



How recommender systems can transform airline offer construction and retailing

Amine Dadoun^{1,2} · Michael Defoin-Platel² · Thomas Fiig³ · Corinne Landra² · Raphaël Troncy¹

Received: 15 July 2020 / Accepted: 18 September 2020 / Published online: 20 March 2021
© The Author(s), under exclusive licence to Springer Nature Limited 2021

Abstract

Recommender systems have already been introduced in several industries such as retailing and entertainment, with great success. However, their application in the airline industry remains in its infancy. We discuss why this has been the case and why this situation is about to change in light of IATA's New Distribution Capability standard. We argue that recommender systems, as a component of the Offer Management System, hold the key to providing customer centricity with their ability to understand and respond to the needs of the customers through all touchpoints during the traveler journey. We present six recommender system use cases that cover the entire traveler journey and we discuss the particular mind-set and needs of the customer for each of these use cases. Recent advancements in Artificial Intelligence have enabled the development of a new generation of recommender systems to provide more accurate, contextualized and personalized offers to customers. This paper contains a systematic review of the different families of recommender system algorithms and discusses how the use cases can be implemented in practice by matching them with a recommender system algorithm.

Keywords Recommender systems · Artificial intelligence · Dynamic offer construction · NDC

Introduction

A recommender system can be seen as an algorithm to compute the probability that a user (customer) would like to interact with an item (product or service). These systems were originally introduced to overcome the problem of information overload that customers face when exposed to a large catalog of products or services. By providing the customers with contextualized and personalized recommendations,

recommender systems aim at narrowing down the search to a manageable subset of products that are relevant to the customer.

Recommender systems have proven to be popular for both customers and sellers, particularly for online retail (Resnick and Varian 1997). The most representative example is Amazon that has become one of the largest retailers in the world because, among other important things such as a large selection of products and a fast and reliable delivery chain, it offers best-of-breed customer experience as a result of an extensive use of recommender systems. Recommender systems result in a more personalized shopping experience, giving customers the feeling of being understood and recognized which contributes in building trust and in maintaining loyalty.

From the seller's point of view, recommender systems offer the possibility to control and to increase the exposure of their catalog by driving customers toward products lacking visibility. Recommender systems are also notoriously good at decreasing bounce rate and at increasing average time spent on a web page for online selling (Taghipour and Kardan 2008). Finally, recommender systems have also proved to be very effective offline in email marketing

✉ Amine Dadoun
amine.dadoun@eurecom.fr

Michael Defoin-Platel
michael.defoinplatel@amadeus.com

Thomas Fiig
thomas.fiig@amadeus.com

Corinne Landra
clandra@amadeus.com

Raphaël Troncy
raphael.troncy@eurecom.fr

¹ Eurecom, Sophia Antipolis, Biot, France

² Amadeus, Sophia Antipolis, Biot, France

³ Amadeus, Kastrup, Denmark



campaigns allowing sellers to run so-called “one-to-one marketing” at scale (Jannach and Jugovac 2019).

Recommender systems are growing in popularity in the travel industry to address the complex set of decisions customers face when booking a flight, selecting a hotel or finding relevant events and activities at their destination. For example, Airbnb is now offering real-time personalization of search rankings within its marketplace (Grbovic and Cheng 2018). Travel agencies or brokers have recently called upon the research community to work further on the particularities of making recommendations in the context of travel. The online travel agency Trivago sponsored the 2019 Recommender Systems Challenge as part of the Association for Computing Machinery (ACM) RecSys conference in order to improve their current recommender system for hotels.

However, despite the successful application of recommender systems across many industries, airline offer construction and retailing remains quite rudimentary with little or no differentiation in how products and services are selected, retailed, or priced across customers. There are several reasons for this. First, in the current airline distribution model, airlines have delegated control of the offer construction to content aggregators, such as global distribution systems (GDSs). Real-time interactions with the airline systems are quite limited, and the pricing function which is used to create offers on behalf of the airline is governed by industry standards that only enable very few parameters to differentiate the content based on who the traveler is. Therefore, airlines cannot provide personalized and contextualized offers in a meaningful way. Second, the responsibility of the offer construction and retailing has historically been managed across separate departments within the airline organization. Offer construction and retailing were therefore never part of a broader and holistic customer experience management strategy.

We believe the current approach is inadequate and that the key to profitability is to manage offers consistently in an integrated Offer Management System (OMS) serving the customer throughout the traveler journey from inspiration to post-trip. However, realizing this vision will require significant advancements in both the science of offer construction and in the distribution capabilities employed across all distribution channels, being direct as well as intermediated.

On the distribution side, this advancement will happen as part of IATA’s New Distribution Capability (NDC), which will allow airlines to move toward customer centric airline retailing. NDC is an enabler for the application of airline OMS including recommender systems. Industry adoption of NDC has continued to grow in recent years. As of August 2020, 40 airlines, 20 aggregators and 10 sellers are NDC certified level 4 (the highest level) covering booking of NDC content as well as supporting changes of the order IATA (2020).

On the science side, the airline industry literature is still underdeveloped in terms of how dynamic offer construction could be designed and implemented. The key contributions of this paper are therefore to detail illustrative examples of recommender system use cases in the airline industry context and to discuss how these use cases could be implemented in practice with the benefits for both airlines and travelers.

The remainder of this paper is structured as follows. In the next section, we present the traveler journey and we identify use cases for recommender systems. Next, we describe the traditional airline distribution model, the new distribution model enabled by NDC, and the airline’s Offer Management System, which will dramatically influence airline offer construction and retailing. We then review the scientific concepts behind each family of recommender systems. Subsequently, we match the use cases with the most appropriate families of algorithms. Finally, we provide some conclusions and we outline some future research directions.

Recommender system use cases throughout the traveler journey

The traveler journey is a key consideration to understand the customer needs and intents (Fig. 1). In their report, Frost & Sullivan (2014) indicate that there “are certain moments when the customer is in a purchasing mind-set and thinking about his trip and what he will need”. For example, at the booking stage, the customer is in a “planning” mind-set. At this stage, the airline can approach the customer with more “expensive” offers such as cabin upgrade, or flexibility options. Close to departure (48 h/24 h), the customer has a different mind-set—making the final preparations for his trip. At this moment, airlines could propose the customer with extra baggage, airport transfer, parking, priority check-in, or fast track access. In this section, we detail some use cases for recommender systems along different phases of the traveler journey.

To provide more in-depth discussion, we focus on recommender systems that are under airline control. These use cases cover customers that actively search and book travel products through the standard distribution channels enabled by NDC—both direct and indirect channels. Thus, use cases for recommender systems regarding customer acquisition through the Internet giants’ web interfaces, social media, and search engines will not be covered, since in these cases, the recommender systems reside outside the airline’s control.



