

Optimizing Email Marketing Campaigns in the Airline Industry using Knowledge Graph Embeddings

Amine Dadoun
EURECOM, Sophia Antipolis, France
Amadeus SAS, Biot, France
amine.dadoun@eurecom.fr

Raphaël Troncy
EURECOM, Sophia Antipolis, France
raphael.troncy@eurecom.fr

Michael Defoin Platel
Amadeus SAS, Biot, France
michael.defoinplatel@amadeus.com

Riccardo Petitti
Amadeus SAS, Biot, France
riccardo.petitti@amadeus.com

Gerardo Ayala Solano
Amadeus SAS, Biot, France
gerardo.ayalasolano@amadeus.com

ABSTRACT

The use of email marketing campaigns in e-commerce is an essential aspect for driving sales and establishing or maintaining good customer service. However, ensuring their success in the airline industry is challenging, and the risk is high that the marketing campaign content is not suited to the needs of specific travelers. The traveler may rapidly feel spammed and the offer content promoted via emails may even be no longer accessed. Consequently, the use of personalised recommender systems for email marketing campaigns becomes essential to build user loyalty and to increase offer conversion. Recently, the use of knowledge graphs has proven to be an effective way to recommend products. In this work, we propose an approach that leverages knowledge graph embeddings to better target the right audience in email marketing campaigns for airline products recommendation. We conduct extensive experiments to compare our approach with the currently in-production rule-based system used by airline marketers and a supervised machine learning model based on handcrafted features as another baseline. The results suggest that the use of knowledge graph embeddings is the most effective approach.

CCS CONCEPTS

• **Computing methodologies** → **Learning latent representations; Boosting**; • **Information systems** → **Recommender systems**.

KEYWORDS

Recommender systems, email marketing campaigns, gradient boosting, knowledge graph, knowledge graph embeddings

ACM Reference Format:

Amine Dadoun, Raphaël Troncy, Michael Defoin Platel, Riccardo Petitti, and Gerardo Ayala Solano. 2021. Optimizing Email Marketing Campaigns in the Airline Industry using Knowledge Graph Embeddings. In *Companion Proceedings of the 2021 World Wide Web Conference (WWW '21 Companion)*, April 19-23, 2021, Ljubljana. ACM, New York, NY, USA, 8 pages.

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '21 Companion, April 19-23, 2021, Ljubljana

© 2021 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

1 INTRODUCTION

With the digital transformation of retail commerce from physical sale points to virtual stores, recommender systems have shown their value in easing the search and purchase decision process of customers who are facing an ever increasing amount of products [5]. In order to embrace e-commerce techniques and boost their revenues, some airlines are using the so-called Amadeus Anytime Merchandizing (AAM) Notification System¹ which is an IT solution that allows airline marketers to effectively define, deploy, monitor and adjust airline email marketing campaigns sent to the travelers in real-time. Customized notifications can be defined and sent to travelers, after booking a flight, to suggest them additional services to buy (e.g. extra baggage, specific meal, preferred seat) called ancillaries. The solution acts as a bridge between the travelers retailing touchpoints and the airline's service and delivery system. As shown on the left side of Figure 1, when using this solution, the airline marketer can choose the appropriate time **when** to send the notification (e.g. 5 days before departure), **what** product to recommend (e.g. Leg Space Seat), **how** to send the offer (e.g. via an Email) and to **whom** this offer should be sent (matching targeting criteria).

Despite all the functionalities included in the AAM Notification System, it is difficult for an airline to find the optimal audience to target for a given offer. We conducted an analysis of historical sales triggered by some notification campaigns during the period 14 May 2019 - 17 Dec 2019 ran by one of our partner airlines and we observed a poor conversion of the notification offers (Section 3.3). This is partly due to the challenging decision-making process that an airline marketer faces when it comes to deciding which values (belonging to broad ranges) are appropriate for the criteria to be used (e.g. sending time, flight itineraries, flight departure point, etc.). Targeting customers with unsolicited notifications can be counter-productive and lead to adversarial effects on customer loyalty if done incorrectly. It is therefore critical to identify the customers that we expect to react positively to an advertised service in order to avoid spamming them with non-personalized emails. This problem can be seen as an inverse recommendation scenario, i.e. recommending a user to an item.

Inspired by recent works that have illustrated the effectiveness of using knowledge graph embeddings [13, 15, 18] and gradient boosting algorithms [10, 17] for item recommendation, we propose

¹<https://amadeus.com/en/portfolio/airlines/anytime-merchandising>

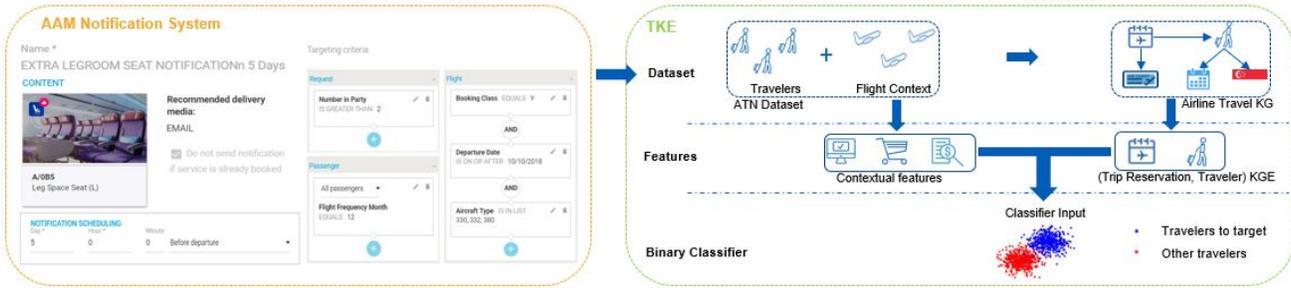


Figure 1: On the left side: AAM Notification System. On the right side: Flowchart of our proposed TKE framework. Notification dataset used in this study is generated from the AAM Notification system. Contextual features include booking context (e.g. number of passengers, date of departure, etc.), notification information (e.g. media used to send the notification, time of notification, etc.).

