Identifying Reliable Recommenders in Users' Collaborating Filtering and Social Neighbourhoods

Dionisis Margaris, Dimitris Spiliotopoulos, and Costas Vassilakis

Abstract Recommender systems increasingly use information sourced from social networks to improve the quality of their recommendations. However, both recommender systems and social networks exhibit phenomena under which information for certain users or items is limited, such as the cold start and the grey sheep phenomena in collaborative filtering systems and the isolated users in social networks. In the context of a social network-aware collaborative filtering, where the collaborating filtering- and social network–based neighbourhoods are of varying density and utility for recommendation formulation, the ability to identify the most reliable recommenders from each neighbourhood for each user and appropriately combine the information associated with them in the recommendation computation process can significantly improve the quality and accuracy of the recommendations offered. In this chapter, we report on our extensions on earlier works in this area which comprise of (1) the development of an algorithm for discovering the most reliable recommenders of a social network recommender system and (2) the development and evaluation of a new collaborative filtering algorithm that synthesizes the opinions of a user's identified recommenders to generate successful recommendations for the particular user. The proposed algorithm introduces significant gains in rating prediction accuracy (4.9% on average, in terms of prediction MAE reduction and 4.2% on average, in terms of prediction RMSE reduction) and outperforms other algorithms. The proposed algorithm, by design, utilizes only basic information from the collaborative filtering domain (user–item ratings) and the social network domain (user relationships); therefore, it can be easily applied to any social network recommender system dataset.

D. Margaris
Department of Digital Systems, University of the Peloponnese, Sparti, Greece
e-mail: margaris@uop.gr

D. Spiliotopoulos · C. Vassilakis (✉)
Department of Management Science and Technology, University of the Peloponnese, Tripoli, Greece
e-mail: dspiliot@uop.gr; costas@uop.gr

M. Çakırtaş, M. K. Ozdemir (eds.), *Big Data and Social Media Analytics*, Lecture Notes in Social Networks, https://doi.org/10.1007/978-3-030-67044-3_3

Keywords Social networks · Recommender systems · Collaborative filtering · Limited information · Near neighbours · Evaluation · Rating prediction

1 Introduction

Nowadays, due to the blistering growth of the available information on the Internet, the task of searching and finding products that may be of interest to users has become an extremely difficult task. Recommender systems (RSs) aim to overcome the information overload problem, by investigating the preferences of online users and suggesting items they might like. Many commercial web services have implemented these systems to recommend products to their clients in a more accurate manner, to increase their profits.

The most widely used approach for making recommendations, stemming from user actions and behaviour, is collaborative filtering (CF). CF synthesizes the informed opinions of people in order to make accurate user rating predictions and personalized recommendations. Since traditional CF relies only on opinions expressed by humans on items, either implicitly (e.g. a user purchases an item or clicks on an advertisement banner, which indicates a positive assessment) or explicitly (e.g. a user submits a specific rating for a particular item), its biggest advantage is that the items' explicit content description is not required [23], since, contrary to content-based RSs, the CF RSs do not recommend items similar to the ones that the users have already experienced (and rated positively). CF works on the assumption that if users had similar tastes on some items in the past (rating assignment, buying, eating, watching, etc.), then they are likely to have similar interests in the future, too [10].

Traditional CF-based RSs assume that users are independent from each other and do not consider the social interactions among them, such as friendship and trust. This approach fails to incorporate important aspects that denote interaction, influence and tie strength among them, which can substantially enhance recommendation quality [4, 11].

Social network-aware RSs take into account both static data sourced from the user profiles (e.g. gender, age and residence), as well as static data sourced from the item profiles (e.g. item price, availability and colour). These features are complemented with dynamic aspects and contextual information stemming from social information, such as user mood and social influence, as well as the item general acceptance and trends in order to supplement the traditional CF data (e.g. the aforementioned static data, as well as user ratings). By taking this information into account in the recommendation process, the social network (SN) RSs manage to achieve more successful and targeted recommendations [16].

However, in some cases the SN- and CF-based information that a RS has at its disposal may be limited: some users may not consent to the use of their SN information for recommendations or may not have SN accounts at all, or the rating data (characteristics and categories of products) may be unavailable for the RS service. And, conversely, in some cases, the CF-based near neighbours (NNs) of a user U may either be limited in number, or have low similarity, or have little utility, in the sense that they have rated very few items that U has not already rated. Generalizing, we can assert the successful combination of SN- and CF-based information effectively depends on identifying which rating prediction information source (the SN relations or the CF NNs) is considered as the most reliable and useful predictor for each individual user in a SN CF dataset. While evaluating available prediction information sources, we should take into account both (a) the characteristics of the information source (e.g. neighbourhood population, degrees of similarity and levels of influence between the user and her neighbourhood) and (b) the dynamics of the recommendation process, considering in particular the fact that—for many users—their SN relationships play an important role in their responses to recommendations, when compared to the CF NNs that traditional CF RSs use [4].

In this chapter, we propose an algorithm that can be applied to any SN-aware RS, which utilizes both the users' social relations (SN-based information) and the users' ratings on items (CF-based information) and combines them effectively to generate more successful rating predictions. The proposed algorithm addresses the issues of limited SN information or limited CF information for some users, by adapting its behaviour, taking into account the density and utility of each user's SN and CF neighbourhoods. In this context, we present and validate seven alternatives for evaluating the importance of each user's SN and CF neighbourhoods and combining the partial predictions produced by each user's SN and CF neighbourhoods.

Through this adaptation, the proposed algorithm achieves considerable improvement in rating prediction accuracy; this is verified by the results of our experiments, in which the performance of the proposed algorithm is evaluated against five contemporary SN CF datasets. In the same experiments, the performance of the proposed algorithm is compared against the performance of the algorithm presented by Margaris et al. [21], which also tackles the same problem; however assuming that all dataset users share the same prediction significance between CF and SN prediction information in RSs.

Notably, in our experiments we used:

1. Both *dense* and *sparse* SN datasets (a SN dataset density refers to the number of relations when compared to the number of users in it—[4])
2. Both *dense* and *sparse* CF datasets (a CF dataset density refers to the number of ratings when compared to the number of users and items in it—Herlocker et al. 2004)
3. Both *undirected edge* (friendships) and *directed edge* (trusts) SN datasets [11]

The experiment results show that the proposed algorithm introduces considerable prediction accuracy gains in terms of rating prediction error under all conditions (4.9% on average, in terms of prediction MAE reduction and 4.2% on average, in terms of prediction RMSE reduction). Since the proposed algorithm requires only basic SN information (user relationships), as well as basic CF information (user ratings on items), it follows that it can be applied to any SN RS dataset. It is also worth noting that the proposed algorithm can be combined with other algorithms that have been proposed for improving prediction accuracy, rating recommendation quality or prediction coverage in CF-based systems, focusing either in traditional CF-based systems (e.g. concept drift and clustering techniques—[9]) or in SN CF-based systems (influence, trust, etc.—[5]).

The rest of the chapter is structured as follows: Section 2 overviews related work, while Sect. 3 presents the SN CF prediction formulation foundations. Section 4 presents the proposed algorithm, as well as the alternatives for combining the partial predictions produced by each user's SN and CF neighbourhoods. Section 5 evaluates the proposed algorithm and, finally, Section 6 concludes the chapter and outlines future work.

2 Related Work

RSs are increasingly utilizing SN data to improve the accuracy of the recommendations offered to their users and augment recommendation variety [4, 11], as well as alleviate the issues of cold start, where it is impossible to provide personalized recommendations due to lack of information, and grey sheep, that is, users whose opinions do not agree with any other user and hence a CF RS cannot produce a recommendation [13].