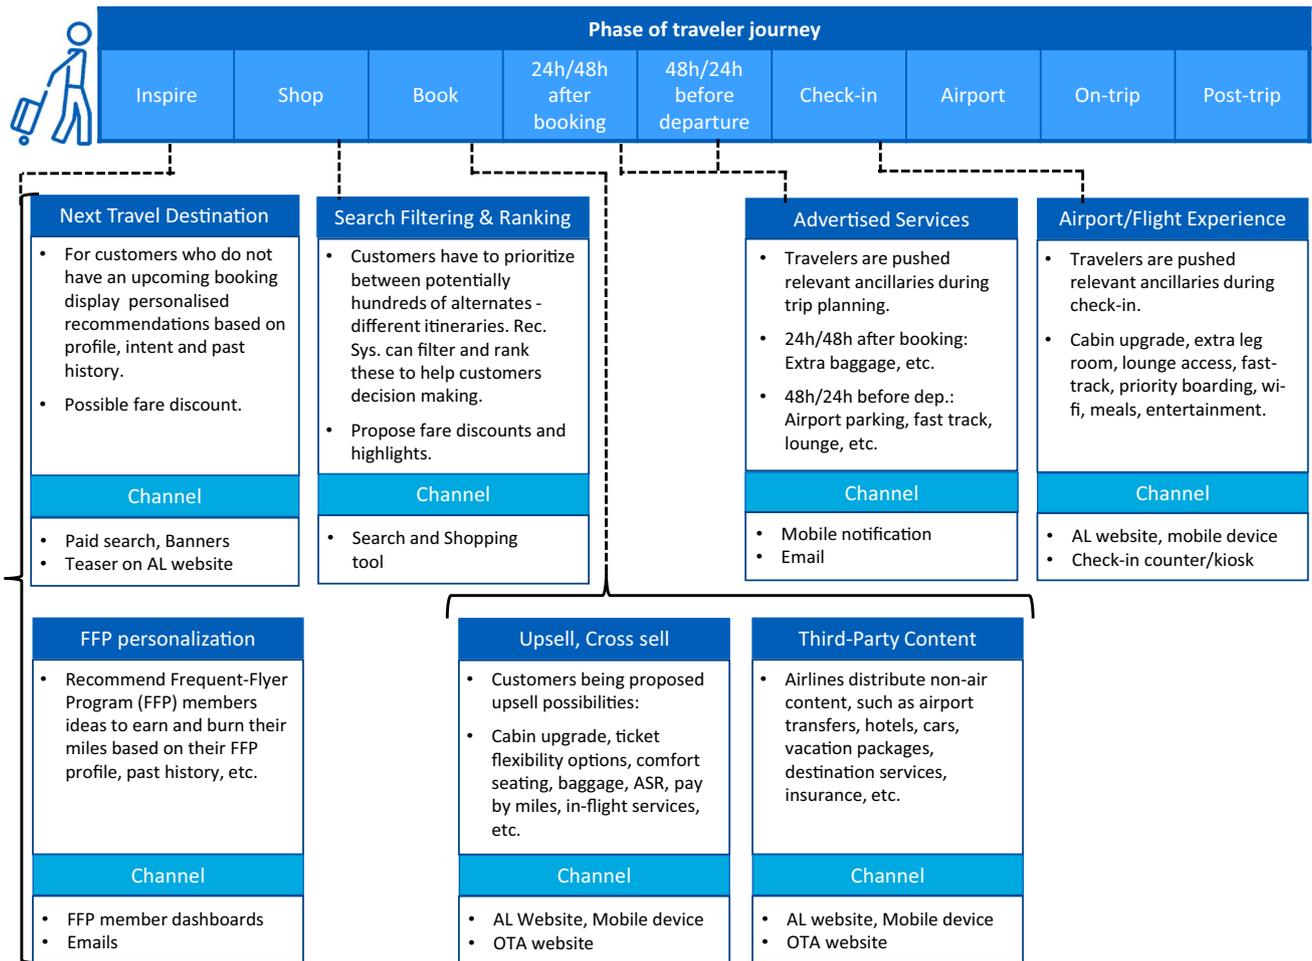


Fig. 1 Recommender system use cases throughout the traveler journey

Next travel destination

The inspiration phase is a key opportunity to influence the customer decision-making process. We distinguish between *passive inspiration* and *interactive inspiration*. The former represents the case when a customer (typically anonymously) lands on a web page and receives travel inspiration simply because some routes are popular in general, while the latter corresponds to the case where the customer interacts with the recommender system by providing personalized search criteria. In the following, to be concrete, we take the assumption that the customer stays anonymous and is engaged in interactive inspiration, providing the recommender system with more leverage.

Affinity shopping tools can be employed to create a personalized shopping experience. Rather than selecting the traditional criteria of origin/destination and calendar dates, these tools enable inspiration based on personalized criteria, such as customers’ budget and interests (events or destination type such as beach, city, etc.). A recommender

system with access to information of upcoming events (e.g., jazz festivals, sport events, exhibitions, etc.), and real-time information about flight prices and promotional fares (campaigns) could be used to recommend the most appropriate destinations and dates that match the customers criteria. Further, it could also recommend how the offers should be retailed using rich format such as infographics, photos and videos. For example, a trip during the summer to Nice Côte d’Azur in France, should have a very different presentation depending on if the customer is interested in beach, nightlife or a culinary experience.

FFP personalization

The frequent-flyer program (FFP) business model is dependent on FFP members having sufficient incentive to earn and burn their points. However, in reality, this may not be so easy. Premium-tier members with large point balances may not be able to find availability on attractive flights or



premium classes due to blackouts or lack of award availability, while low-tier members with small point balances often cannot afford a redemption ticket and see no value in the program.

Recommender systems are in a good position to increase the number of points burned using information about both the member's point balance and the availability of award tickets. For example, the premium-tier member may be offered to burn points for upgrades for his/her family on their annual vacation trip (to mitigate the dilution risk of the award ticket substituting a commercial ticket) or non-air content not readily accessible for purchase on the open market (e.g., backstage passes to concerts, games, etc.). For the low-tier member, recommender systems could offer a "discount" toward the fare of a commercial ticket.

Several other use cases for recommender systems can also be identified, such as incentivizing members to earn points to reach the next tier level or burn points that are close to expiration. In all these cases, the system may be able to increase the value of the program by sending personalized emails to members with the right offer at the right time.

Search filtering and ranking

For a customer who makes searches by comparison shopping, booking air travel can be a daunting experience. He or she must prioritize among potentially hundreds of itineraries, with different prices and product characteristics across multiple partner airlines. As a result, it becomes almost impossible for the customer to make a purchase decision. Today, most search algorithms aim at finding the lowest fares but, in doing so, create irrelevant or unattractive itineraries that distract or overwhelm the customer.

A recommender system can filter the choice set into a manageable number of alternatives and rank them in order of relevancy based on an understanding of the customer's stated criteria. In this way, the recommender system both guides the customer in his decision process and benefits the airline through improved conversion rates. We may also add new customized criteria beyond the usual origin destination, date range, flying time, ground time and overnight stay criteria to incorporate product attributes such as cabin, ticket flexibility, seat reservation and baggage allowance that are not typically considered in comparison shopping requests today.

Upsell, cross-sell and third-party content

When the customer has decided on his preferred itinerary, he enters the booking stage. During the booking stage, the recommender system has ideal information about the customer and his travel party—not only the current trip destination, duration, and already-selected ancillary services, but also the customer's profile and historic purchases. At the booking stage, the customer is in a planning mind-set and this is an ideal opportunity to both increase ancillary revenues for the airlines as well as offer a one-stop shopping experience that covers the customer's full journey.

Examples of products that could be recommended at this stage include upsell offers such as cabin upgrades or ticket flexibility options, as well as cross-sell offers such as baggage, advance seat reservations or in-flights services (e.g., meals). In addition, the airline can also offer third-party content. Based on the customer needs, the commercial relation with the third parties, the prices and availabilities for the relevant resources, the recommender system can propose simple products such as insurance, airport transfers, etc., or even more complex bundled travel such as vacation packages that include hotels and rental cars.

Advertised services

During the post-shopping period, the airline has an opportunity to push offers to customers through unsolicited mail or via notification on a mobile device. This period is a critical phase for the customers' last-minute decisions and preparations for their trip. Customers can be approached with ancillary services such as extra luggage, airport parking, seat selection, priority check-in, etc., and also be informed of availability of cabin upgrades that are aligned with their preferences. Again, the offer and communication would be very different between a family of four traveling long-haul from Frankfurt to New York City in economy class for a two weeks' vacation, versus a business purpose customer traveling the same itinerary and cabin, but staying only for two days. A recommender system would propose not only the most relevant offers but also the most relevant channel and time to push these offers with the benefit of increased adoption rates and customer satisfaction.

Airport/flight experience

During check-in, the customers actively interact with the airline via employees at the check-in counter, the kiosk, or on mobile devices. During this phase, the customer is focusing on the practicalities before takeoff. This may regard logistics



of how to navigate through the airport, but the customer may also wish to indulge themselves with restaurants, lounge access, or cabin upgrades, which could be paid for example using FFP points.