the Travel Knowledge Graph Embeddings for email marketing campaigns (TKE) framework to better target the audience for a service the airline wishes to recommend through email marketing campaigns (Figure 1). In this paper, we made the following contributions:

- (i) We design and develop a Knowledge Graph using Semantic Web technologies to represent the travelers past trips as well as to semantically enrich airline products.
- (ii) We learn vector representations of travel entities via knowledge graph embedding algorithms and we leverage gradient boosting algorithms to compute prediction scores to better target the audience in email marketing campaigns.
- (iii) We perform an empirical comparison of our approach with the current in-production rule-based system as well as with an hypothetical classical machine learning approach using handcrafted features on a real-world production dataset.

The paper is organized as follows. Section 2 provides a literature review of the related work. Section 3 introduces some preliminary concepts, the recommendation problem, the dataset and the knowledge graph used to conduct this work. In Section 4, we describe both our TKE approach as well as a machine learning model using handcrafted features as another baseline. In Section 5, we present the experiments we carried out and we demonstrate the effectiveness of our approach. Finally, in Section 6, we give some conclusions and we outline some future works.

2 RELATED WORK

This section provides a literature review of the research conducted on email marketing campaigns in the e-commerce and tourism domains. We also present state-of-the-art methods of knowledge graph embedding algorithms and the use of gradient boosting algorithms for recommender systems.

2.1 Email Marketing Campaigns

Emails allow marketers to send messages to their customers at very low cost. They generally generate faster responses and create an opportunity for interactive communication with customers [4]. In [16], the authors analyze 70 randomized field experiments and find that email promotions not only increase customers’ average

purchase spending during the promotion window but also carry over to the week after the promotion expires. In our study, we will focus on personalized email marketing campaigns. In [1], the authors performed an empirical comparison of supervised machine learning models based on decision tree and logistic regression algorithms in order to improve the open rate and conversion rate of email marketing campaigns. In [7], the authors propose to use transactional features and instant messaging metadata to train a boosting tree regression algorithm to timely anticipate the needs of consumers in order to increase their level of engagement as well as the rate at which they repurchase products.

In the tourism domain, even if widely used, limited research is conducted on email marketing campaigns. In [19], the authors found that customers’ favourite emails contain special offers, discounts and coupons as well as real-time communication tools. When customers perceived these emails as meeting their personal preferences, they developed a strong relationship with the sender. In [22], the authors demonstrated that the personalization, interactivity and price were important predictors of the possibility of revisiting the same accommodation. To the best of our knowledge, we are the first to propose a supervised machine learning approach to enable personalized email marketing in the airline travel industry.

2.2 Knowledge Graph Embeddings

A knowledge graph embedding (KG embedding) is a representation of a knowledge graph (KG) component into a continuous vector space. The objective of learning those embeddings is to ease the manipulation of graph components (entities, relations) for prediction tasks such as entity classification, link prediction or recommender systems. In [20], the authors classified the learning algorithms into two main categories namely translational distance models and semantic matching models. For the first category, TransE [2] is often mentioned as the most used translational distance model. TransE represents both entities and relations vectors in the same space R^d . Given a fact (h, r, t) , the relation is interpreted as a translation vector r so that the embedded entities h (head) and t (tail) can be connected by r with low error, i.e., $h + r \approx t$ when (h, r, t) holds. TransH [21] introduces relation-specific hyperplanes, each relation r being represented on a hyperplane as w_r its normal vector.

TransR [11] follows the same idea as TransH, but instead of projecting the relations into a hyperplane, it proposes to create a specific space per relation.

In the second category, semantic matching models exploit similarity-based scoring functions. In [12], the authors associate each entity with a vector to capture its latent semantics. Each relation is represented as a matrix that models pairwise interactions between latent factors. The score of a fact (h, r, t) is defined by a bi-linear scoring function minimized through tensor factorization based on ALS optimization technique.

Palumbo et al. [14] propose to use property-specific KG embeddings generated from node2vec algorithm [9] in order to compute relatedness scores between items and users. Similarly, we propose to use translational distance and semantic matching models to generate embeddings and to use them as latent features of a gradient boosting algorithm.

2.3 Gradient Boosting Algorithms

Gradient boosting is a machine learning technique for regression and classification problems, which produces a prediction model in the form of an ensemble of weaker prediction models.

In [10], the authors used multiple additive Regression Trees (Dart) which is an ensemble model that uses boosted regression trees and handles the overspecialization. Their approach for online accommodation recommendation ranked 1st in the famous RecSys 2019 Challenge². In [17], the authors used XGBoost [3] which is an implementation of gradient boosting decision trees to predict tweet engagement, they have intensively used exploratory data analysis to extract and compute relevant features to feed XGBoost algorithm. Their solution ranked 1st in the RecSys 2020 Challenge³. In our work, we use XGBoost algorithm as a supervised machine learning algorithm to better target the audience in email marketing campaigns based on the computed embeddings and contextual features (Figure 1).

