



**HAL**  
open science

## Banner Personalization for e-Commerce

Ioannis Maniadis, Konstantinos N. Vavliakis, Andreas L. Symeonidis

► **To cite this version:**

Ioannis Maniadis, Konstantinos N. Vavliakis, Andreas L. Symeonidis. Banner Personalization for e-Commerce. 15th IFIP International Conference on Artificial Intelligence Applications and Innovations (AIAI), May 2019, Hersonissos, Greece. pp.635-646, 10.1007/978-3-030-19823-7\_53 . hal-02331309

**HAL Id: hal-02331309**

**<https://inria.hal.science/hal-02331309v1>**

Submitted on 24 Oct 2019

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# Banner Personalization for e-commerce

Giannis Maniadis<sup>1</sup>, Konstantinos N. Vavliakis<sup>1,2</sup>, and Andreas L. Symeonidis<sup>1</sup>

<sup>1</sup> Department of Electrical and Computer Engineering, Aristotle University of Thessaloniki, Thessaloniki GR54124, Greece

<https://issel.ee.auth.gr/>

<sup>2</sup> Pharm24.gr, Dafni Lakonias, GR23057, Greece

[www.pharm24.gr](http://www.pharm24.gr)

{johnmaniadis,kvavliak,asymeon}@issel.ee.auth.gr

**Abstract.** Real-time website personalization is a concept that is being discussed for more than a decade, but has only recently been applied in practice, according to new marketing trends. These trends emphasize on delivering user-specific content based on behavior and preferences. In this context, banner recommendation in the form of personalized ads is an approach that has attracted a lot of attention. Nevertheless, banner recommendation in terms of e-commerce main page sliders and static banners is even today an underestimated problem, as traditionally only large e-commerce stores deal with it. In this paper we propose an integrated framework for banner personalization in e-commerce that can be applied in small-medium e-retailers. Our approach combines topic-models and a neural network, in order to recommend and optimally rank available banners of an e-commerce store to each user separately. We evaluated our framework against a dataset from an active e-commerce store and show that it outperforms other popular approaches.

**Keywords:** Banner Personalization · e-Commerce · LDA · ANN.

## 1 Introduction

Personalization allows e-retailers to understand consumers' concerns and propose the right products, at the right time, while not spending exaggerated effort to understand each customer's wills and needs. Personalization techniques became increasingly popular in recent years and are considered key elements in the modern e-commerce world. Many ecommerce stores use large banner or billboard graphics in their designs. Visitors make decisions about a new site in less than a second, so creating a positive first impression and making sure visitors know exactly that they're in the right place is crucial for e-commerce. Banners allow directing customers to the most compelling offers and increase attention to a specific product or brand, resulting to users being showered with hundreds of advertisements on a daily basis. As a result, they often pay little attention to banners appearing on a web page, unless there is increased correspondence between their interests and the subject of the displayed advertisement. But, not

all customers are interested in all offers, as gender, age, past views and purchases are only some of the factors that affect users' interests. Thus, increasing the correlation between user interests and the subject of the displayed banners is necessary. Increasing this correspondence can be achieved by personalizing the banners displayed to each customer, which in turn is expected to increase effectiveness [2]; the right person should receive the right message.

In the last years, personalized advertising, also known as interest-based advertising, has attracted a lot of attention. It is practically the main tool for improving advertising relevance for users and for increasing return-of-investment for advertisers. Corporations like Google, Amazon and Alibaba are making inferences about a user's interests based on the sites and products they visit and applications they use and then target their campaigns according to these interests, providing an improved experience for both users and advertisers. At first glance, matching users with advertisements they like may sound like a simple problem, however the task of developing an efficient recommender system is extremely challenging. Training, optimizing, evaluating and deploying personalization solutions require specialized expertise in analytics, applied machine learning, software engineering, and systems operations to name a few.