In the aforementioned works, Gilbert and Karahalios [11] present a predictive model that maps SN data to tie strength, differentiating between weak and strong ties with relatively good accuracy, and illustrate how the utilization of tie strength may enhance SN design elements, including friend introductions, message routing, information prioritization and privacy controls. On the other hand, Bakshy et al. [4] investigate the effect of social influence on the consumer responsiveness to online advertisements. More specifically, Bakshy et al. [4] analyse how the presence of cues from a user's social neighbourhood affects the user's responsiveness to online advertisements, taking into account the tie strength between the user and the social connections appearing in the cues. With this, they establish the sizable effect from the inclusion of minimal social cues in advertising and quantify the positive relationship between the consumer response rates and the connection strength among users and affiliated peers appearing in social cues. He and Chu [13] mention that even if a user has no past reviews, a RS still can make recommendations to him based on the preferences of his friends, if it integrates with SNs. They designed a framework that makes recommendations based on the user's own preferences, the general acceptance of a target item and his SN friends' opinions.

Other works utilize SN data to improve the accuracy of the recommendations offered to their users. Capdevila et al. [7] present GeoSRS, a hybrid RS for a location-based SN that enables users to write short reviews about places of interest that they visit. The presented RS uses text mining, as well as geographical location information in order to recommend locations. Margaris et al. [20] propose a query personalization algorithm that exploits the browsing and rating information of items by users as well as the influence information from SNs used for personalized query adaptation. The queries were adapted by (re)writing the specification of the query sorting procedure to allow for re-ordering of data based on the projected user interest.

Yan et al. [30] propose an approach for managing the complexity of adding social relation networks to RSs. The proposed method, initially, generates an individual relationship network for each user and item, using a fitting algorithm of relationship networks to control the relationship propagation and contracting. Individual relationship networks are subsequently regularized by taking into account the taste diversity between relationship members, in order to capture the time-evolving nature of tastes and emphasize the aspect of homophily. Finally, the regularized individual relationship networks are fused into a matrix factorization algorithm to generate recommendations. Their method is generalized so it can also be applied to the item–item relationship network via item–user role switching. Pham et al. [25] introduce a social RS using memory-based CF models with user-oriented methods as basic models. This is conducted through analysis on the correlations between social relations and user interest similarities. Additionally, they employ sentiment analysis techniques to identify the top-K favourite products for each user, and this information is exploited by the social RS in the rating prediction computation process. Chamoso et al. [8] propose a relationship RS for business and employment-oriented SN. The proposed RS extracts the relevant information from the SN and utilizes it for recommendation on new contracts and job offers to users. The RS utilizes information scraped from user profiles, user activity and job offer descriptions. Then, metrics are applied to discover new ties that are likely to become relationships.

Seo et al. [26] introduce a method to calculate the friendship strength described by the closeness between users in a social circle. Moreover, they propose a personalized RS based on friendship strength to recommend topics or interests that users might have in order to analyse big social data, using Twitter. The measure that they propose can provide recommendations in multi-domain environments for a variety of topics. Zhao et al. [33] propose a rating prediction method for user-services by exploring the rating behaviour of users of social networks. They predict user-service ratings by focusing on the user rating behaviour and, more specifically, on additional rating information, such as the time the user rated the item, what kind of item it was, the user interest that could be mined from the user rating history and the manner that the user rating behaviour diffuses among the user SN relations. Moreover, they introduce the interpersonal rating behaviour diffusion factor for deep understanding of the users' rating behaviour. For the user-service rating prediction method, four factors are fused into a unified matrix-factorized framework: (a) user personal interest (related to user and item topics), (b) interpersonal interest similarity (related to user interest), (c) interpersonal rating behaviour similarity (related to user rating behaviour habits) and (d) interpersonal rating behaviour diffusion (related to user behaviour diffusions).

Yu et al. [31] present a social recommender that is based on the main idea that likability is reflected by distance. This work employs a distance metric learning approach [29] to derive a distance metric representing the relationships between users and between users and items; these distances are jointly determined by ratings and social relations. This distance metric is combined with matrix factorization, mapping items and users into a unified low-dimensional space and supporting a spatial understanding of the latent factor space and how users and items are positioned inside the space. This approach increases the placement accuracy of users with few ratings, who are ‘pulled’ close to other users that are similar to them. Finally, the learned metrics and positions are used to generate understandable and reliable recommendations.

Mukamakuza et al. [24] examine the existence of observable relationships between rating behaviour and SN connections in social recommenders. More specifically, they investigate publicly available datasets that contain both traces of rating behaviour along with a social graph. Utilizing SN analysis and statistics techniques, they examine the correlation between high rating activity and multiple item feedback. They check whether high correlation leads to SN centrality and vice versa. Ma et al. [18] propose a CF algorithm based on SN relationship and geographic information as complementary conditions for solving fundamental RS problems, such as raw data sparsity and low accuracy/efficiency. Their proposed algorithm introduces the social relation data into the matrix complementation process. This results in reduced sparsity for the original user–item rating matrix and enhances the authenticity of the data complement. Then, the user geographic information is used for filtering the information that is used for the matrix complementation. This approach lowers the data complementation error and improves the data complementation accuracy. The improvement on the recommendation efficiency and accuracy is achieved through conditional selection of the item complements.

Recently, Amato et al. [3] proposed a RS based on a ‘user-centred’ approach for recommendations for big data applications. Their approach works by processing interactions between the users and the multimedia content generated in one or more social media networks. Alahmadi and Zeng [2] present an Implicit Social Trust and Sentiment-based RS framework that mines user preferences from online SNs. Their method utilizes the typically overlooked but widely available information from SNs in RSs. Based on the fact that a user opinion is influenced considerably by the opinions of their trusted SN relations, they present a framework to personalize recommendations through the application of new data sources from mining the short text posts of the users’ friends from microblogs. The resulting Implicit Social Trust and Sentiment-based RS maps converted recommendations to numerical rating scales through three distinct measures: (1) calculation of the implicit trust between friends, based on intercommunication activities, (2) inference of the sentiment reflected from the information from friends’ short posts, called micro-reviews, using natural language for sentiment analysis, enhanced with techniques for handling online social network language features such as emoticons and Internet jargon and (3) quantification of the degree to which the level of trust between friends and sentiment from micro-reviews from friend recommendations impacts each user’s opinion, using machine learning regression algorithms, such as support-vector machines, random forests and linear regression.

All the aforementioned works necessitate the availability of additional information, either regarding the user profile (e.g. location, age and gender), the item description (e.g. price, taxonomical categorization and value for money) or the relationships between users (e.g. tie strength and social influence). In this sense, their applicability is limited when compared to the algorithm proposed in this work that requires the availability of just basic SN-sourced information (i.e. trust relationships or elementary friendship among users).

Notably, the work in Margaris et al. [21] presents an algorithm that also confines its needs for SN-sourced data to trust relationships or elementary friendship among users. It computes SN-aware CF-rating predictions by synthesizing a SN-based prediction with a CF-based one. However, the algorithm presented in Margaris et al. [21] uses the same weight coefficients for the SN- and CF-based predictions in the synthesis step, an approach that does not take into account the particular properties of each user's SN-based as well as CF-based neighbourhoods.

This chapter advances the state-of-the-art, by introducing an algorithm that is able to adapt its behaviour to the features of the users' SN- and CF-based neighbourhoods. More specifically, the proposed algorithm analyses the users' already entered ratings and computes a personalized set of weight coefficients associated with SN- and CF-based predictions for each user. Through this approach, the proposed algorithm significantly leverages prediction accuracy. In this chapter, we also present our experiments and findings that quantify the prediction accuracy improvement and establish that the proposed approach consistently achieves improved accuracy under two correlation metrics and across five contemporary datasets that contain both SN relations and CF ratings.

3 SN CF Prediction Formulation Foundations

In CF, predictions for a user U are computed based on a set of users that have rated items similarly to U, namely U's Near Neighbours (NNs). For the majority of the CF systems, the similarity metric between two users U and V is typically based on either the Pearson Correlation Coefficient (PCC) or the Cosine Similarity (CS) metrics [14].

