possibility to clearly specify the meaning of the data and to establish links between various datasets.

In this paper we described advances to the state of the art in terms of (1) publishing TourMIS data as linked data; (2) interlinking TourMIS data with data from other data sources covering the economic and statistical domains; (3) creating a visual dashboard that explores this integrated data to support cross-domain decision making processes.

We also encountered a suit of issues during our work. This included the necessity to make various changes to the TourMIS dataset before publishing it. Additionally, our solution relies on several tools, mostly research prototypes with poor documentation and support. For example, although several RDB2RDF tools exist almost none of them supports Microsoft SQL Server which acts as a backend for TourMIS. Or, Linked Data publishing tools such as Pubby deliver different performance based on

the underlying triple store – in our particular case, we originally used the Sesame triple store but had to replace it with Fuseki due to its incompatibilities with Pubby. The immaturity of Linked Data tools makes the publishing process time-consuming.

We also encountered licensing problems due to the heterogeneous origin of the TourMIS data set. Although TourMIS is open to anyone who registers to the system, the data comes from different sources and is financed by multiple organizations. Therefore, opening the entire TourMIS data set as Open Data for third parties is an issue which has not yet received a satisfactory answer yet.

Future work includes evaluation of the dashboard with tourism practitioners, inclusion in TourMIS, extension with additional data sources, extension of the dashboard with news media data as well as using the TourMIS linked data for other applications, beyond decision support.

**Acknowledgements.** The work presented in this paper is partly funded by the PlanetData project (FP7:ICT-2009.3.4, #257641). We thank Prof. Karl Wöber and Prof. Arno Scharl for their contribution to the project.

## References

- Berners-Lee, T., Hendler, J., and Lassila, O. (2001) The Semantic Web. *The Scientific American*, May 2001.
- Franke, T. (2006). TourBO: A Prototype of a Regional Tourism Advising System in Germany. In H. W. Daniel R. Fesenmaier, Karl W. Wöber (Ed.), *Destination Recommendation Systems: Behavioural Foundations and Applications* (pp. 257-280). UK: CABI.
- Gruber, T. R. 1993. A translation approach to portable ontology specifications. *Knowl. Acquis.* 5, 2 (June 1993), 199-220.
- Isele, R., Jentzsch, A., Bizer, C. 2010. Silk Server - Adding missing Links while consuming Linked Data. COLD Workshop 2010.
- Mazanec, J. A. (2006). Building Adaptive Systems: A Neural Net Approach. In H. W. Daniel R. Fesenmaier, Karl W. Wöber (Ed.), *Destination Recommendation Systems: Behavioural Foundations and Applications* (pp. 137-170). UK: CABI.
- Ricci, F., Cavada, D., Mirzadeh, N., & Venturini, A. (2006). Case-based Travel Recommendations. In H. W. Daniel R. Fesenmaier, Karl W. Wöber (Ed.), *Destination recommendation Systems: Behavioural Foundations and Applications* (pp. 67-93). UK: CABI.
- Ricci, F., Fesenmaier, D. R., Mirzadeh, N., Rumetshofer, H., Schaumlechner, E., Venturini, A., Zins, A. (2006). DieToRecs: A case-based travel Advisory System. In H. W. Daniel R. Fesenmaier, Karl W. Wöber (Ed.), *Destination Recommendation Systems: Behavioural Foundations and Applications* (pp. 227-239). UK: CABI.
- Sabou, M., Aarsal, I. and Brasoveanu, A.M.P. (2013). TourMISLOD: a Tourism Linked Data Set. *Semantic Web Journal* 4(3): 271-276.
- Ricci, F., & Nguyen, Q. N. (2006). MobyRek: A Conversational Recommender System for On-the-move Travellers. In H. W. Daniel R. Fesenmaier, Karl W. Wöber (Ed.), *Destination Recommendation Systems: Behavioural Foundations and Applications* (pp. 281-294). UK: CABI.
- Wöber, K. 2003. Information supply in tourism management by marketing decision support systems. *Tourism Management*. 24(3):241-255.
- Zins, A., & Grabler, K. (2006). Destination Recommendations Based on Travel Decision Styles. In H. W. Daniel R. Fesenmaier, Karl W. Wöber (Ed.), *Destination Recommendation Systems: Behavioural Foundations and Applications* (pp. 94-120). UK: CABI.