In this paper we propose a compact banner recommendation framework that is appropriate for e-commerce SMSs. Our work focuses to small and medium sized e-commerce sites that, on the one hand have enough traffic for data analysis (0.1-10 million pageviews/month), but on the other hand, have limited processing and manpower resources, thus complex or resource inefficient solutions are not appropriate. Our focus lies especially on banners, sliders or fixed-position billboard graphics, rather than ads such as Google Ads. The main difference is that the available banners in an e-commerce store is limited and usually their position is fixed, whereas the number of available ads is enormous, thus near real-time decisions should be taken for matching thousands or millions of users with millions of advertisements. On the other hand, banner recommendations require matching thousands or in a case of a large e-commerce store, millions of users with a few dozen banners, where the main goal is to display the most interesting banners in the right order (in case of a slider) to users. Towards this end, we propose a banner recommendation framework that clusters products based on their textual representation and user's session-based information like clicks and purchases. The result of this process is used for training a neural network that optimally matches available banners with users. In addition, we propose four new evaluation metrics, to act as key performance indicators (KPIs), as traditional classification and ranking metrics cannot fully describe the effectiveness of banner recommendation methods for SMEs.

The remainder of this paper is structured as follows: Related work on personalization and banner recommendation is discussed in Section 2. Section 3 introduces four KPIs for evaluation in case of banner recommendations, while Section 4 discusses in detail the proposed framework for banner recommendations in e-commerce and evaluation takes place in Section 5. Finally, Section 6 summarizes work done, discusses future work and concludes the paper.

## 2 Related Work

A profusion of papers has been written laying down the fundamental principles of web personalization and surveying the most popular approaches in an effort to define a taxonomy of existing approaches in the field of e-commerce. In [19], Salonen and Karjaluoto discuss areas of interest, research gaps and future directions for web personalization. Their work includes a literature review of the top 20 marketing and information systems journals published during the period of 2005–2015 that shows active research output and the domination of information systems publications. A number of studies have shown that web personalization is well worth the investment in capital and work hours, as it has been observed that website personalization can lead to increased user satisfactions, increased revenue and customer loyalty [20]. Other surveys [1] have indicated that web personalization is a major priority for organizations with an online presence.

Personalization naturally starts with collecting information about the individuals one wishes to reach. This information can be explicit or implicit. Explicit information is what a user offers directly and includes ratings, answers to questionnaires, reviews, personal information etc. Implicit information refers to data that the website gathers based on a user's interactions with the website. This process is a discipline of Web Mining [12] and includes monitoring, storing and processing a user's actions in order to reveal potentially useful underlying data (e.g. browsing and/or purchasing patterns). In this line of research there are many approaches in terms of the data that are collected, the ways by which that is done, and how the collected data may be used [17]. Nevertheless, this type of mining has become controversial in the last years, as it can be regarded intrusive in terms of the amount of user data collected with limited or without user consent. The lack of transparency raises concerns about privacy and security, thus recent regulations such as GDPR [7] try to adjust the excessive use of personal data. Although various aspects of the legislation and their implementation remain uncertain [21], the fact remains that various practices mentioned in previous sources might have to be revised to allow for user consent, or even be rejected altogether as non-compliant with the new regulatory frameworks.

As far as the problem of personalized recommendations is concerned, the main objective is to use available data to create recommendations that promote different content to every individual, according to his or her distinct interests. Several papers have been written that analyze and survey recommender systems in the field of web personalization [5]. Furthermore, various approaches have been proposed in the context of e-commerce, such as statistical approaches [18] and machine learning frameworks. Algorithms such as Collaborative Filtering [24], Deep Learning Neural Networks [10] and SVM models [14] have been widely used in personalized content recommendation systems. Although most of these methods have shown promising results, other methods have focused on the cold-start problem, that is the case where no prior knowledge exists (e.g. the first visit of a new customer). In this case association rules [15] have been used with satisfactory results. Dealing with infrequent items is also another problem that requires special treatment. These items usually are treating as outliers and

excluded in the preprocessing phase [13]. Another interesting problem in the field are Top- $N$  recommenders, which aim to produce  $N$  ordered recommendations, rather than a single one. Top- $N$  recommender algorithms sometimes are considered a distinct problem due to sometimes increased computational load required, as well as differences in the ways they are evaluated [11].

On the other hand, recommending advertisements has attracted increased attention during the past years. Yuan and Tsao [25] proposed a personalized contextualized mobile advertising infrastructure for the advertisement of commercial and non-commercial activities that minimize users' inputs by using implicit browsing behaviors and recommends Top- $N$  scored advertisements using a two-level Neural Network. In a different approach, Wang et al. [23] propose an interactive service recommendation based on contextual search by building an ad domain-based concept hierarchy to make the most of the product details over the e-commerce sites. In [3] Anastasakos et al. explore search engine logs to determine the relevance of an ad document for a search query by using a graph with several million edges. Last but not least, Papadopoulos et al. [16] study the problem of recommending a small set of ads to a user based solely on the currently viewed web page and propose the use of lexical graphs created from web corpora as a means of computing improved content similarity metrics between ads and web pages.