The PCC metric, denoted as $sim\_pcc(U, V)$, is expressed as:

$$sim\,pcc\,(U, V) = \frac{\sum_{\mathrm{k}} \left(\mathrm{r_{U,k}} - \overline{\mathrm{r_U}}\right) * \left(\mathrm{r_{V,k}} - \overline{\mathrm{r_V}}\right)}{\sqrt{\sum_{\mathrm{k}} \left(\mathrm{r_{U,k}} - \overline{\mathrm{r_u}}\right)^2 * \sum_{k} \left(\mathrm{r_{V,k}} - \overline{\mathrm{r_V}}\right)^2}} \tag{1}$$

where k ranges over items that have been rated by both U and V, while $\overline{r_U}$ and $\overline{r_V}$ are the mean values or ratings entered by users U and V.

Similarly, the Cosine Similarity (CS) metric, denoted as $sim\ \_\ cs(U, V)$, is expressed as:

$$simcs\,(U,V) = \frac{\sum_{\mathrm{k}} \mathrm{r_{U,k}} * \mathrm{r_{V,k}}}{\sqrt{\sum_{\mathrm{k}} \left(\mathrm{r_{U,k}}\right)^2} * \sqrt{\sum_{k} \left(\mathrm{r_{V,k}}\right)^2}} \tag{2}$$

Afterwards, for user U, his NN users, NN_U, are selected out of the ones whom a positive similarity has been computed with. Then, the rating prediction $p_{U,i}$ for the rating of user U on item i is computed. The computation is expressed as:

$$p_{U,i} = \overline{r_u} + \frac{\sum_{\mathrm{V} \in NN_u} sim_{CF}\,(U,\mathrm{V}) * \left(r_{\mathrm{V},i} - \overline{r_{\mathrm{V}}}\right)}{\sum_{\mathrm{V} \in NN_u} sim\,(U,\mathrm{V})} \tag{3}$$

where the $sim_{CF}(U, V)$ denotes the similarity metric that the particular CF system implementation has selected.

The work in Margaris et al. [21] introduced the concept of *SN NN*s of a user U: user V is considered to be U's SN NN, if a social relation, such as friendship or trust, has been established between them in the context of a SNS.

The set of SN NNs of user U will be denoted as SN_NN_U and is formally expressed as:

$$SNNN_U = \{V \in users(S) : r\,(U,V) \in S_r\} \tag{4}$$

where *users(S)* is the set of users within social network S, r is a social relationship within S and S_r is the extension of relationship r in the context of S. Similarly, we denote the initial CF NNs of a user U as CF_NN_U.

Moreover, the algorithm presented in Margaris et al. [21] follows a metasearch score combination algorithm [19] in order to combine the two partial prediction scores. One score is based on the SN-based near neighbourhood of the user (SN_NN_U), while the second is based on the traditional CF near neighbourhood of the user (CF_NN_U). The score from the SN-based near neighbourhood is denoted as $\mathrm{p}_{U,i}^{SN}$ and computed as:

$$\mathrm{p}_{U,i}^{SN} = \frac{\sum_{\mathrm{V} \in SNNN_{U,i}} sim_{SN}\,(U,V) * \left(r_{\mathrm{V},i} - \overline{r_{\mathrm{V}}}\right)}{\sum_{\mathrm{V} \in SNNN_{\mathrm{U}}} sim_{SN}\,(U,V)} \tag{5}$$

As far as the computation of the $sim_{SN}(U, V)$ quantity is concerned, which represents the SN-based user similarity, in this work we adopt the following approach:

- If the SN dataset provides values representing the strength/weight of the relationship between users U and V, $sim_{SN}(U, V)$ is set to this value.
- If the SN dataset does not provide such values, then $sim_{SN}(U, V)$ is fixed to 1.0, for all user pairs *(U, V)* for which a relationship is established within the SN.

Similarly, the CF near neighbourhood-based score is denoted as $\mathrm{p}_{U,i}^{CF}$, computed as:

$$p_{U,i}^{CF} = \frac{\sum_{V \in CF_{NNU,i}} sim_{CF}(U,V) * \left(r_{V,i} - \overline{r_V}\right)}{\sum_{V \in CFNN_U} sim_{CF}(U,V)} \tag{6}$$

$SN\_NN_{U,i}$ and $CF\_NN_{U,i}$ denote the SN- and CF-based NNs of user *U*, which have rated item *i*, respectively.

In previous research, certain features of SN structure and/or interaction among users have been shown to denote the strength of relationships between users. Such features include the number of common/mutual relations, tie strength, intimacy of message content and others [4, 20]. In our future work, we plan to investigate methods for exploiting these features, in order to compute or refine the $sim_{SN}(U, V)$ metric.

The partial predictions $p_{U,i}^{CF}$ and $p_{U,i}^{SN}$ are combined and the result is adjusted by the mean value of ratings entered by *U, U* ($\overline{r_U}$) U ($\overline{r_U}$), in order to formulate the rating prediction $p_{U,i}$, as shown in Eq. (7; [21]):

$$p_{U,i} = \begin{cases} \overline{r_U} + p_{U,i}^{CF}, & if\, SNNN_{U,i} = \varnothing \\ \overline{r_U} + p_{U,i}^{SN}, & if\, CFNN_{U,i} = \varnothing \\ \overline{r_U} + w_{CF} * p_{U,i}^{CF} + w_{SN} * p_{U,i}^{SN}, & if\, SNNN_{U,i} \neq \varnothing \wedge \\ & CFNN_{U,i} \neq \varnothing \end{cases} \tag{7}$$

In Eq. (7), the w_{CF} parameter corresponds to the weight assigned to the (partial) prediction computed by considering only the *CF_NNs*. The w_{SN} parameter, which is complementary to the w_{CF} parameter ($w_{SN} + w_{CF} = 1.0$), denotes the weight assigned to the prediction computed by considering only the *SN_NNs* of *U*, respectively. If no *CF_NNs* of *U*'s who have rated item *i* exist, then the prediction is based exclusively on the ratings of the user's *SN_NNs* and vice versa.

As shown in Eq. (7), the algorithm presented in Margaris et al. [21] uses the same w_{SN} and w_{CF} values (weights) to combine the partial predictions ($p_{U,i}^{CF}$ and $p_{U,i}^{SN}$) for all users within each dataset. However, such strategy may be suboptimal, since the properties of the CF- and SN-based neighbourhoods of each user *U* may vary significantly, necessitating the use of personalized weight assignments. For instance, a user U_1 may have a SN-based neighbourhood of high cardinality and a CF-based one of low cardinality, indicating that the w_{SN} for this particular user should be assigned a higher value than the respective value of w_{CF} for the same user.

In the following section, an algorithm that tackles the aforementioned problem is proposed. The proposed algorithm is able to adapt its behaviour by taking into account the density and utility of each user's SN and CF neighbourhoods. Furthermore, seven alternatives for combining the CF and SN partial predictions, calculated by Eqs. (5) and (6), are presented, which target to effectively replace the combination formula (Eq. 7) that the algorithm presented in Margaris et al. [21, 22] uses in order to produce the combined rating prediction value.

4 The Proposed Algorithm and the Partial Prediction Combination Alternatives

This section describes the proposed algorithm, as well as the seven alternatives for combining the CF and SN partial predictions. Finally, the time and space complexity of the proposed algorithm are assessed.

4.1 The Proposed Algorithm

The algorithm proposed in this chapter modifies Eq. (7) that was presented in the previous section, by catering for the use of personalized weights for the two partial predictions, $p_{U,i}^{SN}$ and $p_{U,i}^{CF}$, for the combination step. More specifically, the third case of Eq. (7), which corresponds to the condition (SN_{NNU}, i $\neq \varnothing \wedge CF_{NNU}$, i $\neq \varnothing$), is modified as shown in Eq. (8):

$$\overline{r_U} + w_U^{CF} * p_{U,i}^{CF} + w_U^{SN} * p_{U,i}^{SN} \tag{8}$$

where w_U^{SN} and w_U^{CF} denote the personalized weights for the SN- and the CF-based predictions, respectively.