3 DATA & PROBLEM FORMULATION

In this section, we first provide definitions of some useful concepts. Then, we define the problem and research questions addressed in this work. Finally, we present the dataset and the design of the knowledge graph used for the experiments.

3.1 Preliminaries

We define a Knowledge Graph (KG) similarly to what is done in [6].

Definition 1. A knowledge graph is a set $K = (E, R, O)$, where E is the set of entities, $R \subset E \times \Gamma \times E$ is a set of typed relations between entities and O is an ontology. The ontology O defines the set of relation types ('properties') Γ .

Definition 2. Travelers are a subset of the entities of the KG, $t \in T \subset E$.

Definition 3. Ancillaries are all products offered by the airline beyond air tickets. They can be flight-related (e.g. extra baggage,

preferred seat, etc.) or standalone (e.g. lounge access) services. Ancillaries are advertised in the email marketing campaigns. Ancillaries are a subset of the entities of the KG, $a \in A \subset E$.

Definition 4. A notification campaign is a set of notifications sent to an audience of travelers within a given period of time and under some criteria to recommend an ancillary product.

Definition 5. We define the conversion rate of a notification campaign as follows:

$$CR = \frac{1}{N_o} \sum_{i=1}^{N_o} hit(N_i) \quad (1)$$

where N_o is the number of notifications sent through the notification campaign, and $hit(N_i)$ is equal to 1 if the notification N_i triggers a purchase. In our work, we focus on optimizing the conversion rate.

3.2 Problem Formulation

Given a notification campaign aimed at a large audience of travelers who have already booked a flight in a given context, we aim to target the relevant travelers among all the travelers that the notifications will reach. More specifically, we address the following research questions:

RQ 1: How to extract the relevant sample of travelers to target for a given notification campaign?(Figure 2).

RQ 2: How does a supervised machine learning approach perform compared to a rule-based approach to target the relevant audience for a notification campaign?

RQ 3: How does the use of KG embeddings compare to the use of handcrafted features as input of a supervised machine learning model trained to target the relevant audience for a notification campaign?

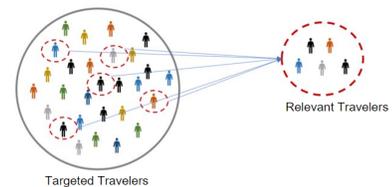


Figure 2: The task is to extract the relevant travelers among the whole set of travelers that were initially targeted by the notification campaign through AAM Notification System.

3.3 Notification Campaign Analysis

We analyzed three notification campaigns summing up to approximately 8.2 million notifications sent by one of our partner airlines to its travelers between 14 May 2019 and 17 December 2019 in order to understand the behavior of travelers in response to the notification campaigns, and to compute the conversion rates of these notification campaigns.

As shown in Table 1, there are three different types of ancillaries that are advertised in three notification campaigns. By analyzing airline sales data over the same period, we can see that only 3 out of 34 different types of purchased ancillaries were offered in the

²<https://recsys.acm.org/recsys19/challenge/>

³<https://recsys.acm.org/recsys20/challenge/>

notification campaigns. This shows an untapped sales potential. Moreover, we observed that 50% of sales triggered by a notification happens on the same day (< 24 hours) after receiving the notification. This demonstrates the effect of a notification on the purchases.

Table 1: Conversion rates of notification campaigns: rule-based approach

Notification Campaign	Notification time	Date Range	Number of Notifications	Sales	CR
Extra leg room seat	5 days before Departure	19 May - 23 December 2019	~355 K	~2.8K	0.8%
Prepaid baggage	2 days before Departure	14 May - 17 December 2019	~7.5 M	~11K	0.15%
Lounge	Right after air ticket purchase	16 October - 17 December 2019	~338 K	104	0.03%
All Notifications	-	-	-	~13.8K	0.18%

While the prepaid baggage notification campaign was aimed at all travelers who booked a flight during the period indicated in Table 1, the notification campaigns for Lounge access and Extra leg room seat contain a number of filtering criteria that explain the large discrepancy in terms of the number of notifications sent out. Indeed, for these two notification campaigns, the airline marketer chose to send the notification to a quite restrictive audience by combining a number of criteria (fare family, aircraft type, no chargeable seat in their booking, etc.)

3.4 Airline Travel Notification Dataset

We conducted experiments on a real-world production dataset of bookings from the T-DNA database⁴. Each booking contains one or several air ticket purchases, and is stored using Passenger Name Record (PNR) information. The PNR is created at reservation time by airline reservation system and contains information about the purchased air ticket (e.g. travel itinerary, payment information), traveler demographics and ancillaries information if purchased (comprised in the EMD ticket). The considered dataset contains approximately 2.33 million bookings for approximately 2.85 million unique travelers.

The Airline Travel Notification (ATN) dataset is produced by joining the notification dataset and the historical bookings dataset from T-DNA. This dataset contains information about the shopping and booking context (e.g. search date, number of passenger, departure date, etc.) and information about travelers (e.g. demographics and loyalty membership information). In total, the dataset contains 42 columns and ~ 8.2 million rows. For our experiments, the dataset was broken down into three different sub-datasets that correspond to the three different notification campaigns (Table 1).