Obviously, a lot of progress has been made in personalization and recommendation systems; however, most of the existing solutions focus on product or ads recommendation. In our case, we deal with the problem of e-commerce banner recommendation, especially focusing in small and medium e-commerce sites. In this case the pool of available recommendations is limited and available banner positions are fixed, while available data and resources are also limited. Our proposed solution is a Top- $N$  recommendation framework for selecting the most appropriate banners in e-commerce. Advancing the state-of-the-art we propose an integrated solution that combines topic analysis and neural networks, a combination that not only achieves better results, but is also more adaptable to data changes.

### 3 Proposed Evaluation Metrics

To measure the effectiveness of recommender systems and compare different approaches, various metrics have been applied such as Classification Accuracy, Logarithmic Loss, Confusion Matrix, Area under Curve, Recall, F1 Score, Mean Absolute Error and Mean Squared Error [22]. Nevertheless, our problem presents some differences from typical classification or ranking problems, as in our case there are difficulties in separating good and bad recommendations, due to the fact that available classes and/or banners for recommendation may be limited and may not match the user's interests. Thereafter, in our case we may have to recommended the best of a few bad choices. For this reason, we propose four new metrics to act as KPIs (Key Performance Indicators) in our problem: a) Recall-Top $N$ , b) Recall-Purchases, c) Min-1 and d) Pos-Error. We argue that

these new metrics are more suitable for evaluating the performance of banner recommendations systems and we believe can improve the evaluation procedure.

**a) Recall-TopN** In the field of Information Retrieval, Recall (also known as sensitivity) is the fraction of relevant instances that have been retrieved over the total amount of relevant instances. We propose a variation of the Recall metric that we call Recall-TopN (Equation 1), where  $N$  is the number of recommendations,  $S$  the total number of sessions,  $P_i$  the number of categories that the system predicted correctly (true positives) for session  $i$  and  $C_i$  the number of categories the user showed interest for (relevant elements).

$$Recall - TopN = \frac{1}{S} * \sum_{i=1}^S \frac{P_i}{\max(C_i, N)} \quad (1)$$

This adjustment is necessary, as even if the user is interested in multiple categories, there will be only  $N$  recommendations regardless of the number of relevant items. Traditionally, recall would penalize this behavior as in our case the number of true positives cannot never exceed  $N$ .

**b) Recall-Purchases** Next, we propose Recall-Purchases (Equation 2), a metric that intends to evaluate the percentage of recommendations that potentially can lead to actual conversions. In Equation 2,  $M$  is the number of sessions that lead to a purchase,  $B_i$ , is the number of categories of which products were purchased for session  $i$ , and  $P_i$  is the number of categories recommended. Recall-Purchases may be more important than Recall-TopN, as conversions is the definite target in e-commerce.

$$Recall - Purchases = \frac{1}{M} * \sum_{i=1}^M \frac{P_i}{\max(B_i, N)} \quad (2)$$

**c) Min-1** Another KPI (Key Performance Indicator) that is necessary in e-commerce is determining where a recommendation set is useful or completely irrelevant. Towards this end we define Min-1, which is the percentage of sessions where at least one of Top- $N$  recommended categories was interesting for the user. While, not a particularly strict criterion, it is important in determining how many recommendation sets are even minimally useful.

**d) Pos-Error** Recommendation systems should not only provide relevant recommendations, but also rank them properly. Thus, we propose a metric we call Pos-Error that calculates the degree to which recommendations are not only relevant, but also optimally ranked. An intuitive description would be that Pos-Error increases when all banners belonging to user's top- $N$  interests are optimally ranked. Pos-Error is defined in Algorithm 1, where  $S$  is the number of

sessions,  $C$  is the number of categories,  $P$  is a  $NXC$  matrix with the predictions of the recommender regarding user interest per session per tag, and  $C$  is a  $NXC$  matrix with the actual user interested displayed in the corresponding sessions.