Considering the personas mentioned before, the family of four returning from their vacation in New York City may have excess baggage, while the business purpose customer returning from New York City on a red-eye flight may be looking for an upgrade to the business cabin. These examples serve to illustrate that customers' needs may vary significantly and that the airline has an opportunity to approach the customers with relevant offers based on a deep understanding of their needs, preferences and intent.

Toward a new distribution capability for the airline industry

In this section, we first detail the traditional airline distribution model. This will provide the necessary background for understanding the objectives behind the new distribution standards, known as the new distribution capability (NDC), which we discuss subsequently. We demonstrate that NDC is an enabler for the application of the airline OMS including recommender systems.

Traditional distribution model

Figure 2 shows how a customer's request for an itinerary is passed from a retailing platform (airline retailing platform, or other retailing platforms), possibly through a distributor, and to the airline's Inventory system for evaluation, using the distribution model in place today. For the direct channel (direct connect), the airline fully controls the shopping and pricing flow. However, for the indirect channels, the current distribution paradigm relies on a two-step process. First, the

airline files fares with data distributors such as ATPCO or SITA. These filed fares drive the construction and pricing of the products that can be offered to the customers. Then, the availability computation within the airline's Inventory system (flight execution) determines which of the filed fares are made available for sale. The airlines control the availability computation via their Revenue Management Systems (RMS), which essentially can be performed using offline optimization (airline planning).

Other retailing platforms may interact directly with the airline's flight execution layer via proprietary interfaces. Distributors such as the GDSs acquire the filed fares content and have the authorization to build offers on behalf of the airlines (delegated shopping & pricing). The distributors then poll the airline's availability to determine which fare products are available for sale. Consistency across indirect channels is enabled by highly standardized content and associated processing logic that the GDSs adopt and implement when accepting airline content and developing their shopping and pricing engines. This means that there is a limited ability for customer-specific information to be used in the indirect distribution channel. In principle, even if the airlines could create contextualized and personalized offers in the direct channel, this would create inconsistency that cannot be resolved among the distribution channels.

New distribution capability (NDC)

The new distribution capability (NDC) is a set of new technical communication standards that was initiated almost a decade ago by the International Air Transport Association (IATA). The vision with NDC is to modernize airline distribution and enable airlines to have better control of their offers and their retailing. We list below the most important benefits for airlines that are adopting NDC, which are of

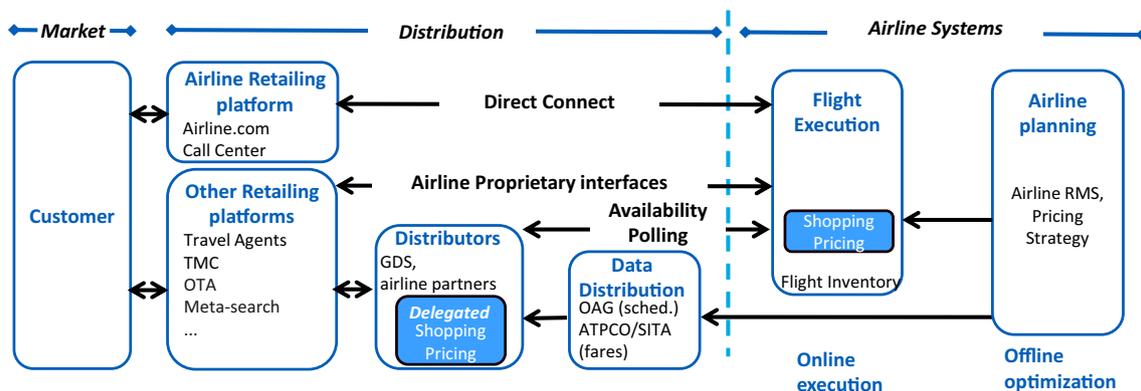


Fig. 2 Traditional distribution model



particular relevance for this paper. For further information on the objectives and benefits of NDC, we refer the reader to (Hoyles 2015).

- **Personalized and contextualized offers.** The airlines will have access to customer and contextual information in a shopping or booking request, which will allow for personalized and contextualized offers.
- **Dynamic offers.** The airlines will be able to create, distribute, and fulfill dynamic offers as described in the next section.
- **Dynamic pricing.** The airlines can employ dynamic pricing using a continuous price.
- **Retailing.** The airlines can provide the retailing platforms with product description that encompasses retailing preferences and information. For instance, rich media content that further complements their offers using visual elements, such as infographics, photos, videos, etc.
- **Merchandising.** The airlines will be able to employ merchandising techniques to affect customers purchase behavior.

In the NDC environment, airlines still make the decision of distributing via direct channels and/or via indirect channels with third-party intermediation. However, delegation of the offer creation to intermediaries no longer exists. Instead, each customer shopping request in an agent’s front-office system is passed to the airline OMS, either directly in the case of NDC direct connect distribution, or via an aggregator in the case of NDC Intermediated distribution. Note that the airline proprietary interfaces and availability polling arrows in Fig. 2 have been replaced by NDC direct connect and NDC intermediated arrows in Fig. 3, enabling a cost efficient deployment at scale for the distribution network actors. The airline’s OMS creates a set of one or more offers that are returned to the customer. Each offer is individually tagged with an offer ID that can be used in any subsequent request on that offer. If the customer accepts an offer, the offer is converted into an order and the contract with the customer is established.

Figure 3 shows how airlines are aspiring to take control of the offer creation, at scale and across all distribution channels.

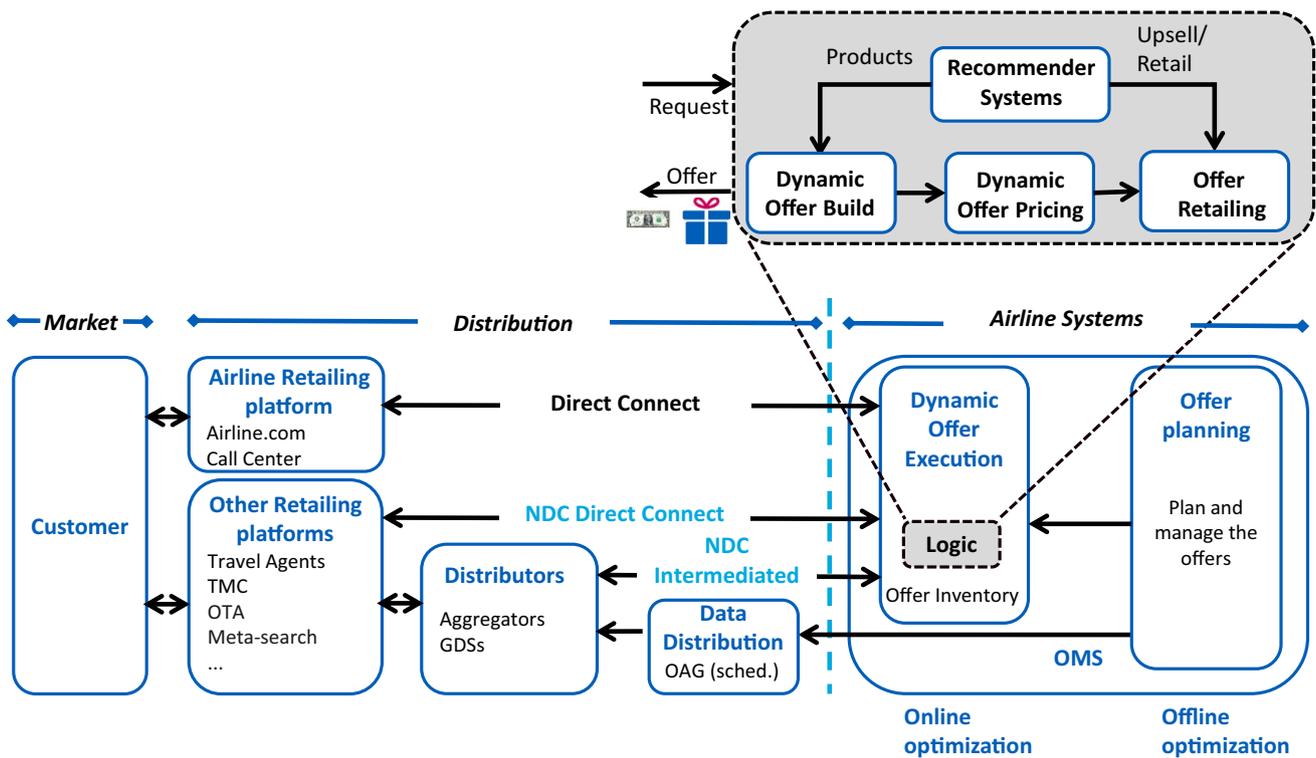


Fig. 3 Distribution model using NDC



The offer management system (OMS)

As seen in Fig. 3, the airline OMS controls the offer construction and retailing for both the direct channel and the indirect channel in NDC. We can think about OMS as an extension of the airline's RMS in several dimensions.