⁴T-DNA: Traveler DNA is a database that contains bookings of travelers over a dozen of airlines

3.5 Airline Travel Knowledge Graph

The Airline Travel KG is constructed from the T-DNA database. We develop an ontology which is available in the Turtle format at https://gitlab.eurecom.fr/amadeus/tke4rec/-/blob/master/ontology/ams_ontology.ttl. To design the KG, we have defined 7 classes corresponding to top level entities and based on the various tables available in the T-DNA database:

- **Traveler:** A traveler is identified by a T-DNA id. A traveler has a booking history (PNRs) that contain purchase history (air tickets, EMD tickets). An instance of traveler is a schema:Person⁵.
- **Trip Reservation:** A trip reservation (PNR) represents the booking of all travelers contained in the PNR. It contains information such as the number of passengers, the destination, etc.
- **Journey:** A journey is linked to a trip reservation. Each journey has a stay duration, departure and arrival airports.
- **Air Ticket:** An air ticket is contained in a PNR and contains flight and transactional information. A PNR can have multiple air tickets because of the different flight legs (e.g. Nice-Paris, Paris-New York) or/and the number of passengers.
- **EMD Ticket:** An Electronic Miscellaneous Document (EMD) ticket is linked to an air ticket. It contains information on the ancillary purchased by the traveler (e.g. ancillary type, ancillary price, etc.).
- **Ancillary:** An ancillary is a service purchased by a traveler (associated to a flight) in addition to the air ticket. It is identified by a sub-code (RFISC), labelled by a commercial name, defined by ATPCO⁶. It belongs to a group of ancillaries (Group, RFIC). We propose to model the different ancillaries as SKOS concepts and we create an ancillary thesaurus represented as a concept scheme.
- **Airport:** It represents the airport where the traveler travels to. An airport serves one or several cities.

The KG used to tackle our use case contains 41 different properties (Figure 3), ~ 80 million edges and ~ 9 million nodes.

We present in Table 2 some statistics about the Airline Travel KG. In Figure 4, an excerpt of the KG is depicted, where a Malaysian

Table 2: Statistics of subgraphs

Subgraph	#Edges	#Nodes	#travelers	#PNRs
Extra leg room seat	7M	800K	67K	205K
Prepaid baggage	64M	7.6M	572K	2.2M
Lounge	6.7M	789K	42K	203K

traveler identified by T21354, born on "1988-05-05" has booked a one way flight for two people from Kuala Lumpur to Melbourne. The EMD ticket identified by 23143 and linked to the air ticket 21563 represents the purchase of an ancillary (a seat).

⁵The prefix schema is used for concepts defined by <https://schema.org>

⁶ATPCO Ancillary description: <https://www.atpco.net/resource/optional-services-industry-sub-codes>

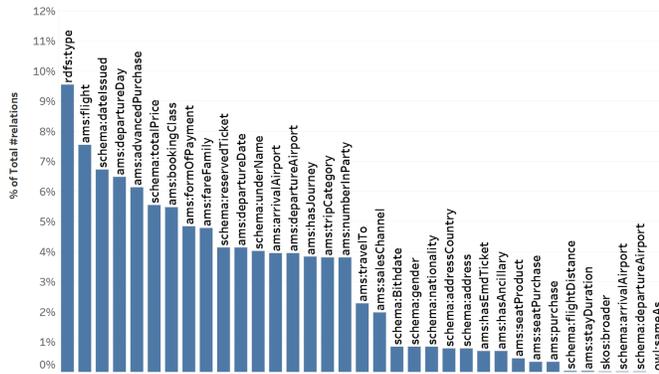


Figure 3: Distribution of #relations of properties in the Airline Travel KG. All prefixes can be found in the ontology definition.

4 APPROACH

Our proposed framework TKE can be seen as a two-stage approach as presented in Figure 1. In the first stage, we extract contextual features from the ATN dataset and compute KG embeddings of travelers and trip reservations from the Airline Travel KG. In the second stage, contextual features and KG embeddings are used as input of an XGBoost classifier in order to predict, for a given user, whether the notification should be sent or not.

We also propose a supervised machine learning model that includes as input contextual features and additional handcrafted travelers’ features that capture travelers’ preferences which we think could be particularly significant for model accuracy (hypothesis proven in the ablation study) as another baseline to compare with. We describe below the handcrafted travelers’ features as well as the KG embeddings used in TKE:

- **Handcrafted travelers’ features:** Features designed to capture travelers’ preferences for ancillaries, destinations, points of sale, etc. and also customer lifetime value (Section 4.1).
- **Knowledge graph embeddings:** Latent features representation of travelers and trip reservations computed based on KG embedding algorithms such as TransE [2] (Section 4.2).

4.1 Handcrafted Travelers Features

In this section, we present the handcrafted travelers features designed to compare their use with that of KG embeddings. We compute several features based on travelers purchase history, such as preferred ancillary, preferred destination, etc. We list below the features computed that leverage travelers’ history:

- **Bookings count:** Number of bookings already purchased by the traveler with the airline.
- **Average flight revenue:** The average booking price tickets for all historical bookings of the traveler.
- **Preferred ancillary:** This feature corresponds to the most sold ancillary to the traveler.
- **Preferred destination:** This feature corresponds to the most visited destination (airport) by the traveler.
- **Preferred seat characteristic:** This feature represents the seat characteristic that is the most purchased by the traveler.

There are three types of seat characteristic namely Upper deck, Exit Row, Leg Space.