---

**Algorithm 1** Pos-Error
 

---

```

Error=[]
For i=1 to S do
  Ps = argsort ( P[i,] )
  Cs = argsort ( C[i,] )
  For j=1 to N do
    p = arg ( Ps== Cs[j] )
    if j > p do
      Append to error log(P(c,s)-R(c,s) )/log(R(C,s))
Pos-Error = [mean(Error)]-1

```


---

According to Algorithm 1, Pos-Error increases as the distance between predicted and the ground truth positions increases. Pos-Error captures the fact that prediction mismatches about categories the user is more interested are more important than prediction mismatches for less preferred categories. Logarithmic values are used to smooth out results.

## 4 Proposed Framework

In this Section we propose an integrated framework for banner recommendations as depicted in Figure 1. Our framework processes product textual information, user history and categories (classes) that users express interest for. Latent Dirichlet allocation (LDA) [4] is used employed for clustering products based on their textual description (name, description etc.). Then products clusters and user history data are processed in order to calculate user’s session-based interest to specific categories. Finally, we use a neural network that processes the session-based interests and produces banner recommendations.

In every e-commerce system there are data that consist of user ids (customer ids or cookie ids in case of registered or unregistered visitors respectively), products, user pageviews, purchases and categories. Moreover, products are categorized based on brand, attributes or general interest, so by categories (classification classes) we mean the general user interests (brand, category or class). First, we apply LDA for extracting semantic information from products and cluster them into topic themes, based on the similarities of their textual descriptions. Topic modeling is based on the assumption that each document  $d_i$  is described as a random mixture of topics and each topic as a focused multinomial distribution over terms. LDA builds a set of thematic topics from terms that tend to co-occur in a given set of documents. The result of the process is a set of  $N_\Theta$  topics, each expressed with a set of  $N_W$  terms. The number of topics  $N_\Theta$  and the number of



**Fig. 1.** Proposed Framework for Banner Recommendations

terms per topic  $N_W$  have to be defined in advance. These two parameters can be used to adjust the degree of specialization of the latent topics.

LDA discovers a mixture of topics  $P(\theta|d)$  for each document  $d_i$ , where each topic is described as a mixture of terms  $P(w|\theta)$  following another probability distribution as given in Equation 3. The probability of the  $w_i$  term describing a given document  $d$  is  $P(w_i|d)$ , where  $\theta_i$  is the latent topic, and  $P(w_i|\theta_{i=j})$  is the probability of  $w_i$ , within topic  $j$ . The probability of picking a term from topic  $j$  in a document is  $P(\theta_i = j|d)$ .

$$P(w_i) = \sum_{j=1}^{N_\theta} P(w_i|\theta_{i=j})P(\theta_i = j|d) \quad (3)$$

LDA estimates the topic-term distribution  $P(w|\theta)$  and the document-topic distribution  $P(\theta|d)$  from an unlabeled corpus of documents using Dirichlet priors for the distributions and a fixed number of topics. The Gibbs sampler [6] iterates multiple times over each term  $w_i$  in document  $d_i$  and samples a new topic  $j$ .

$$P(\theta_{i=j}|w_i, d_i, \theta_{-i}) \propto \frac{C_{W_{ij}}^{W\theta} + \beta}{\sum_w C_{W_j}^{W\theta} + W\beta} x \frac{C_{d_{ij}}^{D\theta} + \alpha}{\sum_\theta C_{d_i\theta}^{D\theta} + \theta\alpha} \quad (4)$$

The Gibbs sampler sets the complexity of topic modeling to  $O(N_\Theta NI)$ , where of  $N_\Theta$  is the number of topics,  $N$  the number of documents (movies) and  $I$  is the number of the sampler's iterations. The most common way to evaluate a probabilistic model is to measure the log-likelihood of a held-out test set, thus we use perplexity [4] for determining the optimal number of clusters, which is monotonically decreasing in the likelihood of the test data, and is algebraically equivalent to the inverse of the geometric mean per-word likelihood. A lower perplexity score indicates better generalization performance. More formally, for a test set of  $N$  documents, the perplexity is described in Equation 5.

$$perplexity = \exp \left\{ - \frac{\sum_{i=1}^N \log(p(w_i))}{\sum_{i=1}^N d_i} \right\} \quad (5)$$



After we assign each product to a topic cluster via LDA, we group user actions into sessions. We use an arbitrary time limit of  $T$  hours, which means that a session consists of actions of the same user, all of which happened within  $T$  hours of the previous one. For each session we count the instances a category was viewed. Thus, for each session a vector of length  $C$  is created, with each slot containing the number of times the respective category or product was viewed. In case of product views, we consider that products are assigned to one or more categories. Thus, a matrix  $Vcat$  is created with dimensions  $(N, C)$ , where  $N$  is the number of sessions examined and  $C$  is the number of categories. Following the same process, purchase statistics are calculated and matrix  $Pcat$  is created (Algorithm 2).