The main extensions are as follows. First, RMS optimizes only the prices (actually the availabilities) of the pre-filed flight products, while OMS optimizes both product components (flight products, ancillaries, third-party content) and prices. Second, unlike RMS which provides the same price to all customers for a given flight and fare product, OMS may differentiate among customers and construct personalized and contextualized offers. Third, and not considered by RMS, OMS may construct one or multiple offers in a so-called *offer set* that will be displayed together as options. For further information, we direct readers to (Fiig et al. 2018).

Finally, because RMS does not differentiate among customers, the price computation can essentially be pre-computed during the offline optimization processes and the online process is a lightweight execution logic. For OMS, this is not the case, as computing personalized and contextual offers is designed to be a real-time decision and the optimization logic must be moved to the online domain. This has significant ramifications for the IT system design of the OMS, which we will discuss below.

The online optimization logic of the OMS is comprised of the following components, which is illustrated in the inset in Fig. 3. In particular, we would like to draw attention to the role of recommender systems in guiding both the Dynamic Offer Build and the Offer Retailing, which has also been exemplified with the recommender system use cases presented.

- **Dynamic Offer Build.** This module makes the determination of the relevant set of products (flights, ancillaries, and third-party content) to be returned at the individualized customer level.
- **Dynamic Offer Pricing.** This module takes as input the offers that were built by “Dynamic Offer Build” and determines for each of these the selling price that maximizes the contribution considering both customer and contextual information.
- **Offer Retailing.** This module aims to increase conversion rates by applying merchandizing techniques to affect the customer's purchasing behavior.

In the description above, we have seen the different functional steps of an OMS to dynamically construct, price and retail an offer. However, we also need to consider the ecosystem that will trigger and support this process. In

particular, online search engines have strict performance requirements. As these engines generate thousands of search transactions per booking, these IT systems need to be extremely cost-effective, scalable and resilient, to provide real-time dynamic offer construction and retailing while providing consistency across all distribution channels. Recent advancements in technology and infrastructure capabilities can enable airlines and system providers to accomplish these goals. For example, cloud infrastructure and real-time worldwide data synchronization and processing power allow data centers across continents to host and run local instances of the online optimization logic, accessible to any distribution channel, while continuously being under airline control.

The science of recommender systems

Introduction to recommender systems

In the terminology of recommender systems, the customers are referred to as *users* and the products in the catalog are referred to as *items*. Hence, a recommender system can be seen as a way to compute the probability that a user would like to interact with an item and use this probability to recommend the most relevant subset of items to him. Depending on the context, an interaction would correspond to the act of searching, buying, visiting, watching, etc.

In its most simple form, a recommender system is typically built in three consecutive steps: *information collection*, *learning* and *recommendation* (Isinkaye et al. 2015). The information collection phase consists in building a weighted graph $G = (U, I, E, w)$, where U , the set of users, and I , the set of items, are the nodes in the graph and E corresponds to the set of edges. These edges represent the past interactions between users and items. There are no edges between the users nor the items, hence the graph is bipartite. The strength of these past interactions is given by the function $w: E \mapsto [0, 1]$.

In the learning phase, a machine learning (ML) algorithm is used to train a model \mathcal{W} that approximates w in G . Finally, in the recommendation phase, the trained model is used to predict, for every possible pair $(u, i) \in (U \times I)$, the strength of the interaction between user u and item i . From these predictions, it is then possible to derive the list of items that could be recommended to the users.

From tapestry (Goldberg et al. 1992), introduced in the early 90s that is considered as the first example of a working collaborative filtering algorithm, to the massive usage of deep learning algorithms (Zhang et al. 2019), the research on recommender systems is now one of the most prolific topics in the artificial intelligence (AI) literature.



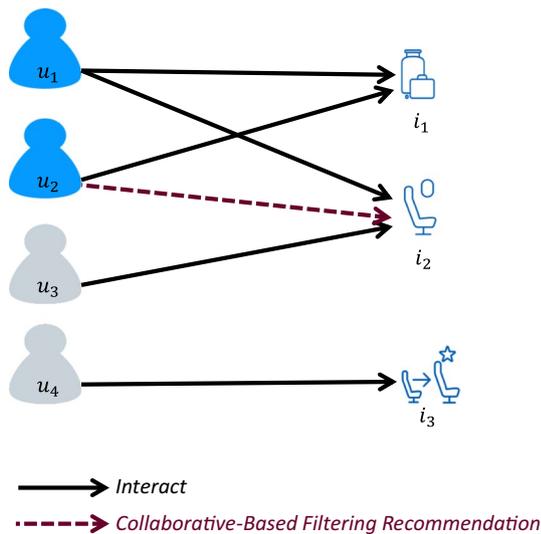


Fig. 4 CF recommender systems: bipartite graph between users and items showing how item i_2 is recommended to user u_2 through a CF algorithm

Machine learning models to predict user-item interactions have evolved from using simple models such as linear and logistic regression to deep neural network models that endow them non-linearity, and thus allow them to find non-linear patterns in the data. However, each of these approaches has its own specificities and it is important to understand their strengths and limitations when addressing a particular use case. In this section, we review the main families of recommender systems.

Collaborative filtering recommender systems (CF)

Collaborative filtering (CF) algorithms are among the most widely used algorithms in the field of recommender systems (Sarwar et al. 2001) and have been applied in industries such as e-commerce or online entertainment to recommend the most relevant products or movies to their customers. In the original formulation, a CF algorithm relies only on the interactions present in the graph G without any additional knowledge or information about the items or the users.

Figure 4 shows an illustrative example of the bipartite user-item graph G for ancillary products. The graph contains interactions between users (travelers) and items (seat, baggage, etc.) represented by the solid arrows, while the dashed arrows represent the recommendations obtained from CF algorithms. Let us consider the item i_1 (baggage) for example. Users u_1 and u_2 both purchased this item. Furthermore, user u_1 also purchased item i_2 , thus item i_2 is recommended to user u_2 .

We can divide CF algorithms into two different classes of methods, the first one relying on *Matrix Factorization* techniques (Hu et al. 2008) and the second one, named *Neighborhood Methods* (Sarwar et al. 2001), relying on computing the similarity between users or items.

Over the years, significant progress has been made to improve CF algorithms, for example in terms of learning speed (He et al. 2016) or accuracy (Rendle et al. 2009; He et al. 2017). Nevertheless, despite their proven overall effectiveness and usability, CF algorithms are still limited especially when users interact with a restricted number of items (*data sparsity*) or when new users or items frequently enter the system and, consequently, past interactions are not available (the user or item *cold start problem*).

Content-based filtering recommender system (CB)

The content-based (CB) filtering method (Lieberman 1995) aims at building user preference profiles based not only on historical user-to-item interactions but also on a form of description of these items that is often represented by a set of keywords or properties. Conversely, it is also possible to associate items to user profiles by looking at the description of the users interacting with them.