Listings 1–3 present the aforementioned algorithm in detail; more specifically, Listings 1 and 2 present the computation of the CF- and SN-based partial pre-

```
FUNCTION compute_CF_Prediction(User X, Item it)
/* Implementation of equation (6), used for the CF prediction computation based on X's
    CF_NNs.
     INPUT: X is the user for whom the partial prediction will be computed; it is the respective
         item.
     OUTPUT: The CF partial rating prediction, or NULL if no CF_NN of X has rated it, and
         hence a prediction cannot be computed. */
       predictionNumerator = 0.0
       predictionDenominator = 0.0
       FOREACH Y∈ CF_NN_X
              IF (r_Y,it ≠ NULL) THEN
                     predictionNumerator += sim_CF(X,Y) * (r_Y,i − r̄_Y)
                     predictionDenominator += sim_CF(X,Y)
              ENDIF
       END /* FOREACH */
       IF (predictionDenominator = 0) THEN /* No CF_NN of X has rated it. */
              RETURN NULL
       ELSE /* At least one CF_NN of X has rated it. */
              return r̄_X + (predictionNumerator / predictionDenominator)
       END
END /* FUNCTION */
```

Listing 1 Computation of the CF-based partial rating prediction

```
FUNCTION compute_SN_Prediction(User X, Item it)
/*  Implementation of equation (5), used for the SN prediction computation based on X's
SN_NNs.
    INPUT: X is the user for whom the partial prediction will be computed; it is the respective
item.
    OUTPUT: The SN partial rating prediction, or NULL if no SN_NN of X has rated it, and
hence a prediction cannot be computed. */
      predictionNumerator = 0.0
      predictionDenominator = 0.0
      FOREACH Y ∈ SN_NN_X
            IF (r_Y,it ≠ NULL)THEN
                  predictionNumerator += sim_SN(X,Y)* (r_Y,i − r̄_Y)
                  predictionDenominator += sim_SN(X,Y)
            ENDIF
      END /* FOREACH */
      IF (predictionDenominator = 0) THEN /* No SN_NN of X has rated it. */
            RETURN NULL
      ELSE /* At least one SN_NN of X has rated it. */
            return r̄_X + (predictionNumerator / predictionDenominator)
      END
END /* FUNCTION */
```

Listing 2 Computation of the SN-based partial rating prediction

```
FUNCTION compute_Prediction (User U, Item i)
/* INPUT: U is the user for whom the prediction will be computed; i is the respective item.
   OUTPUT: The prediction computed or NULL if no NN (CF or SN) of U has rated i. */
      p^CF_U,i = compute_CF_Prediction (U, i)
      p^SN_U,i = compute_SN_Prediction (U, i)

/* If no NN (CF or SN) of U has rated i, return NULL. */
      IF (p^CF_U,i = NULL && p^SN_U,i = NULL) THEN
            RETURN (NULL)
      ELSE IF (p^CF_U,i = NULL)
            /* If no CF NN exists, the prediction will be based only on the SN_NNs. */
            RETURN (p^SN_U,i)
      ELSE IF (p^SN_U,i = NULL)
            /* If no SN NN exists, the prediction will be based only on the CF_NNs. */
            RETURN (p^CF_U,i)
      ELSE /* Both CF and SN NNs of U have rated it. */
            RETURN (r̄_U + w^CF_U * p^CF_U,i + w^SN_U * p^SN_U,i)
      ENDIF
END /* FUNCTION */
```

Listing 3 Synthesizing the CF- and SN-based partial predictions into a comprehensive rating prediction dictions, respectively, while Listing 3 presents the synthesis of the two partial predictions into a comprehensive prediction.

In the following subsection, the alternatives for combining the CF and SN partial predictions are presented and analysed, while Sect. 5 presents the experimentally deduced optimal combination and the evaluation.

4.2 Alternatives for Combining the CF and SN Partial Predictions

Regarding the computation of the personalized weights w_U^{SN} and w_U^{CF}, in this chapter we test the seven following alternatives:

1. The prediction is based only on the part (CF or SN) where each user has the largest number of NNs; In case of a tie, the weight is equally split between the two neighbourhoods. According to the above, the weights w_U^{SN} and w_U^{CF} are formulated as follows:

$$\mathrm{w}_U^{CF} = \begin{cases} 1, & if\ \left|CF_{NN_{U,i}}\right| > \left|SN_{NN_{U,i}}\right| \\ 0, & if\ \left|CF_{NN_{U,i}}\right| < \left|SN_{NN_{U,i}}\right| \\ 0.5, & if\ \left|CF_{NN_{U,i}}\right| = \left|SN_{NN_{U,i}}\right| \end{cases}$$
$$\mathrm{w}_U^{SN} = 1 - \mathrm{w}_U^{CF} \tag{9}$$

and, effectively, the prediction for the rating that user U would assign to item i is computed as shown in Eq. (10):

$$p_{U,i} = \begin{cases} \overline{r_{\mathrm{U}}} + \mathrm{p}_{U,i}^{CF}, & if\ \left|CF_{NN_{U,i}}\right| > \left|SN_{NN_{U,i}}\right| \\ \overline{r_{\mathrm{U}}} + \mathrm{p}_{U,i}^{SN}, & if\ \left|CF_{NN_{U,i}}\right| < \left|SN_{NN_{U,i}}\right| \\ \overline{r_{\mathrm{U}}} + 0.5 * \mathrm{p}_{U,i}^{CF} + 0.5 * \mathrm{p}_{U,i}^{SN}, & if\ \left|CF_{NN_{U,i}}\right| = \left|SN_{NN_{U,i}}\right| \end{cases} \tag{10}$$

This alternative will be denoted as *max_NNs*.

2. The w_U^{CF} weight is computed as the relative number of the CF_NNs to the number of all the NNs taken into account for the prediction computation (CF_NNs and SN_NNs):

$$\mathrm{w}_U^{CF} = \begin{cases} 0, & if\ \left|CF_{NN_{U,i}}\right| = 0 \\ 1, & if\ \left|SN_{NN_{U,i}}\right| = 0 \\ \frac{CF_{NN_{U,i}}}{CF_{NN_{U,i}} + SN_{NN_{U,i}}}, & if\ \left|CF_{NN_{U,i}}\right| > 0 \wedge \left|SN_{NN_{U,i}}\right| > 0 \end{cases}$$
$$\mathrm{w}_U^{SN} = 1 - \mathrm{w}_U^{CF} \tag{11}$$

This alternative will be denoted as *w_NNs*.

3. The prediction is based only on the part (CF or SN) for which the user has the largest cumulative similarity weight produced by his NNs; in case of a tie, the

weight is equally split between the two neighbourhoods. According to the above, the weights w_U^{SN} and w_U^{CF} are formulated as follows:

$$w_U^{CF} = \begin{cases} 1, & if \sum_{V \in CFNN_{U,i}} sim_{CF}(U,V) > \sum_{V \in SNNN_{U,i}} sim_{SN}(U,V) \\ 0, & if \sum_{V \in CFNN_{U,i}} sim_{CF}(U,V) < \sum_{V \in SNNN_{U,i}} sim_{SN}(U,V) \\ 0.5, & if \sum_{V \in CFNN_{U,i}} sim_{CF}(U,V) = \sum_{V \in SNNN_{U,i}} sim_{SN}(U,V) \end{cases}$$
$$w_U^{SN} = 1 - w_U^{CF} \tag{12}$$

and, effectively, the prediction $p_{U,i}$ for the rating that user U would assign to item i is computed as shown in Eq. (13):

$$p_{U,i} = \begin{cases} \overline{r_U} + p_{U,i}^{CF}, & if \sum_{V \in CFNN_{U,i}} sim_{CF}(U,V) > \sum_{V \in SNNN_{U,i}} sim_{SN}(U,V) \\ \overline{r_U} + p_{U,i}^{SN}, & if \sum_{V \in CFNN_{U,i}} sim_{CF}(U,V) < \sum_{V \in SNNN_{U,i}} sim_{SN}(U,V) \\ \overline{r_U} + 0.5 * p_{U,i}^{CF} + 0.5 * p_{U,i}^{SN}, & if \sum_{V \in CFNN_{U,i}} sim_{CF}(U,V) = \sum_{V \in SNNN_{U,i}} sim_{SN}(U,V) \end{cases} \tag{13}$$

This alternative will be denoted as *max_sim*.