- **Extra leg room seat already purchase:** For each seat characteristic, we create a binary feature that represents if a traveler has already purchased an Extra leg room seat or not.
- **Seat sales count:** This feature represents the number of times a seat has been purchased by the traveler.
- **Prepaid baggage sales count:** This feature represents the number of times a prepaid baggage has been purchased by the traveler.
- **Lounge sales count:** This feature represents the number of times a lounge access has been purchased by the traveler.
- **Notification response rate:** This feature is equal to the number of sales divided by the number of notifications sent to the traveler (regardless of the recommended service).

The handcrafted features and the features available in the ATN dataset are used as input of a gradient boosting decision tree classifier. We use the official implementation⁷ of XGBoost [3] to train a binary classifier to predict if the notified travelers will convert or not. In Section 5, we give more details about the hyper-parameters used in XGBoost.

4.2 TKE Features

In this section, we explain how the KG embeddings used in TKE are computed. We use translational distance models to compute travelers and trip reservations embeddings as shown in Figure 1. More formally, we learn the KG embeddings based on a link prediction task, where some links of ancillary purchases and seat products are hidden in the training set, and put in the test set. Translational distance models are trained under the closed world assumption [20] using a pairwise loss that penalizes negative instances. More concretely, ancillaries that were not purchased by a traveler are considered as negative instances under the closed world assumption. Translational distance models are evaluated using ranking metrics such as hit rate or mean reciprocal rank. Hence, these models will return a high similarity score (low euclidean distance) for the ancillaries that are close in the graph embedding space to the embeddings of the ancillaries historically purchased by the travelers. As an example, we obtain a hit rate of ~ 0.42 with the TransE algorithm on the Airline Travel KG. In addition to translational distance models, we implemented a single-hidden multi-layer perceptron (MLP) as proposed in [8], where each relation (as well as entity) is associated with a single vector. More specifically, given a fact (h, r, t), the vector embeddings of h, r, and t are concatenated in the input layer, and mapped to a non-linear hidden layer. The score is then generated by a linear output layer. The generated embeddings are used as input of XGBoost classifier in addition to the contextual features as shown in Figure 1. We carry out a thorough empirical comparison of the aforementioned KG embedding algorithms and select the KG embeddings that allow the classifier to predict with the highest accuracy.

5 EXPERIMENTS

The objective of the experiments is to compare the use of handcrafted features (a) with the use of KG embeddings (b). (a) helps in

⁷XGBoost: <https://XGBoost.readthedocs.io/>

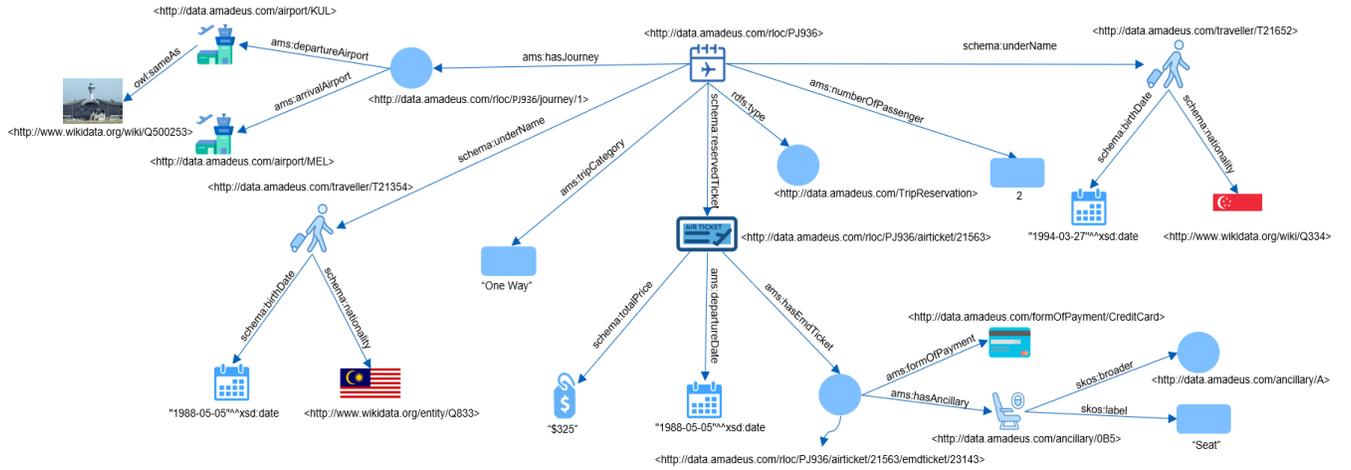


Figure 4: Excerpt of the knowledge graph representing the travelers included in a Trip reservation through the property `schema:underName`, as well as other properties and relations to other entities. Literals are represented in blue rectangle, whereas other entities are represented in blue circle. In this depiction, some properties that links travelers, trip reservations, air tickets and emd tickets are represented as an example, but more properties are included in the graph.

interpreting the results and predictions obtained by the algorithm, while (b) lacks interpretation (latent features), but is easier to compute and maintain. We publish our code as open source in order to ease reproducibility⁸.

5.1 Experimental Setup

In this section, we present the different settings and the evaluation protocol (evaluation metrics and split of the dataset) used to conduct the experiments.

Dataset: We experiment both approaches (a) and (b) with the three datasets presented in Table 1. We use the Airline Travel KG presented in Section 3.5 to generate the KG embeddings useful for our main approach TKE.