In Fig. 5, we present the graph G enriched with the item properties needed for the use of CB recommender systems. Each item (ancillary product) is characterized by a set of properties: for example, the baggage item has the value "C" for the Reason for Issuance Code (RFIC) and the value "A" for the Reason for Issuance Code (RFIC) and the value "A"

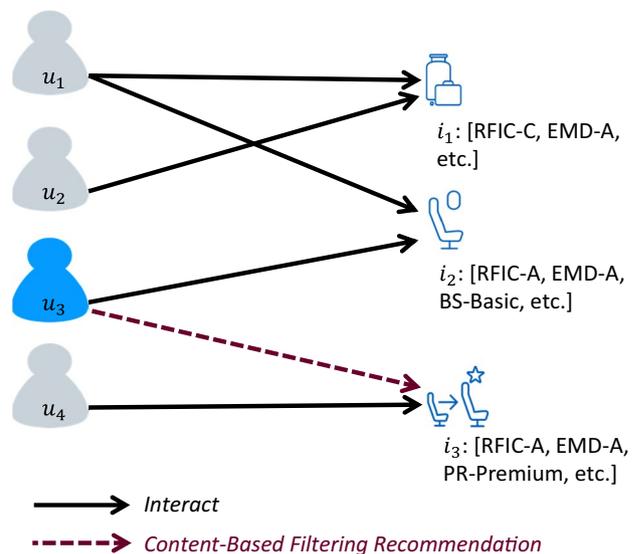


Fig. 5 CB recommender systems: bipartite graph between users and items enriched with item descriptions showing how item i_3 is recommended to user u_3 through CB algorithm



for the Electronic Miscellaneous Document (EMD) category, as it is a flight-associated product. In this example, the CB algorithm recommends item i_3 (premium seat) to user u_3 because item i_3 has the same characteristics of item i_2 which user u_3 has interacted with (added in his cart) in the past.

With CB filtering, even new items without any previously observed interactions will have at least a description that can be used by the system to provide recommendations. Hence, the problem of item *cold start* is mitigated. Nevertheless, CB filtering methods also have some shortcomings. For example, building and maintaining relevant representations for every item can turn into a heavy feature engineering task. Also, introducing novelty into what is being recommended to a given user is not possible since the system works only by looking at content associated with the user's past interactions.

One of the alternatives to deal with the above mentioned limitations such as the lack of novelty consists in mixing CB and CF techniques in what is referred to as Hybrid recommender systems in the literature (Melville et al. 2002; Khrouf and Troncy 2013).

Context-aware recommender system (CA)

CF or CB algorithms model the users' behavior by relying on past user-item interactions or on the content of the items. However, to better capture the complex decision-making process that the users are following when exposed to a selection of items (e.g., the offer set construction by OMS), it is

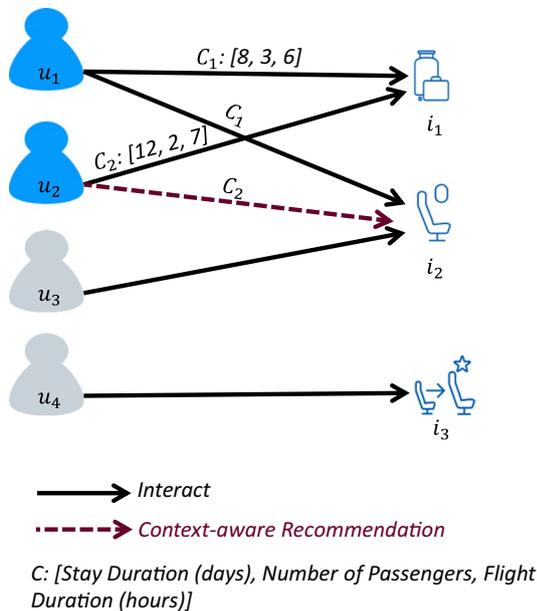


Fig. 6 CA recommender systems: bipartite graph between users and items enriched with contextual information showing how item i_2 is recommended to user u_2 through CA algorithm

crucial to consider the overall context of this process. For instance, a user who wants to travel during summer with four people for two weeks (likely leisure travel) will not have the same needs when traveling alone for two days during a winter week (likely business travel).

A context-aware (CA) recommender system should first be able to collect contextual information and then make use of it to better tailor the offers depending on the circumstances. In Fig. 6, we present the graph G enriched with contextual information. As an illustration, let us consider that the user u_1 who purchased both items i_1 (baggage) and i_2 (seat) for his trip to Paris which will last 8 days with a flight duration of 6 hours. On the other hand, we consider the user u_2 that will travel from New York to Paris on a similarly long flight (7 hours) for 12 days and purchased item i_1 in addition to the flight ticket. Item i_2 is being recommended to user u_2 , as contexts C_1 & C_2 are closely related.

Several initiatives have been conducted to enrich existing recommendation approaches with contextual information. We can categorize them into three different groups (Adomavicius and Tuzhilin 2015): (i) Contextual Pre-filtering (Adomavicius and Tuzhilin 2005) where the contextual information is used only to filter out the graph of user-item interactions to keep only the data pertaining to a particular context; (ii) Contextual Post-filtering (Panniello et al. 2009) where the context is used to produce contextualized recommendations on top of what a traditional recommender system suggests; and finally (iii) Contextual Modeling (Karatzoglou et al. 2010; Rendle 2010; Xiao et al. 2017) where the context itself is considered by the model as input information together with the user-item interaction graph.

Knowledge-aware recommender system (KA)

According to Paulheim (2017), a Knowledge Graph (KG) (i) mainly describes real-world entities and their interrelations, organized in a graph, (ii) defines possible classes and relations of entities in a schema, (iii) allows for potentially interrelating arbitrary entities with each other and (iv) covers various topical domains.

KGs became an increasingly popular research direction toward cognition and human-level intelligence, and are now used in many AI applications such as semantic search or automatic fraud detection. In recent years, KGs have also been introduced in Knowledge-Aware (KA) recommender systems (Palumbo et al. 2017) in order to enrich the graph of user-item interactions with more complex and structured information about the users, the items, and the interactions themselves.

In Fig. 7, an example of a KA recommender system is shown. Beyond the simple lists of properties already