4. The w_U^{CF} weight is computed as the ratio of the sum of similarities of U to her CF-neighbourhood to the sum of (a) the similarities of U to her CF-neighbourhood and (b) the similarities of U to her SN-neighbourhood, that is:

$$w_U^{CF} = \begin{cases} 0, & if \left|CF_{NN_{U,i}}\right| = 0 \\ 1, & if \left|SN_{NN_{U,i}}\right| = 0 \\ \frac{\sum_{V \in CFNN_{U,i}} sim_{CF}(U,V)}{\sum_{V \in CFNN_{U,i}} sim_{CF}(U,V) + \sum_{V \in SNNN_{U,i}} sim_{SN}(U,V)}, & if \left|CF_{NN_{U,i}}\right| > 0 \wedge \left|SN_{NN_{U,i}}\right| > 0 \end{cases}$$
$$w_U^{SN} = 1 - w_U^{CF} \tag{14}$$

This alternative will be denoted as *prop_sim*.

5. The w_U^{CF} weight is computed as the ratio of the average similarity of U to her CF-neighbourhood to the sum of (a) the average similarity between U and the members of her CF-neighbourhood and (b) the average similarity between U and the members of her SN neighbourhood. According to the above, the weights w_U^{SN} and w_U^{CF} are computed as follows:

$$w_U^{CF} = \begin{cases} 0, & if \left|CF_{NN_{U,i}}\right| = 0 \\ 1, & if \left|SN_{NN_{U,i}}\right| = 0 \\ \frac{\frac{\sum_{V \in CFNN_{U,i}} sim_{CF}(U,V)}{\left|CF_{NN_{U,i}}\right|}}{\frac{\sum_{V \in CFNN_{U,i}} sim_{CF}(U,V)}{\left|CF_{NN_{U,i}}\right|} + \frac{\sum_{V \in SNNN_{U,i}} sim_{SN}(U,V)}{\left|SN_{NN_{U,i}}\right|}}, & if \left|CF_{NN_{U,i}}\right| > 0 \wedge \left|SN_{NN_{U,i}}\right| > 0 \end{cases}$$
$$w_U^{SN} = 1 - w_U^{CF} \tag{15}$$

This alternative will be denoted as *prop_avg sim*.

6. The prediction is based only on the part (CF or SN) where each user has the largest ratio of NNs considering the item *i* (the item for which the prediction is computed) to the number of his overall NNs (this amount will be denoted as $rel_{NN_{U,i}}$). In case of a tie, the weight is equally split between the two neighbourhoods. According to the above, the weights w_U^{SN} and w_U^{CF} are formulated as follows:

$$\begin{aligned} w_U^{CF} &= \begin{cases} 1, & if \left| CFrel_{NN_{U,i}} \right| > \left| SNrel_{NN_{U,i}} \right| \\ 0, & if \left| CFrel_{NN_{U,i}} \right| < \left| SNrel_{NN_{U,i}} \right| \\ 0.5, & if \left| CFrel_{NN_{U,i}} \right| = \left| SNrel_{NN_{U,i}} \right| \end{cases} \\ w_U^{SN} &= 1 - w_U^{CF} \end{aligned} \tag{16}$$

where $CFrel_{NN_{U,i}} = \frac{\left| CF_{NN_{U,i}} \right|}{\left| CF_{NN_U} \right|}$ and $SNrel_{NN_{U,i}} = \frac{\left| SN_{NN_{U,i}} \right|}{\left| SN_{NN_U} \right|}$.

This alternative will be denoted as *max_rel_NNs*. Notably, in this variant the weights are tailored not only to the user for whom the recommendation is formulated for but also to the specific item through the consideration of item-specific neighbourhoods.

7. The w_U^{CF} weight is computed as the ratio of $CFrel_{NN_{U,i}}$ to the sum of $CFrel_{NN_{U,i}}$ and $SNrel_{NN_{U,i}}$, that is:

$$w_{U,i}^{CF} = \begin{cases} 0, & if\ CF_{NN_{U,i}} = 0 \\ 1, & if\ SN_{NN_{U,i}} = 0 \\ \frac{CFrel_{NN_{U,i}}}{CFrel_{NN_{U,i}} + SNrel_{NN_{U,i}}}, & if\ CF_{NN_{U,i}} > 0 \wedge SN_{NN_{U,i}} > 0 \end{cases} \tag{17}$$

while the value of the $w_{U,i}^{SN}$ weight is supplementary to the above (1 - $w_{U,i}^{CF}$).

This alternative will be denoted as *w_rel_NNs*. This alternative, similarly to the previous one, tailors the weights to both the user for whom the prediction is generated for and the item for which the rating is predicted.

In the next section, we will assess the performance of the aforementioned alternatives, in terms of prediction accuracy.

4.3 Complexity Analysis

In this subsection, we assess the complexity of the algorithms presented in the previous paragraphs.

The procedure computing the CF-based partial rating prediction presented in Listing 1 iterates over the user's CF neighbourhood and therefore its complexity is $O(|CF\_NN_U|)$, where CF_NN_U is the collaborative filtering-based near neighbourhood of the user for whom the rating is computed. Wang et al. [28] conclude that the consideration of the 8 members of the user's CF neighbourhood having the highest similarity with the user suffices to compute accurately this partial prediction, since no notable effect on the rating prediction is observed when more than 8 members are considered. Therefore, we can consider an upper bound for the complexity of this step.

Similarly, the procedure computing the SN-based partial rating prediction presented in Listing 2 iterates over the user's SN neighbourhood and therefore its complexity is $O(|SN_U|)$, where $\mathrm{SN_U}$ is the social neighbourhood of the user for whom the rating is computed. The work in Margaris et al. [19] asserts that considering the 8–10 members of each user's social neighbourhood with the strongest influence on the user (as influence is quantified by tie strength [4]) suffices to compute this metric accurately, since considering more members of the social neighbourhood has a negligible effect on the recommendation formulation. Wang et al. [28] concur this finding, further limiting the number of social neighbours that need to be considered to 6. Therefore, we can consider an upper bound for the complexity of this step.

Finally, the partial score synthesis presented in Listing 3 does not involve any iterations, and hence its complexity is equal to *O(1)*.

Taking into account all the above, we conclude that the complexity of the proposed algorithm is

$$\mathrm{O}\left(|\mathrm{CFNN_U}| + |\mathrm{SN_U}| + 1\right) \tag{18}$$

With $|\mathrm{CF\_\ NN_U}|$ and $|\mathrm{SN_U}|$ being capped by values 8 and 10, respectively.

Regarding space complexity, the overhead introduced by the proposed algorithm is negligible, compared to a plain CF-SN algorithm, since the additional information required by the proposed algorithm is the social network-based similarity between each user U and each member of U's social neighbourhood. Since, according to the discussion presented above, it suffices to maintain only up to 10 members of the social neighbourhood, the space overhead introduced by the algorithm is also capped to up to 10 real numbers per user.

5 Experimental Evaluation

In this section, we report on the experiments that were designed for the quantification of the achieved rating prediction improvement, from the deployment of the proposed algorithm. The results are compared against the results from:

1. The SN RS algorithm presented in Margaris et al. [21], denoted as *same weights,* which utilizes the same w_U^{SN} and w_U^{CF} weights for all users in each dataset. The *same weights* dataset has been shown to achieve improvements ranging from 1.35% to 3.25% for the dataset listed in Table 1, notably however the optimal values for the w_U^{SN} and w_U^{CF} weights are dataset-specific (e.g. the optimal value

Table 1 Datasets summary

Dataset name	#Users	#Items	#Ratings	Avg. #Ratings /User	Density	#Social Relations	Avg. #Social Relations /User	Type of items	Type of relations
Ciao [12]	30 K	73 K	1.6 M	53.4	0.07%	40K	1.3	General	Trust
FilmTrust [12]	1.5 K	2.1 K	35 K	23.5	1.13%	1.8 K	1.2	Movies	Trust
Epinions [27]	49 K	134	665 K	13.5	0.01%	487 K	9.9	General	Trust
LibraryThings [6]	83 K	506 K	1.7 M	20.5	0.004%	130 K	1.6	Books	Trust
Dianping SocialRec 2015 [17]	148 K	11 K	2.1 M	14.5	0.13%	2.5 M	17	Restaurants	Friendship

for w_U^{SN} is 0.3 for the Ciao dataset, while for the Epinion dataset the respective value is 0.6), hence a training phase must be executed for each dataset.
2. The plain CF algorithm [14], which does not take into account the SN relations.