Training & Test Sets: The three datasets corresponding to the three notification campaigns are split using the same strategy. Each dataset is sorted temporally, and 80% of the first rows of each dataset are used as training/validation sets. We use a cross-fold validation to train and validate all models ($k=5$, a split of 80% for training and 20% for validation). The remaining 20% are used as test set to evaluate the model. The split between training and validation set is performed randomly in order to avoid a seasonality effect that is usually occurring in the travel industry. KG embedding algorithms are often designed to solve a link prediction task. We consider it is appropriate to split the KG by removing some edges that are included in the set of properties that link travelers with ancillaries and consider them as test sets, in order to evaluate the quality of the computed embeddings.

Evaluation metrics: The output of our approach is the probability of purchasing the recommended ancillary a included in the notification N :

$$P(\text{purchase} = a|N) = P(\text{purchase}|\text{Context}, TE, RE) \quad (2)$$

⁸<https://gitlab.eurecom.fr/amadeus/tke4rec>

where, TE and RE are the Traveler and Trip reservation embeddings.

To evaluate and compare, the different approaches implemented, we used the conversion rate defined at definition 5 and the three metrics defined as follows:

- **TPR:** The true positive rate is the percentage of correct positive predictions. It represents the ratio of travelers that the algorithm suggests to send the notification and effectively purchase the ancillary. TPR is defined as follows:

$$TPR = \frac{TP}{(TP + FN)} \quad (3)$$

- **TNR:** The true negative rate is the percentage of correct negative predictions. It represents the ratio of travelers that the algorithm suggest to not send the notification and effectively do not purchase the ancillary. TNR is defined as follows:

$$TNR = \frac{TN}{(TN + FP)} \quad (4)$$

- **ROC-AUC:** The area under ROC curve (FPR, TPR) helps to choose what is the optimal probability threshold that maximizes the CR and TPR and is defined as follows:

$$ROC-AUC = \int_0^1 TPR d(FPR) \quad (5)$$

where, $FPR = 1 - TPR$ is the false positive rate

It is noteworthy that the conversion rate was measured offline as well as all the metrics based on the test set. According to equation 1, N_o represents the number of predicted positives and each hit hit_i corresponds to a true positive prediction.

Implementation Framework & Parameter Settings: For KG embedding algorithms, we use the deep learning framework pytorch⁹ to implement MLP [8] and the library pykg2vec [23] for all the other KG embedding algorithms. The hyper-parameters of

⁹Pytorch:<https://pytorch.org/>

Table 3: Evaluation results of the different approaches. (a) represents the results of XGBoost for different inputs; (b) represents the results of the TKE approach for different KG embedding algorithms. The average standard deviation (by varying the seed when splitting the dataset) of each metric is as follows: $AUC - ROC : \pm 0.02$, $TPR : \pm 3\%$, $TNR : \pm 2\%$, $CR : \pm 0.1\%$

Model	Extra leg room seat				Prepaid baggage				Lounge			
	AUC-ROC	TPR	TNR	CR	AUC-ROC	TPR	TNR	CR	AUC-ROC	TPR	TNR	CR
Rule-based	-	-	-	0.8%	-	-	-	0.15%	-	-	-	0.03%
(a) ADS	0.75	78%	58%	2.2%	0.83	80%	71%	0.38%	0.76	80%	62%	0.18%
(a) HDS	0.79	81%	60%	2.37%	0.85	82%	74%	0.4%	0.84	86%	67%	0.22%
(a) AHDS	0.83	85%	65%	2.8%	0.88	86%	74%	0.56%	0.89	88%	65%	0.36%
(b) TransE	0.85	86%	69%	3.1%	0.91	92%	65%	0.6%	0.90	89%	78%	0.35%
(b) TransH	0.84	85%	67%	3%	0.90	91%	65%	0.59%	0.95	96%	85%	0.59%
(b) TransR	0.84	85%	67%	2.9%	0.90	91%	65%	0.6%	0.92	92%	80%	0.52%
(b) MLP	0.87	88%	69%	3.2%	0.92	94%	65%	0.62%	0.91	90%	81%	0.56%

all the models were tuned using a combination of random-search and grid-search algorithms. We apply grid-search algorithm on the implemented algorithms using the following values: the embedding size $k \in \{32, 64, 96, 128, 256\}$, the batch size $\in \{128, 256, 512, 1024\}$, the number of epochs $\in \{50, 100, 200\}$, the learning rate $lr \in \{0.001, 0.003, 0.01, 0.03, 0.1, 0.3\}$ and negative samples $Ns \in [2, 10]$ for MLP algorithm. We also optimize the following hyper-parameters of XGBoost classifier: the max depth of a tree $\in [5, 50]$, the number of trees $\in [10, 100]$, the sub-sample of each tree $\in [0.65, 0.85]$ and the col-sample of each tree $\in [0.65, 0.85]$. In addition to these hyper-parameters, we compute a weighted score (ratio of number of negative class to the positive class) that we use in XGBoost to approach the problem as a cost-sensitive learning problem due to the high class imbalance between positive (purchase) and negative (no purchase) classes (Table 1).

5.2 Results and Discussion

In this section, we discuss the results obtained from the experiments. Results of the experiments conducted are presented in table 3. TPR, TNR and ROC-AUC metrics are not provided for the rule-based approach implemented in AAM Notification System. The reason behind this is that the dataset used in the experiments is generated by the AAM notification system, which is different from the original dataset that contains all travelers used for the rule-based approach to identify the travelers matching the targeting criteria.