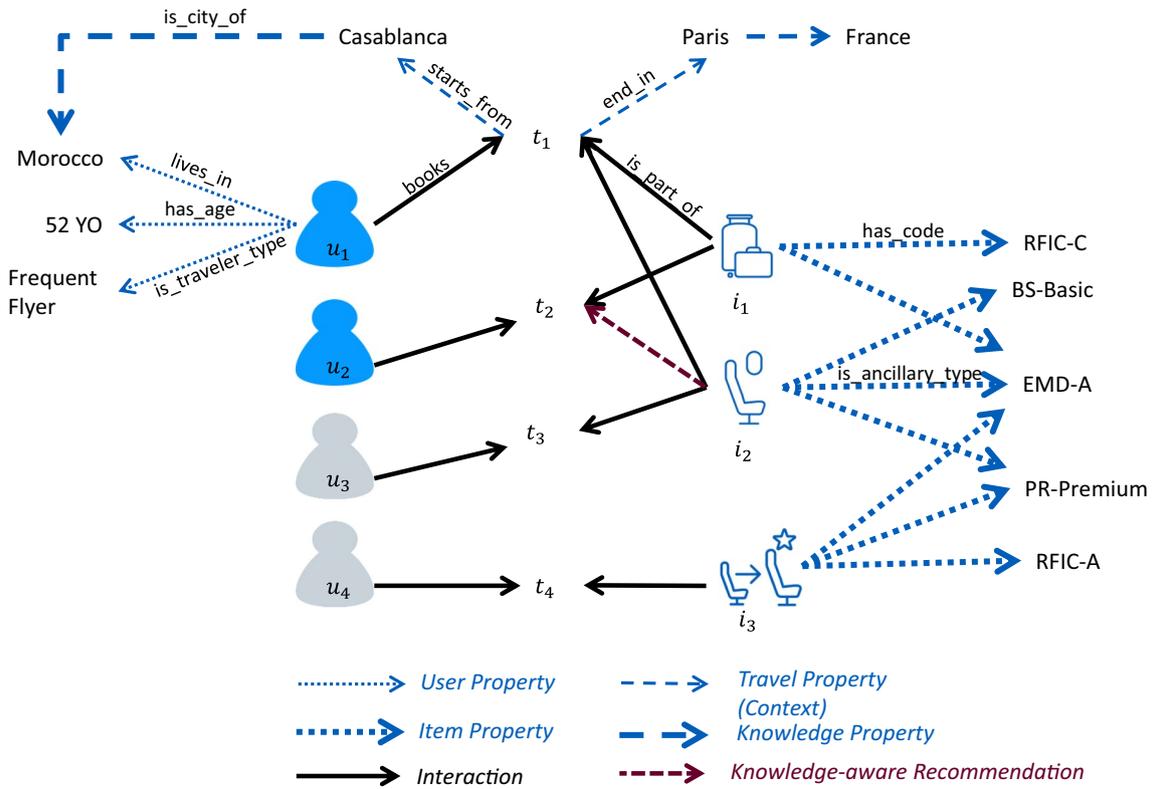


Fig. 7 KA recommender systems: knowledge graph representing user-item interactions in addition to information about users, items and the context of each interaction showing how item i_2 is recommended to the user u_2 via KA algorithm over the knowledge graph

managed by previous versions of recommender systems, KGs represent and leverage semantically rich relations between entities. We see that travel t_1 booked by user u_1 starts from Casablanca, a city in Morocco, which is also the country where user u_1 lives. By construction, KGs can easily be linked between each other. For example, it would be straightforward to extend the graph from Fig. 7 to include cities’ main Points of Interest (Monti et al. 2018).

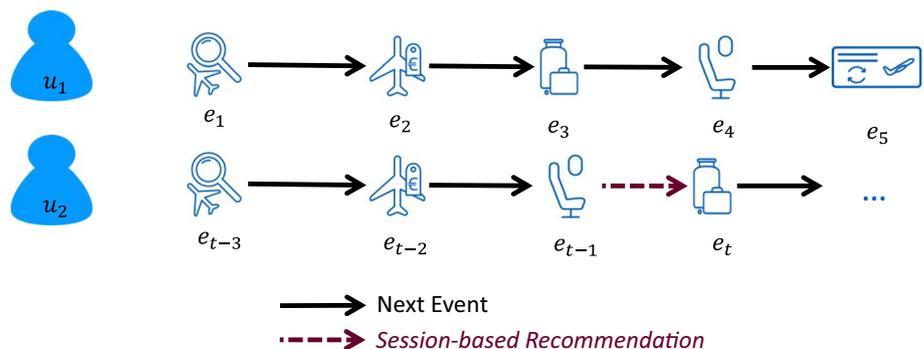
One remarkable thing about KA recommender systems is their ability to make use of the KG structure to provide better recommendations (Sun et al. 2018). Using deep learning, and in particular, graph embedding techniques (Zhang et al. 2016; Palumbo et al. 2017), it is now possible to turn

virtually any type of information into a vector that the system can learn. Dadoun et al. (2019) proposed to use the Semantic Trail Dataset (Monti et al. 2018) that contains users’ check-ins in many cities around the world to build location embeddings for travel recommendations.

Session-based recommender system (SB)

Recommender system approaches based on historical user-item interactions are very powerful because they are able to exploit long-term user profiles (Ludewig and Jannach 2018). However, in real-world applications such as e-commerce platforms, many new users visit the system every day for

Fig. 8 Session-based recommender systems: sequence of user events (interaction with the catalog), user u_2 is being recommended a bag at t through SB algorithm



which no historical information is available (the user *cold start* problem).

It is therefore necessary to analyze users’ live sequence of actions (for instance, their sequence of clicks) to identify patterns and generate recommendations (Linden et al. 2003). This approach can range from simply detecting frequently co-occurring actions (Agrawal et al. 1993) to a more in-depth modeling of the sequence itself with deep learning techniques (Hidasi et al. 2016).

In Fig. 8, user u_1 starts a browsing session looking for a flight (event e_1), then chooses his flight (e_2) and adds it to the shopping cart, and he decides to add two ancillaries (seat and baggage) which represent events e_3 and e_4 , to finally make his booking e_5 . On the other hand, user u_2 follows the same path as u_1 for his first two events and decides at $t - 1$ to add a seat to his shopping cart. Since adding seat and baggage in the same shopping cart are two co-occurring events, the session-based (SB) recommender system will propose to user u_2 to add baggage to his cart.

Beyond the different families of algorithms described in this section, the field of recommender systems is in constant evolution with more and more complex approaches being regularly proposed to address the limitations of the previous generations. As an example, a promising research direction mixing reinforcement learning (Sutton and Barto 2018) and

recommender systems (Rohde et al. 2018; Zhao et al. 2018) is being explored with the ambition to focus on long-term returns and break the pernicious feedback loop of recommendation as described in (Chaney et al. 2018).

Adapting recommender systems for offer construction and retailing

In this section, we revisit the use cases introduced in the Section “Recommender system use cases throughout the traveler journey” and we discuss how they can be implemented in practice using the families of recommender system algorithms described in the previous section. We identify the most appropriate algorithms given the non-functional requirements, such as (i) the available input data, (ii) the output data, (iii) the chosen objectives, and (iv) the operational constraints (e.g., response times). For each use case, we also provide relevant metrics that could be used to assess the quality of each recommender system. Figure 9 provides a summary of this analysis.

		Next Travel Destination	FFP Personalization	Search Filtering & Ranking	Upsell, Cross sell & Third Party content	Advertised Services	Airport/Flight Experience
Input Data	Past user-item interactions	-	✓	-	✓	✓	✓
	User information	-	✓	-	✓	✓	✓
	Item Information	✓	✓	✓	✓	✓	✓
	Context Information	✓	✓	✓	✓	✓	✓
	Knowledge Graph	-	✓	-	✓	✓	✓
	Live interactions	✓	-	✓	✓	-	-
	Extra information	RMS, Interests, Budget, Upcoming events	FFP, RMS	RMS	Third Party	-	FFP, RMS
Output Data	Offer build	destination, date range	action to burn points	ranking of offers	ranking of offers	ancillary proposition	ranking of offers
	Offer retail	presentation, infographics, description	offers timing	offers highlight	offers highlight	offer timing	presentation, infographics, description
Objectives	Travelers’ Loyalty	-	✓	-	✓	✓	✓
	Air product conversion	✓	✓	✓	-	-	-
	Ancillary product Conversion	-	-	-	✓	✓	✓
	Third Party Conversion	-	-	-	✓	-	-
	Miles burned	-	✓	-	-	-	-
Specifics	Response Time	-	-	✓	✓	-	✓
	Data Acquisition	-	-	✓	✓	-	✓
Algorithms Family		CB, CA	CA, KA	CB, CA, SB	SB, CA, KA	KA	CA, KA

Fig. 9 Summary of recommender system algorithms for each use case given the input data, outputs, objectives and constraints. Algorithms in brackets are feasible, while the algorithms without bracket are preferred



Next travel destination

We take the assumption that the customer (user) is anonymous at this stage of the traveler journey. Hence, for this use case, we cannot rely on the past interactions of the user and we discard the use of sophisticated algorithms such as KA that are most effective with this information. Instead, we consider using CA algorithms in a post-filtering fashion starting with CB or SB algorithms to rank destinations based on either the content of the destinations (CB) or the user's clicks through his live interactions (SB). The outputs of the CB/SB algorithms can then be filtered according to the criteria specified by the user from the search tool. Metrics used to evaluate the recommendations could be click-through rate and conversion rate.