---

**Algorithm 2** Calculate session-based Category interest
 

---

```

For i=1 to Number of Sessions do
  V = Category Pageviews related to session i
  P = Category Purchases related to session i
  For j=1 to C do
    Vcat [i,j] = Number of V entries related to category j
    Pcat [i,j] = Number of P entries related to category j
  
```

---

Having created matrices  $Vcat$  and  $Pcat$  for pageviews and purchases respectively, we repeat the process for views and purchases of products with regard to their assignments to LDA clusters. Each session is assigned two vectors of zeros, one for views and one for purchases, where each slot is related to an LDA group. For each view or purchase, the respective LDA vector is added to the total views and purchases of the accordingly. Thus, two more matrixes  $Vlda$  and  $Plda$  are created according to Algorithm 3.

---

**Algorithm 3** Calculate Category interest based on LDA clusters
 

---

```

For i=1 to Number of Sessions do
  Pv = Product Pageviews related to session i
  Pb = Product Purchases related to session i
  For j=1 to Sizeof Pv do
    Lda_var = LDA memberships for product Pv [j]
    Vlda [i] += Lda_var
  For j=1 to Sizeof Pb do
    Lda_var = LDA memberships for product Pb [j]
    Plda [i] += Lda_var
  
```

---

Next, we normalize each matrix in order to contain values in the range  $[0,1]$  according to the relevant interest degree for each category in every session. After that,  $Vcat$  and  $Pcat$  are aggregated according to Equation 6 where  $\alpha$  and  $\beta$

are weight factors as purchases are a stronger signal than product views. Then, matrices  $Scat$  and  $Slda$  are concatenated to form matrix  $S$ .

$$\text{For each session } i : \overline{Scat}_i = \alpha * \overline{Vcat}_i + \beta * \overline{Pcat}_i \quad (6)$$

Having calculated session-based category interest for each user, we estimate user interests taking into account the entire user-history as the mean category interest for each session, which results to a matrix  $H$  that we use as input of the neural network for predicting the output matrix  $Scat$ . In other words, the proposed framework will predict user interests at session  $s$  using the available information of that user’s actions before session  $s$ , as described by their historical data in array  $H$ .

## 5 Evaluation

In this Section we evaluate the proposed solution using a dataset coming from a live e-commerce store. Our evaluation dataset contains anonymous data of 4,292,495 unique clicks, 10,167 products, 83,467 unique customers and 146,803 product purchases. After some preprocessing, we grouped available data into sessions, where in our case each session consisted of actions of the same user for a period of  $T = 2$  hours. We also clustered products into 50 topic-clusters using LDA after evaluating the optimal number of clusters using the perplexity metric. In our case we found out that 13 out of the available 68 categories (classes) were too generic, covering more than 80% of the products, so we treated these 13 categories as outliers and removed them, the number of products remained unchanged. Then, we performed a five-fold cross validation evaluation test for determining the optimal hyper-parameters of our neural network. Our neural network, built with Python and Tensorflow, used the Relu activator and neuron weights initialized with the Glorot initialization function [9]. We tested hyperparameter values in order to find the optimal combination using various combinations of the following hyperparameter values:

- Learning Rate = {0.1, 0.01, 0.001, 0.0001, 0.00001, 0.000001}
- Layers = {2, 3, 4, 5, 6, 10}
- Neurons per Layer = {236, 354, 472, 590, 708}
- Regularization Factor = {0.1, 0.01, 0.001}
- Loss Function = {Mean Squared Error, Categorical Cross Entropy}
- Optimizer = {Stochastic Gradient Descent, Adam}
- Batch Size = {64, 32, 16}

A sample of our experiment set is available in Table 1, the duration of each ANN training run was between 90 and 120 minutes, depending on the learning rate, number of layers and neurons per layer, on an Intel i7 6700K computer with a GeForce GTX 970. The following hyperparameters: Learning Rate=0.0001, Layers=3, Neurons=590, Loss Function=Categorical Cross

entropy, Optimizer=Adam, Batch Size=32 and Regularization Factor=0.1 resulted in Recall Top- $N = 0.333$ , which was the best result in our tests, thus we used these values in the next step.