Ablation Study: Table 3 shows that using the features from the ATN dataset in addition to the travelers handcrafted features (AHDS: ATN + Handcrafted features) as input of XGBoost performs better than using only one of them (ADS: ATN features or HDS: Handcrafted features) as input for all notification campaigns. We also observe that using only travelers handcrafted features as input information of XGBoost gives better results than using the entire ATN dataset for all notification campaigns. We compute the most important features of the model (a) AHDS for each notification campaign and we report below the three most important ones with their respective information gain:

- Extra Leg Room Seat: {Preferred Seat Characteristic: 0.31, Preferred ancillary: 0.12, Ticket amount: 0.08}.

- Prepaid Baggage: {Preferred destination: 0.21, Destination: 0.12, Prepaid Baggage sales Frequency: 0.10}.
- Lounge: {Average Flight Revenue: 0.22, Destination: 0.20, Age: 0.15}.

Knowledge graph embeddings: We observe in Table 3 that using KG embeddings (concatenation of traveler and reservation KG embeddings) with contextual features as input of XGBoost performs better than using travelers handcrafted features regardless of the algorithm used to compute the embeddings or the notification campaign. Moreover, KG embeddings computed from **MLP** shows to perform better than KG embeddings computed from translational distance models except for the lounge notification campaign, where the use of KG embeddings computed from **TransH** model gives the best results. In Table 4, we present the values of the hyper-parameters that lead to the best results given in Table 3.

Table 4: Best performing hyper-parameters and model for our knowledge graph embedding approach

Notification Campaign	Model	k	lr	Ns	batch size	epochs
Extra leg room seat	MLP	128	0.003	4	512	100
Prepaid baggage	MLP	128	0.001	4	512	50
Lounge	TransH	96	0.03	-	128	50

6 CONCLUSIONS AND FUTURE WORK

In this work, we have presented a two stage approach to address the problem of audience targeting for email marketing campaigns: first, we compute KG embeddings of travelers and reservations; second, we use these embeddings in addition to contextual features as input of a XGBoost classifier to learn what is the relevant audience to target for a given notification campaign. We conducted several experiments to address our research questions:

RQ 1: The results of the experiments presented in Table 3 show that extracting the relevant audience for a given notification campaign is not an easy task. Indeed, despite the fact that the conversion rate increases significantly with our approach, it remains relatively small. However, thanks to our approach, notification campaigns are better targeted and we manage to avoid recommending an unsuitable service to at least 65% of passengers.

RQ 2: Experiments have shown that the handcrafted features based supervised machine learning approach (a) gives better results than the rule-based one. Indeed, in Table 3, we can observe that the conversion rate is multiplied by more than 3 for Extra Leg Room Seat, almost 4 for Prepaid Baggage, and 12 for Lounge. Hence, we prove the benefit of using supervised machine learning over a simpler rule-based approach while it is the currently adopted mechanism used by airline marketers. It should be noted that the list of possible criteria available in AAM Notification System (Figure 1) is the same as the list of features used in the supervised machine learning approach.

RQ 3: Experiments show that regardless of the KG embedding algorithm tested, the KG embedding approach is better than the handcrafted features approach. This is very interesting from a scientific point of view, as it shows the added value of having a KG in the travel domain that could be used not only for ancillary recommendation task, but also other recommendation tasks (e.g. Trip recommendation) as the same KG embeddings could be used. It is worth noticing that when dealing with a cold-start problem (new user or item) for on-line usability, a rule-based approach is more appropriate.

Finally, as future work, we expect to tackle the task of personalized ancillary ranking in email marketing campaigns. More specifically, the goal of our future work will be to answer what is the most appropriate service to recommend in a notification campaign. In addition to addressing and optimizing what to recommend to a traveler, it would be interesting to optimize the time when to send the notification as it is an important decision making factor, especially in the airline travel industry [5].