FFP personalization

In this use case, the customer identity is known and we can therefore leverage on individual FFP data—such as tier level, point balance, point expiration dates, recency, frequency, and monetary value—but also on price/point conversion rates for the recommended itineraries and services in order to produce meaningful recommendations. The algorithm must also be able to mix this information with a variety of other data from different sources, ranging from the product catalog of air and non-air products, the customer travel history, and the product availability and prices provided by the RMS.

Hence, because of their data integration capabilities, KA algorithms appear to be the natural choice for this complex use case. Moreover, as demonstrated in Yao et al. (2015), KA can be extended to include contextual information allowing the algorithm to capture the travel intent of the user. Metrics used to evaluate the recommendations could be conversion rate and FFP points burned.

Search filtering & ranking

We take the assumption that the customer (user) is anonymous during this stage. In this situation, the recommender system will have to rely on stated criteria (origin-destination, date range, stops, etc.), the context of the search (search time and date, type of the device being used, etc.), product attributes (cabin, flexibility, baggage allowance, etc.), and possible extended criteria depending on the capabilities of the search tool. The recommender system may also employ user navigation behavior to better understand the travel intent. Given the input data available, CA/SB recommender systems (Rendle 2010; Sarwar et al. 2001) seem to be judicious choices provided that session

data can be acquired and response time kept within acceptable limits. Metrics used to evaluate the recommendations could be click-through rate, conversion rate, and sales.

Upsell, cross-sell and third-party content

At this stage, the customer identity is known. However, the customer travel history will, in many cases, still be absent or rather limited. In this case, SB/CA algorithms could be considered. On the other hand, when customer travel history is present, hybrid approaches integrating personalized recommendations could be investigated using for example the KA algorithms. Response time and data acquisition are important specifics of this use case and must be taken into consideration before the preferred algorithm is chosen. Of note, the SB algorithms have a very fast execution time compared to CA and KA, which may impact the choice. Metrics used to evaluate the recommendations could be Conversion Rate, ancillary/third-party revenue and adoption rates.

Advertised services

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. This problem can be seen as an inverse recommendation scenario—recommending a user to an item.

This problem is well-suited for the KA algorithm. Indeed, in this use case where the customer identity is known, the algorithm can take advantage of a diverse set of data: collaborative information (e.g., historical ancillary purchases), user-related information (e.g., number in party), item-related information (e.g., product descriptions), and context-related information (e.g., attributes of the current order). Additionally, other ML approaches such as contextual multi-armed bandits (Li et al. 2010) could also be employed to find the best timing and channel for sending the notifications. Metrics used to evaluate the recommendations could be click-through rate, conversion rate, and incremental revenue.

Airport/flight experience

The time period spent at the airport or during the flight itself is a particularly favorable window of opportunity for the airlines to approach the traveler with personalized and contextualized offers. The algorithms of choice could be CF or CB given their ability to learn the preferences of the travelers and provide near real-time recommendations, especially when the product catalog is rather limited. Alternatively,



the CA algorithm should be also considered, since this algorithm is able to capture travel intent which may well be of importance in this use case. The conversion rate, incremental revenue, FFP points burned are the most appropriate metrics to evaluate how these algorithms perform.

Conclusion and future research directions

Recommender systems have already been introduced in several industries such as retailing and entertainment, where their capability to display personalized and contextualized recommendations have provided benefits to customers and sellers alike. However, their application in the airline industry remains in its infancy. In this paper, we explain that this is primarily a result of the limitations of IT systems that delegate airline control of offer creation to content aggregators. The traditional distribution paradigm relies on a two-step process—fare filing which drives the product and price construction, followed by the availability computation—which provides airlines with limited control over offer construction and retailing. Further, the airlines are unaware of the customer's identity and therefore unable to generate personalized recommendations.

NDC is an enabler for the airlines to provide contextualized and personalized offers, thereby opening the door for the application of recommender systems via the airlines offer management systems (OMS). We believe that recommender systems hold the key to customer centricity with their ability to understand and respond to the needs of the customers throughout all touchpoints during the traveler journey, which we have exemplified with airline-specific recommender system use cases.

We have explained how recent advances in AI have enabled the development of a new generation of recommender systems to provide more accurate, contextualized and personalized offers to users. However, choosing one family of algorithms over another can be a complex task for a travel industry expert because of the large number of algorithms described in the literature and the particularities of the travel domain. Therefore, we have for each of the use cases, provided guidance by identifying the preferred algorithms.

While we have discussed how the application of recommender systems can provide "short-term" (or transactional) benefit to the airline through increased ancillary adoption rates and revenue, we believe that recommender systems may have an even greater opportunity for improving customer experience and increasing customer loyalty by enabling airlines to understand their customers' needs, preferences and intent. The impacts of effective recommendations and retailing on customer loyalty in the airline industry have yet to be explored.

We propose three main areas for our future research directions.

- **Empirical Study.** The next logical step is to perform an empirical study of the performance of the algorithms using actual airline data. This requires to partner with airlines in order to acquire real life data.
- **Explainability.** One of the main challenges for the AI community is to bring explainability to decision-making algorithms. Indeed, it is crucial to understand why an algorithm has recommended a specific item. One popular method of explainability arises from Neighborhood Methods that can state, for example, that "a customer that bought this item, also bought these items". The KA recommender systems are also ideally suited for this purpose, as this algorithm constructs an explainable path within the knowledge graph that lead to the item recommendation (Song et al. 2019). Moreover, performing an ablation study on algorithm inputs, where an input of a model is removed to assess the effect on algorithm performance, would allow us to understand what input data are the most beneficial for an accurate prediction.
- **Industry disruptions.** The Covid-19 pandemic has disrupted the airline industry in an unprecedented way. The industry may not experience a smooth recovery but rather in waves as different countries open/close for air traffic in response to pandemic evolution. This raises questions of the performance and robustness of the different algorithms in the presence of sparse, scattered, and constantly evolving data.

References

- Adomavicius, G., and A. Tuzhilin. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transaction on Knowledge and Data Engineering* 17 (6): 734–749.
- Adomavicius, G., and A. Tuzhilin. 2015. *Context-Aware Recommender Systems*, 191–226. Berlin: Springer.
- Agrawal, R., T. Imieliński, and A. Swami. 1993. Mining association rules between sets of items in large databases. In *ACM SIGMOD International Conference on Management of Data (ICDM)*, pp. 207–216. Washington, D.C., USA.
- Chaney, A. J. B., B. M. Stewart, and B. E. Engelhardt. (2018). How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *12th ACM Conference on Recommender Systems (RecSys)*, pp. 224—232, Vancouver, British Columbia, Canada.
- Dadoun, A., R. Troncy, O. Ratier, and R. Petitti. (2019). Location embeddings for next trip recommendation. In *9th International Workshop on Location on the Web (LocWeb)*, pp. 896–903, San Francisco, USA.
- Fiig, T., R. Guen, and M. Gauchet. 2018. Dynamic pricing of airline offers. *Journal of Revenue and Pricing Management* 17: 04.