**Table 1.** Evaluation of various hyperparameters for our neural network.

Learning Rate	Layers	Neurons per Layer	Regularization Factor	Loss Function	Optimizer	Batch Size	Recall Top-N
0.100000	2	236	0.100	CCE	SGD	32	0.3179
0.100000	3	354	0.010	CCE	SGD	16	0.3097
0.100000	4	472	0.100	MSE	Adam	32	0.0665
0.010000	6	708	0.100	CCE	Adam	32	0.2849
0.010000	6	708	0.001	CCE	Adam	16	0.2377
0.001000	2	236	0.001	CCE	SGD	32	0.2908
0.001000	2	354	0.001	MSE	Adam	16	0.3118
0.000100	4	472	0.100	CCE	Adam	32	0.3290
0.000010	3	354	0.100	MSE	Adam	32	0.2876
0.000001	3	354	0.100	CCE	Adam	32	0.2840

Next, we evaluated our proposed framework against four approaches. First, we wanted to test more machine learning approaches, so we replaced the neural network in the processing step with a random forest (with 100 estimators and a maximum depth of 10), as well as  $k$ -nearest neighbors (KNN) [8] algorithm. Finally, we evaluated our approach against two baseline solutions, we set  $N = 5$  most popular categories and then randomly selected  $N = 5$  categories. According to the results of our tests which are depicted in Table 2, our approach achieves higher scores in all four evaluation metrics, up to 17.45% in case of random forests and up to 77.36% in case of randomly showing banners, which is the case for the majority of existing e-commerce sites.

**Table 2.** Evaluation of various hyperparameters for our neural network.

	NN		Ran. Forest		KNN		Popular		Random	
Metric	Value	Value	% Diff.	Value	% Diff.	Value	% Diff.	Value	% Diff.	
Recall-TopN	0.3803	0.3640	-4.29%	0.3153	-17.08%	0.2806	-26.21%	0.1016	-73.28%	
Recall-Purch.	0.4597	0.3795	-17.45%	0.3474	-24.44%	0.2417	-47.42%	0.1041	-77.36%	
Min-1	0.5284	0.5205	-1.48%	0.4756	-9.99%	0.4441	-15.95%	0.18386	-65.21%	
Pos-Error	0.5391	0.5326	-1.20%	0.4959	-8.01%	0.4875	-9.57%	0.4022	-25.39%	

## 6 Conclusion

In this paper, we introduced a compact integrated system for personalized banner recommendations that is specially targeted to e-commerce stores. We proposed

a framework that employs latent dirichlet allocation and neural networks for processing user history and product’s textual information in order to produce and rank recommendations from a limited pool of available banners that can be displayed in specific banner positions of an e-commerce store. We evaluated the proposed approach on a dataset coming from an active e-commerce store against other approaches and showed that our solution can provided up to 17.45% better results than random forest methods and improve up 77.36% than baseline solutions, which is the case for most small and medium e-commerce stores. Our future work includes further experimentation on discovering the optimal parameters for LDA and our neural network. We also plan to evaluate our approach against more recommendation algorithms, as well as test the proposed solution live in an e-commerce store via A/B testing.

## Acknowledgements

This work was partially funded by an IKY scholarship funded by the “Strengthening of Post-Academic Researchers” Act from the resources of the OP “Human Resources Development, Education and Lifelong Learning” with Priority Axes 6,8,9 and co-funded by the European Social Fund ECB and the Greek government.

## References

1. Trends in personalization. Tech. Rep. MSU-CSE-06-2, Researchscape International (2018), <https://www.evergage.com/wp-content/uploads/2018/04/Evergage-2018-Trends-in-Personalization-Survey.pdf>
2. Adams, R.: Intelligent advertising. *AI & SOCIETY* **18**(1), 68–81 (Jan 2004)
3. Anastasakos, T., Hillard, D., Kshetramade, S., Raghavan, H.: A collaborative filtering approach to ad recommendation using the query-ad click graph. In: Proceedings of the 18th ACM Conference on Information and Knowledge Management. pp. 1927–1930. CIKM ’09, ACM, New York, NY, USA (2009)
4. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. *J. Mach. Learn. Res.* **3**, 993–1022 (Mar 2003)
5. Bobadilla, J., Ortega, F., Hernando, A., Gutiérrez, A.: Recommender systems survey. *Knowledge-Based Systems* **46**, 109 – 132 (2013)
6. Casella, G., George, E.I.: Explaining the gibbs sampler. *The American Statistician* **46**(3), 167–174 (1992)
7. Dato, A.: Data in the post-gdpr world. *Computer Fraud & Security* **2018**(9), 17 – 18 (2018)
8. Dudani, S.A.: The distance-weighted k-nearest-neighbor rule. *IEEE Transactions on Systems, Man, and Cybernetics* **SMC-6**(4), 325–327 (April 1976)
9. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: Teh, Y.W., Titterton, M. (eds.) Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research, vol. 9, pp. 249–256. PMLR, Chia Laguna Resort, Sardinia, Italy (13–15 May 2010)