For all cases, the plain CF algorithm is used as a baseline. In order to quantify the rating prediction accuracy of the contending algorithms, we have used two well-established error metrics, namely the Mean Absolute Error (MAE) and the Root Mean Squared Error (RMSE) that amplifies the importance of large deviations.

To compute the MAE and the RMSE, we employed the standard 'hide one' technique [32], where each time one rating in the database was hidden. Then, based on the ratings of other non-hidden items, its numeric value was tried to be predicted. Furthermore, in our experiments both the PCC and the CS metrics were used.

For hosting the datasets and running the rating prediction algorithms, we used a laptop equipped with a quad core Intel N5000 @ 1.1GHz CPU, 8GB of RAM and one 256GB SSD with a transfer rate of 560 MBps.

In the experiments, we have used five datasets that exhibit the following properties:

1. They contain both user-item ratings, as well as SN user relations.
2. They vary with respect to the type of dataset item domain (music, books, movies, restaurants, etc.), CF-density and SN-density, and size.
3. They are widely used for benchmarking in SN CF research and they are up to date; published the last 10 years.

Table 1 summarizes the basic properties of the considered datasets.

In our first experiment, random ratings from each user are hidden (5 rating selections per user) and then their values are predicted. To further validate our results, we conduct an additional experiment in every dataset containing the timestamps of the ratings, where the last rating from each user in the database is hidden and then its value is predicted. The results of these two experiments were in close agreement (less than 2.5% difference in results) and, therefore, we report only on the results of the first experiment, for conciseness.

In the remainder of this section, we present and discuss the results obtained from applying the algorithm presented in the previous section to the five datasets, using the two aforementioned errors metrics, as well as the two similarity metrics (PCC and CS). From the presentation of the results, we have excluded the variant *prop_avg sim*, since it was found to yield lower rating prediction accuracy in comparison to the baseline algorithm.

5.1 Prediction Accuracy Experiments Using the PCC as the Similarity Metric

Figure 1 illustrates the performance regarding the MAE reduction when the PCC similarity metric is used to quantify the similarity between two users. We can

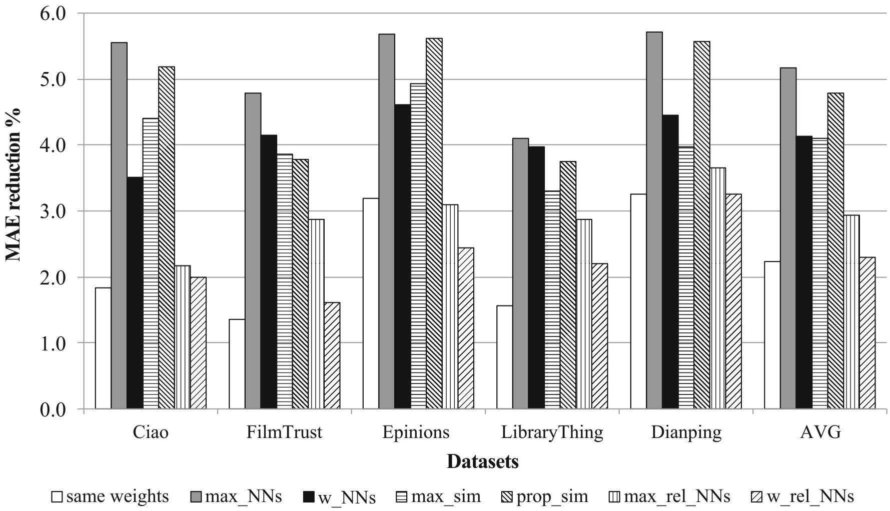

Fig. 1 MAE reduction for all datasets, using the PCC as the similarity metric

observe that the proposed algorithm, using the first alternative (*max_NNs)*, is the one achieving the best results for all five datasets tested. It achieves an average MAE reduction over all datasets equal to 5.2%, surpassing by approximately 2.3 times the corresponding improvement achieved by the *same weights* algorithm (2.2%) presented in [21], which uses the same weights for the CF NNs and SNNNs of all users. At the individual dataset level, the performance edge of the proposed algorithm against the *same weights* algorithm ranges from over 75% for the 'Dianping SocialRec 2015' dataset to 260% higher for the 'FilmTrust' dataset.

It has to be mentioned that the lowest MAE improvement for the proposed algorithm is observed for the 'Filmtrust' and the 'LibraryThing' datasets, which have relatively low #Social Relations / #Users ratio among the five datasets (1.2 and 1.6, respectively). In contrast, the highest MAE improvements for the proposed algorithm are observed for the 'Epinions' and the 'Dianping SocialRec 2015' datasets, which have the highest #Social Relations / #Users ratio among the five datasets (9.9 and 17, respectively). This fact clearly demonstrates the power of the proposed algorithm to exploit the available SN information to improve the RS prediction accuracy.

Considering the other alternatives for computing the weights for the CF and SN neighbourhoods, we can observe that the *prop_sim* algorithm is the runner up, achieving an average improvement of 4.8% against the baseline algorithm, lagging this behind the *max_NNs* algorithm by 0.4%. Interestingly, the biggest gap between the performance of *max_NNs* and *prop_sim* is observed for the 'Filmtrust' and the 'LibraryThing' datasets, which have relatively low #Social Relations / #Users ratios, indicating that the *prop_sim* algorithm achieves good results in more dense social neighbourhoods, but its performance in sparse social neighbourhoods declines.

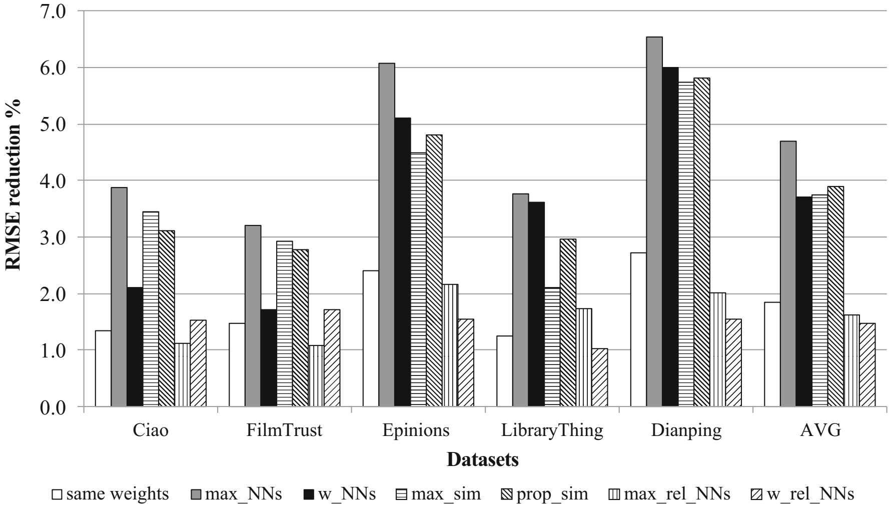

Fig. 2 RMSE reduction for all datasets, using the PCC as the similarity metric

Interestingly, the *max_rel_NNs* and *w_rel_NNs* alternatives, which are the two methods that consider item-specific neighbourhoods, tailoring the weights *both* to the user for whom the recommendation is formulated for *and* to the specific item for which the prediction is formulated, are found to achieve the lowest improvements among all variants discussed in this section. This indicates that for any individual user *U*, the effect of *U's* CF and SN neighbourhoods on the rating predictions formulated for *U* is generally uniform across all items, and the consideration of item-specific neighbourhoods merely adds noise to the rating prediction procedure. This issue will be investigated further in our future work.