- Frost and Sullivan. 2014. *Thinking like a retailer*. <https://amadeus.com/documents/en/blog/pdf/2014/12/report-thinking-like-a-retailer-airline-merchandising.pdf>.
- Goldberg, D., D. Nichols, B.M. Oki, and D. Terry. 1992. Using collaborative filtering to weave an information tapestry. *Communication of the ACM* 35 (12): 61–70.
- Grbovic, M., and H. Cheng. 2018. Real-time personalization using embeddings for search ranking at Airbnb. In 24th ACM SIGKDD international conference on knowledge discovery and data mining (KDD).
- He, X., H. Zhang, M.-Y. Kan, and T.-S. Chua. 2016. Fast matrix factorization for online recommendation with implicit feedback. In 39th International ACM conference on research and development in information retrieval (SIGIR), pp. 549–558.
- He, X., L. Liao, H. Zhang, L. Nie, X. Hu, and T. Chua. 2017. Neural collaborative filtering. In 26th International conference on world wide web (WWW), pp. 173–182.
- Hidasi, B., A. Karatzoglou, L. Baltrunas, and D. Tikk. (2016). Session-based recommendations with recurrent neural networks. arXiv 1511.06939. <https://arxiv.org/abs/1511.06939>.
- Hoyle, Y. 2015. New distribution capability (NDC)—together, let's build airline retailing. <https://www.iata.org/contentassets/6de4dce5f38b45ce82b0db42acd23d1c/guide-ndc-registration-certification-program.pdf>.
- Hu, Y., Y. Koren, and C. Volinsky. 2008. Collaborative filtering for implicit feedback datasets. In 8th IEEE international conference on data mining (ICDM), pp. 263–272.
- IATA. 2020. Together, let's build airline retailing ndc program update january. <https://www.iata.org/contentassets/6de4dce5f38b45ce82b0db42acd23d1c/ndc-standard-presentation.pdf>.
- Isinkaye, F., Y. Folajimi, and B. Ojokoh. 2015. Recommendation systems: Principles, methods and evaluation. *Egyptian Informatics Journal*. <https://doi.org/10.1016/j.eij.2015.06.005>.
- Jannach, D., and M. Jugovac. 2019. Measuring the business value of recommender systems. *ACM Transactions on Management Information Systems*. <https://doi.org/10.1145/3370082>.
- Karatzoglou, A., X. Amatriain, L. Baltrunas, and N. Oliver. 2010. Multiverse recommendation: N-dimensional tensor factorization for context-aware collaborative filtering. In 4th ACM conference on recommender systems (RecSys), pp. 79–86, Barcelona, Spain.
- Khrouf, H., and R. Troncy. 2013. Hybrid event recommendation using linked data and user diversity. In 7th ACM conference on recommender systems (RecSys), pp. 185–192.
- Li, L., W. Chu, J. Langford, and R. E. Schapire. 2010. A contextual-bandit approach to personalized news article recommendation. In 19th International conference on world wide web, pp. 661–670, Raleigh, North Carolina, USA.
- Lieberman, H. 1995. Letizia: An agent that assists web browsing. In 14th International joint conference on artificial intelligence (IJCAI).
- Linden, G., B. Smith, J. York. 2003. Amazon.Com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1):76–80
- Ludewig, M., and D. Jannach. 2018. Evaluation of session-based recommendation algorithms. *User Modeling and User-Adapted Interaction* 28 (4–5): 331–390.
- Melville, P., R. Mooney, and R. Nagarajan. 2002. Content-boosted collaborative filtering for improved recommendations. In 18th National conference on artificial intelligence (AAAI), pp. 187–192.
- Monti, D., E. Palumbo, G. Rizzo, R. Troncy, and M. Morisio. 2018. Semantic trails of city explorations: How do we live a city. arXiv 1812.04367. <http://arxiv.org/abs/1812.04367>.
- Palumbo, E., G. Rizzo, and R. Troncy. 2017. Entity2Rec: Learning user-item relatedness from knowledge graphs for top-N item recommendation. In 11th ACM conference on recommender systems (RecSys), pp. 32–36, Como, Italy.
- Panniello, U., A. Tuzhilin, M. Gorgoglione, C. Palmisano, and A. Pedone. 2009. Experimental comparison of pre- vs. post-filtering approaches in context-aware recommender systems. In 3rd ACM conference on recommender systems (RecSys), pp. 265–268.
- Paulheim, H. 2017. Knowledge graph refinement: A survey of approaches and evaluation methods. *Semantic Web* 8 (3): 489–508.
- Rendle, S. 2010. Factorization machines. In IEEE international conference on data mining (ICDM), pp. 995–1000.
- Rendle, S., C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In 25th Conference on uncertainty in artificial intelligence (UAI), pp. 452–461.
- Resnick, P., and H.R. Varian. 1997. Recommender systems. *Communication of the ACM* 40 (3): 56–58.
- Rohde, D., S. Bonner, T. Dunlop, F. Vasile, and A. Karatzoglou. 2018. RecoGym: A reinforcement learning environment for the problem of product recommendation in online advertising. In International workshop on offline evaluation for recommender systems (REVEAL).
- Sarwar, B., G. Karypis, J. Konstan, and J. Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In 10th International conference on world wide web (WWW), pp. 285–295.
- Song, W., Z. Duan, Z. Yang, H. Zhu, M. Zhang, and J. Tang. 2019. Explainable knowledge graph-based recommendation via deep reinforcement learning. arXiv 1906.09506. <https://arxiv.org/abs/1906.09506>.
- Sun, Z., J. Yang, J. Zhang, A. Bozzon, L.-K. Huang, and C. Xu. 2018. Recurrent Knowledge Graph Embedding for Effective Recommendation. In 12th ACM conference on recommender systems (RecSys), pp. 297–305, Vancouver, British Columbia, Canada.
- Sutton, R.S., and A.G. Barto. 2018. *Reinforcement Learning: An Introduction*. A Bradford Book. Cambridge: The MIT Press.
- Taghipour, N., and A. Kardan. 2008. A hybrid web recommender system based on q-learning. In ACM symposium on applied computing (SAC), pp. 1164–1168, Fortaleza, Ceara, Brazil.
- Xiao, J., H. Ye, X. He, H. Zhang, F. Wu, and T.-S. Chua. 2017. Attentional factorization machines: Learning the weight of feature interactions via attention networks. In 26th International joint conference on artificial intelligence (IJCAI), pp. 3119–3125.
- Yao, W., J. He, G. Huang, J. Cao, and Y. Zhang. 2015. A Graph-based model for context-aware recommendation using implicit feedback data. *World Wide Web* 18 (5): 1351–1371.
- Zhang, F., N. J. Yuan, D. Lian, X. Xie, and W.-Y. Ma. 2016. Collaborative Knowledge Base Embedding for Recommender Systems. In 22nd ACM SIGKDD international conference on knowledge discovery and data mining, pp. 353–362.
- Zhang, S., L. Yao, A. Sun, and Y. Tay. 2019. Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys* 52: 1. <https://doi.org/10.1145/3285029>.
- Zhao, X., L. Xia, L. Zhang, Z. Ding, D. Yin, and J. Tang. 2018. Deep reinforcement learning for page-wise recommendations. In 12th ACM conference on recommender systems (RecSys), pp. 95–103, Vancouver, British Columbia, Canada.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Amine Dadoun is an industrial PhD student in Data Science at EURECOM in collaboration with Amadeus. His research focuses on recommender systems applied to the travel domain, including Deep Learning based Recommender Systems, Knowledge Graph & Word



Embeddings for Recommender Systems, and Session Based Recommender Systems.

Michael Defoin-Platel is Head of Machine Learning Services at Amadeus. He received a PhD in Artificial Intelligence from the University of Nice. Michael has been conducting academic research in fundamental and applied AI during 10 years before moving to the industry. He is now responsible for the development of AI in Amadeus products.

Thomas Fiig is the Director, Chief Scientist at Amadeus. He holds a Ph.D. in theoretical physics and mathematics from the University of Copenhagen, Denmark. He has more than 20 years of revenue management experience and has published several papers on forecasting, offer optimization, reinforcement learning.

Corinne Landra is Associate Director Expert at Amadeus. She has more than 10 years's experience on shopping and pricing solutions serving both the direct and indirect channels. In her current role,

Corinne is focusing on designing the future of Airline Offer Management System, defining as part of it, the architecture around pricing and product dynamicity. She represents Amadeus at ATPCO and IATA New Distribution Capability (NDC) industry fora.

Raphaël Troncy is an Associate Professor at EURECOM. He got his PhD from Grenoble University (France) in 2004 and had research positions at INA (France), CNR (Italy) and CWI (The Netherlands). He published nearly 200 scientific papers in various journals, top tier conferences and associated workshops and held key organisational roles in multiple world-leading conferences. He will be General Chair of The Web Conference in 2022 He is an expert in knowledge engineering, information extraction and recommender systems. He is the primary investigator of many National and European projects where semantic technologies and information extraction are used together to build knowledge graphs, exploratory search engines and recommender systems.