10. Hidasi, B., Karatzoglou, A.: Recurrent neural networks with top-k gains for session-based recommendations. In: Proceedings of the 27th ACM International Conference on Information and Knowledge Management. pp. 843–852. CIKM '18, ACM, New York, NY, USA (2018)
11. Hu, J., Liang, J., Kuang, Y., Honavar, V.: A user similarity-based top-n recommendation approach for mobile in-application advertising. *Expert Systems with Applications* **111**, 51 – 60 (2018)
12. Kosala, R., Blockeel, H.: Web mining research: A survey. *SIGKDD Explor. Newsl.* **2**(1), 1–15 (Jun 2000)
13. Lin, K.L., Hsu, J.Y., Huang, H.S., Hsu, C.N.: A recommender for targeted advertisement of unsought products in e-commerce. In: Seventh IEEE International Conference on E-Commerce Technology (CEC'05). pp. 101–108 (July 2005)
14. Lin, K.L., Hsu, J.Y.j., Huang, H.S., Hsu, C.N.: A recommender for targeted advertisement of unsought products in e-commerce. In: Proceedings of the Seventh IEEE International Conference on E-Commerce Technology. pp. 101–108. CEC '05, IEEE Computer Society, Washington, DC, USA (2005)
15. Osadchiy, T., Poliakov, I., Olivier, P., Rowland, M., Foster, E.: Recommender system based on pairwise association rules. *Expert Systems with Applications* **115**, 535 – 542 (2019)
16. Papadopoulos, S., Menemenis, F., Kompatsiaris, Y., Bratu, B.: Lexical graphs for improved contextual ad recommendation. In: Boughanem, M., Berrut, C., Mothe, J., Soule-Dupuy, C. (eds.) *Advances in Information Retrieval*. pp. 216–227. Springer Berlin Heidelberg (2009)
17. Pierrakos, D., Paliouras, G., Papatheodorou, C., Spyropoulos, C.D.: Web usage mining as a tool for personalization: A survey. *User Modeling and User-Adapted Interaction* **13**(4), 311–372 (Nov 2003)
18. Qudah, D.A.A., Cristea, A.I., Bazdarevic, S.H., Al-Saqqa, S., Rodan, A., Yang, W.: Personalized e-advertisement and experience: Recommending user targeted ads. In: 2015 IEEE 12th International Conference on e-Business Engineering (ICEBE). vol. 00, pp. 56–61 (Oct 2015)
19. Salonen, V., Karjaluoto, H.: Web personalization: The state of the art and future avenues for research and practice. *Telematics and Informatics* **33**(4), 1088 – 1104 (2016)
20. Senecal, S., Nantel, J.: The influence of online product recommendations on consumers' online choices. *Journal of Retailing* **80**(2), 159 – 169 (2004)
21. Sirur, S., Nurse, J.R.C., Webb, H.: Are we there yet? understanding the challenges faced in complying with the general data protection regulation (GDPR). *CoRR* **abs/1808.07338** (2018)
22. Tsoumakas, G., Katakis, I.: Multi-label classification: An overview. *Int J Data Warehousing and Mining* **2007**, 1–13 (2007)
23. Wang, J., Wang, B., Duan, L.y., Tian, Q., Lu, H.: Interactive ads recommendation with contextual search on product topic space. *Multimedia Tools and Applications* **70**(2), 799–820 (May 2014)
24. Wang, T., Ren, Y.: Research on personalized recommendation based on web usage mining using collaborative filtering technique. *WSEAS Trans. Info. Sci. and App.* **6**(1), 62–72 (Jan 2009)
25. Yuan, S.T., Tsao, Y.: A recommendation mechanism for contextualized mobile advertising. *Expert Systems with Applications* **24**(4), 399 – 414 (2003)