Figure 2 demonstrates the performance regarding the RMSE reduction when similarity between users is measured using the PCC.

We can observe that the proposed algorithm, again using the first alternative (*max_NNs)*, achieves the best results for all five datasets tested, with an average RMSE reduction over all datasets equal to 4.7%, surpassing the improvement achieved by the *same weights* algorithm (1.8%), by approximately 2.6 times. At the individual dataset level, the performance edge of the proposed algorithm against the *same weights* algorithm ranges from 120% for the 'FilmTrust' dataset to 300% higher for the 'LibraryThing' dataset.

Furthermore, we can again clearly see that the 'Epinions' and the 'Dianping SocialRec 2015' datasets, having the highest #Social Relations / #Users ratio among the five datasets tested, achieve the highest RMSE reduction, while the other three datasets ('Ciao', 'Filmtrust' and 'LibraryThing'), which have the lowest #Social Relations / #Users ratio achieve the lowest MAE improvement. This fact, again, clearly demonstrates the power of the proposed algorithm to exploit the available SN information, in order to improve the RS prediction accuracy.

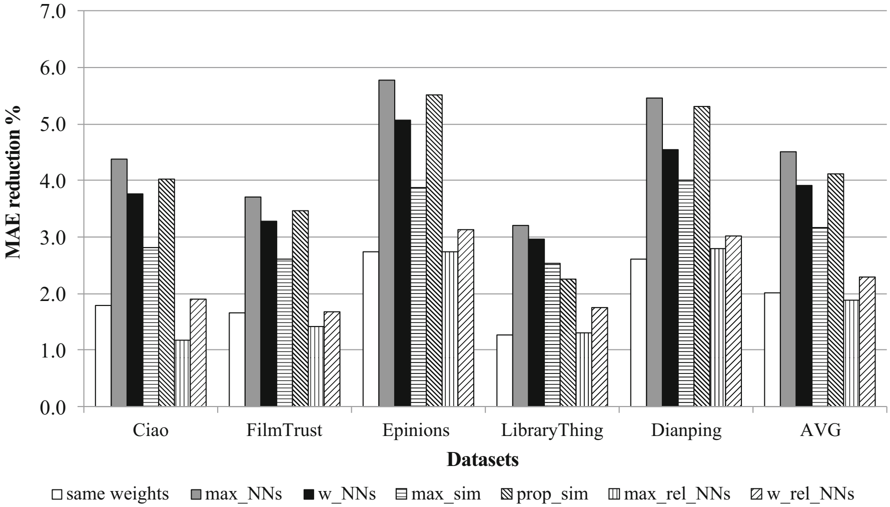

Fig. 3 MAE reduction for all datasets, using the CS as the similarity metric

Regarding the other alternatives for computing the weights for the CF and SN neighbourhoods, we can observe that the *prop_sim* algorithm is again ranked second, albeit with a wider margin from the *max_NNs* alternative than the one observed in the reduction of the MAE (0.8% against 0.4%). This indicates that the *max_NNs* alternative corrects more large errors than the *prop_sim* variant (the RMSE metric penalizes more severely large errors). Again, the two alternatives that consider item-specific neighbourhoods achieve the lowest improvements to the RMSE among all alternatives discussed in this section.

5.2 Prediction Accuracy Experiments Using the CS as the Similarity Metric

Figure 3 illustrates the performance regarding the MAE reduction when the CS metric is used to quantify the similarity between users. We can observe that the proposed algorithm, using the first alternative (*max_NNs)*, is again the one achieving the best results for all five datasets tested, with an average MAE reduction over all datasets equal to 4.5%, surpassing the improvement achieved by the *same weights* algorithm (2%), by approximately 2.2 times. At the individual dataset level, the performance edge of the proposed algorithm against the algorithm that sets the same weights for all users in each dataset ranges from 110% for the ‘Dianping SocialRec 2015’ dataset to 150% higher for the ‘LibraryThing’ dataset.

Yet again, that accuracy improvement achieved by the proposed algorithm is positively correlated to the density of available SN relations. Similarly to the case

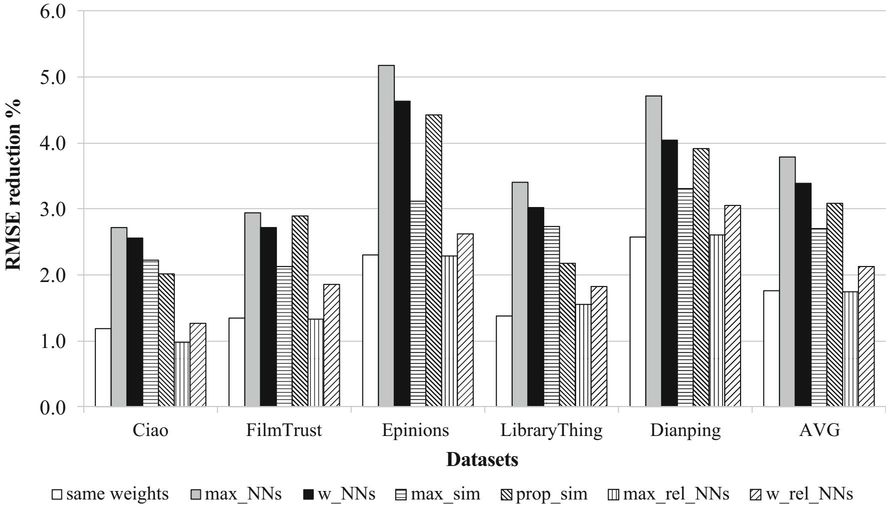

Fig. 4 RMSE reduction for all datasets, using the CS as the similarity metric

of using the PCC similarity metric, the *prop_sim* algorithm is the runner up: it achieves a MAE improvement of 4.1% against the baseline, lagging behind the performance of the *max_NNs* by 0.4%. Moreover, the two alternatives that consider item-specific neighbourhoods achieve the lowest improvements to the RMSE among all alternatives discussed in this section.

Finally, Fig. 4 illustrates the performance regarding the RMSE reduction when similarity between users is measured using the CS metric.

We can observe that the proposed algorithm, using the first alternative (*max_NNs)*, is the one achieving the best results for all five datasets tested, with an average RMSE reduction over all datasets equal to 3.8%, surpassing the improvement achieved by the *same weights* algorithm (1.8%) by approximately 2.1 times. At the individual dataset level, the performance edge of the proposed algorithm against the *same weights* algorithm ranges from 80% higher for the 'Dianping SocialRec 2015' dataset to 150% higher for the 'LibraryThing' dataset.

Furthermore, we can clearly see that the 'Epinions' and 'Dianping SocialRec 2015' datasets, having relatively high #Social Relations / #Users ratio among the five datasets tested, achieve the highest RMSE reduction, while the other three datasets ('Ciao', 'Filmtrust' and 'LibraryThing'), which have the lowest #Social Relations / #Users ratio, achieve the lowest MAE improvement. This clearly demonstrates the power of the proposed algorithm to exploit the available SN information, in order to improve the RS prediction accuracy.

Considering the other alternatives, we can observe that in this case the *w_NNs* variant is ranked second, while the *prop_sim* algorithm is ranked third. Again, the *prop_sim* algorithm appears to mostly remedy prediction errors of smaller magnitude, while the *max_NNs* and *w_NNs* algorithms manage to correct larger errors, and hence the latter two algorithms surpass *prop_sim* in this case. Once more, the two alternatives that consider item-specific neighbourhoods achieve the lowest improvements to the RMSE among all alternatives discussed in this section.