REFERENCES

- [1] Redouan Abakouy, El Mokhtar En-Naimi, and Anass El Haddadi. 2017. Classification and Prediction Based Data Mining Algorithms to Predict Email Marketing Campaigns. In *Proceedings of the 2nd International Conference on Computing and Wireless Communication Systems (ICWC'S'17)*. Association for Computing Machinery, New York, NY, USA, Article 34, 5 pages. <https://doi.org/10.1145/3167486.3167520>
- [2] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating Embeddings for Modeling Multi-relational Data. In *Advances in Neural Information Processing Systems 26*, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (Eds.). Curran Associates, Inc., Red Hook, NY, USA, 2787–2795. <http://papers.nips.cc/paper/5071-translating-embeddings-for-modeling-multi-relational-data.pdf>
- [3] Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A Scalable Tree Boosting System. In *22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*. Association for Computing Machinery, San Francisco, California, USA, 785–794.
- [4] L. Chittenden and R. Rettie. 2003. An evaluation of e-mail marketing and factors affecting response. *Journal of Targeting, Measurement and Analysis for Marketing* 11 (2003), 203–217.
- [5] Amine Dadoun, Thomas Fiig, Michael Defoin-Platel, Corine Landra, and Raphaël Troncy. 2021. How Recommender Systems can Transform Airline Offer Construction and Retailing. *Journal of Revenue and Pricing Management, Special issue on AI* (2021).
- [6] Amine Dadoun, Raphaël Troncy, Olivier Ratier, and Riccardo Petitti. 2019. Location Embeddings for Next Trip Recommendation. In *9th International Workshop on Location on the Web (LocWeb)*. Association for Computing Machinery, San Francisco, USA, 896–903.
- [7] Alexandros Deligiannis, Charalampos Argyriou, and Dimitrios Kourtesis. 2020. Predicting the Optimal Date and Time to Send Personalized Marketing Messages to Repeat Buyers. *International Journal of Advanced Computer Science and Applications* 11, 4 (2020). <https://doi.org/10.14569/IJACSA.2020.110413>
- [8] Xin Dong, Evgeniy Gabrilovich, Geremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmann, Shaohua Sun, and Wei Zhang. 2014. Knowledge Vault: A Web-Scale Approach to Probabilistic Knowledge Fusion. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '14)*. Association for Computing Machinery, New York, NY, USA, 601–610. <https://doi.org/10.1145/2623330.2623623>
- [9] Aditya Grover and Jure Leskovec. 2016. Node2Vec: Scalable Feature Learning for Networks. In *22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, New York, NY, USA, 855–864.
- [10] Paweł Jankiewicz, Liudmyla Kyryshchuk, Paweł Sienkowski, and Magdalena Wójcik. 2019. Boosting Algorithms for a Session-Based, Context-Aware Recommender System in an Online Travel Domain. In *Proceedings of the Workshop on ACM Recommender Systems Challenge (RecSys Challenge '19)*. Association for Computing Machinery, New York, NY, USA, Article 1, 5 pages.
- [11] Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning Entity and Relation Embeddings for Knowledge Graph Completion. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI'15)*. AAAI Press, Austin, Texas, 2181–2187.
- [12] Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A Three-Way Model for Collective Learning on Multi-Relational Data. In *Proceedings of the 28th International Conference on International Conference on Machine Learning (ICML'11)*. Omnipress, Madison, WI, USA, 809–816.
- [13] Enrico Palumbo, Diego Monti, Giuseppe Rizzo, Raphaël Troncy, and Elena Baralis. 2020. entity2rec: Property-specific knowledge graph embeddings for item recommendation. *Expert Syst. Appl.* 151 (2020), 113235. <https://doi.org/10.1016/j.eswa.2020.113235>
- [14] Enrico Palumbo, Giuseppe Rizzo, and Raphaël Troncy. 2017. Entity2Rec: Learning User-Item Relatedness from Knowledge Graphs for Top-N Item Recommendation. In *Eleventh ACM Conference on Recommender Systems*. ACM, New York, NY, USA, 32–36.
- [15] Enrico Palumbo, Giuseppe Rizzo, Raphaël Troncy, Elena Baralis, Michele Osella, and Enrico Ferro. 2018. Translational Models for Item Recommendation. In *The Semantic Web: ESWC 2018 Satellite Events*, Aldo Gangemi, Anna Lisa Gentile, Andrea Giovanni Nuzzolese, Sebastian Rudolph, Maria Maleshkova, Heiko Paulheim, Jeff Z Pan, and Mehwish Alam (Eds.). Springer International Publishing, Cham, 478–490.
- [16] Navdeep S. Sahni, Dan Zou, and Pradeep K. Chintagunta. 2017. Do Targeted Discount Offers Serve as Advertising? Evidence from 70 Field Experiments. *Manage. Sci.* 63, 8 (Aug. 2017), 2688–2705. <https://doi.org/10.1287/mnsc.2016.2450>
- [17] Benedikt Schifferer, Gilberto Titericz, Chris Deotte, Christof Henkel, Kazuki Onodera, Jiwei Liu, Bojan Tunguz, Even Oldridge, Gabriel De Souza Pereira Moreira, and Ahmet Erdem. 2020. GPU Accelerated Feature Engineering and Training for Recommender Systems. In *Proceedings of the Recommender Systems Challenge 2020 (RecSysChallenge '20)*. Association for Computing Machinery, New York, NY, USA, 16–23. <https://doi.org/10.1145/3415959.3415996>
- [18] Zhu Sun, Jie Yang, Jie Zhang, Alessandro Bozzon, Long-Kai Huang, and Chi Xu. 2018. Recurrent Knowledge Graph Embedding for Effective Recommendation. In *12th ACM Conference on Recommender Systems*. ACM, New York, NY, USA, 297–305.
- [19] Kar Yan Tam and Shuk Ying Ho. 2005. Web Personalization as a Persuasion Strategy: An Elaboration Likelihood Model Perspective. *Info. Sys. Research* 16, 3 (Sept. 2005), 271–291. <https://doi.org/10.1287/isre.1050.0058>
- [20] Q. Wang, Z. Mao, B. Wang, and L. Guo. 2017. Knowledge Graph Embedding: A Survey of Approaches and Applications. *IEEE Transactions on Knowledge and Data Engineering* 29, 12 (Dec 2017), 2724–2743. <https://doi.org/10.1109/TKDE.2017.2754499>
- [21] Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge Graph Embedding by Translating on Hyperplanes. In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence (AAAI'14)*. AAAI Press, Québec City, Québec, Canada, 1112–1119.
- [22] Kiseol Yang, Ji Hye Min, and Krishna Garza-Baker. 2019. Post-stay email marketing implications for the hotel industry: Role of email features, attitude, revisit intention and leisure involvement level. *Journal of Vacation Marketing* 25, 4 (2019), 405–417.
- [23] Shih-Yuan Yu, Sujit Rokka Chhetri, Arquimedes Canedo, Palash Goyal, and Mohammad Abdullah Al Faruque. 2021. Pykg2vec: A Python Library for Knowledge Graph Embedding. *Journal of Machine Learning Research* 22, 16 (2021), 1–6. <http://jmlr.org/papers/v22/19-433.html>