6 Conclusions and Future Work

Nowadays, where the available information on the Internet is chaotic, the task of recommending interesting products to the users is more difficult than ever. The core task RSs is the investigation of the preferences of online users and suggestion of items they might. CF, which is the most widely used RSs method, synthesizes the people's opinions to make accurate user rating predictions, which will lead to personalized recommendations, under the assumption that if users liked (bought, ate, listened, etc.) common items in the past, they are likely to do so in the future, as well.

However, traditional CF-based RSs assume that users are independent from each other and do not take into account the social interactions among them, such as trust and friendship. As a result, they fail to incorporate important aspects that denote influence and interaction among the users, which can enhance recommendation quality.

The aforementioned drawback has been recently overcome by SN-aware RSs, which take into account information derived from the user profiles and from the item profiles, as well as dynamic aspects and contextual information stemming from social information. With this information in hand, the SN-aware RSs achieve more targeted and hence more successful recommendations.

However, the success of a SN-aware RS greatly depends on the combination of the SN- and the CF-based information, in the sense of identifying which rating prediction information source (the SN relations or the CF NNs) is considered as the most reliable and useful predictor each time.

In this chapter, we proposed an algorithm that effectively combines SN information, specifically user social relations, with CF information, that is user ratings for items. The proposed algorithm formulates two partial prediction scores, from the SN and the CF neighbourhood, and then combines the two partial predictions using a weighted average metascore combination approach.

In contrast to the algorithm presented in Margaris et al. [21], which set the same weights to the two partial predictions for all the users within each dataset, the algorithm proposed in this chapter sets personalized weights for each individual user, based on the density and utility of each individual user's SN- and CF-based neighbourhoods.

In this direction, we have tested seven weight calculation alternatives. The one based only on the partial result, either SN or SF, where each user has the largest number of NNs for the item whose rating is about to be predicted, proved to be the optimal.

The proposed algorithm was validated through a set of experiments, aiming to quantify the improvement obtained in prediction accuracy, due to the consideration of the SN NNs in the recommendation process. In these experiments, five datasets containing both SN information (user–user relation) and CF information (user–item rating) were used. Measurements were taken under the two similarity metrics most widely used in RSs, namely the PCC and the CS. Additionally, two types of social relations, friendship (undirected) and trust (directed), were considered, in order to examine the behaviour of the proposed algorithm under several settings commonly encountered in SN RSs. The algorithm was proven to be highly adaptive to the characteristics of the datasets, yielding promising results in all cases.

The evaluation results have shown that the proposed algorithm may provide substantial improvement on rating prediction quality, across all datasets. The MAE decreases by 5.2% and the RMSE declines by 4.7%, on average, when the PCC metric is used (the respective reductions of the algorithm proposed in Margaris et al. [21] were 2.2% and 1.8%), and by 4.5% and 3.8%, respectively, when the CS metric is used (the respective reductions of the algorithm proposed in Margaris et al. [21] were 2% and 1.8%). In both cases, the performance of the plain CF algorithm is taken as a baseline. Furthermore, the proposed algorithm outperforms the algorithm proposed in Margaris et al. [21] for all cases, by an average of 2.3 times.

Since the proposed algorithm takes into account only each user's CF NN cardinality, as well as his SN NN cardinality, it does not introduce any extra overhead to the prediction calculation procedure, compared to the *same weights* algorithm; on the contrary, we can note that while the algorithm presented in Margaris et al. [21] always calculates both the CF- and the SN-based partial prediction, the algorithm proposed in this chapter can only calculate one partial prediction being thus more efficient (except for the case that the SN and CF neighbourhoods of the user have the same cardinality, where both partial scores need to be computed, thus the prediction formulation cost is similar to that of the algorithm presented in Margaris et al. [21]).

Moreover, the fact that the proposed algorithm achieves the highest error improvement for the datasets that have the highest #Social Relations / #Users ratios among the five datasets tested, under both metrics (PCC and CS), clearly proves the capacity of proposed algorithm to successfully exploit the available SN information to improve the RS prediction accuracy.

The proposed algorithm requires the availability of typical CF information (i.e. a user-rating matrix) and elementary social relation information (bidirectional friendships or unidirectional trusts). However, due to the fact that no additional information, such as users' demographic information (age, gender, nationality, location, etc.), items' characteristics (price, category, reliability, etc.) or SN's contextual information (tie strength, influence, etc.) is required, it can be applied to any SN CF dataset, standalone or in combination with other algorithms that have been proposed for improving rating prediction accuracy and/or coverage, such as matrix factorization models [15] and concept drift techniques [9].

This study has two limitations. Additional SN information (such as users' influence, tie strength, common–mutual relations, demographic data, contextual information, etc.) is not taken into account for tuning the $sim(U, V)_{SN}$ parameter. Furthermore, the proposed algorithm does not take the age of each user rating into account, in the sense that aged user ratings may not accurately reflect the current state of users regarding their likings and tastes, which may produce inaccurate predictions, due to the concept drift phenomenon [9].

Our future work will focus on investigating the computation-tuning of the $sim(U, V)_{SN}$ parameter value, considering additional information derived from the *SNs* domain. Furthermore, we are planning to evaluate the presented algorithm under additional user similarity metrics, such as the Euclidean Distance, the Hamming Distance, and the Spearman Coefficient [14], for the cases which those metrics are proposed by the literature as more suitable for the additional information. The above can also be utilized in broader applications of prediction methods that utilize social media data, such as textual reviews [22] or user-contributed data for the creation of detailed user profiles [1]. Finally, the combination of the proposed method with concept drift techniques [9] will also be investigated.

References

1. Aivazoglou, M., Roussos, A., Margaris, D., Vassilakis, C., Ioannidis, S., Polakis, J., & Spiliotopoulos, D. (2020). A fine-grained social network recommender system. *Social Network Analysis and Mining, 10*(20), 1–18.
2. Alahmadi, D., & Zeng, X. (2016). ISTS: Implicit social trust and sentiment based approach to recommender systems. *Expert Systems with Applications, 42*(22), 8840–8849.
3. Amato, F., Moscato, V., Picariello, A., & Piccialli, F. (2019). SOS: A multimedia recommender System for Online Social networks. *Future Generation Computer Systems, 93*, 914–923.
4. Bakshy, E., Eckles, D., Yan, R., & Rossen, I. (2012a). Social influence in social advertising: Evidence from field experiments. In *Electronic Commerce 2012 13th ACM Conference* (pp. 146–161). ACM. New York, NY, USA.
5. Bakshy, E., Rosenn, I., Marlow, C., & Adamic, L. (2012b). The role of social networks in information diffusion. In *World Wide Web (WWW) 2012 21st International Conference* (pp. 519–528). ACM. New York, NY, USA.
6. Cai, C., He, R., & McAuley, J. (2017). SPMC: Socially-aware personalized Markov chains for sparse sequential recommendation. In *Artificial Intelligence (IJCAI '17) 2017 International Joint Conference* (pp. 1476–1482).
7. Capdevila, J., Arias, M., & Arratia, A. (2016). GeoSRS: A hybrid social recommender system for geolocated data. *Information Systems, 57*, 111–128.
8. Chamoso, P., Rivas, A., Rodríguez, S., & Bajo, J. (2018). Relationship recommender system in a business and employment-oriented social network. *Information Sciences, 433–434*, 204–220.
9. Gama, J., Zliobaite, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys, 46*(4), 1–37.
10. Ge, M., Delgado-Battenfeld, C., & Jannach, D. (2010). Beyond accuracy: Evaluating recommender systems by coverage and serendipity. In *Recommender systems (RecSys) 2010 4th ACM International Conference on* (pp. 257–260). ACM. New York, NY, USA.
11. Gilbert, E., & Karahalios, K. (2009). Predicting tie strength with social media. In *Human Factors in Computing Systems 2009 SIGCHI Conference* (pp. 211–220). ACM. New York, NY, USA.

12. Guo, G., Zhang, J., Thalmann, D., & Yorke-Smith, N. (2014). ETAF: An extended trust antecedents framework for trust prediction. In *Advances in Social Networks Analysis and Mining (ASONAM) 2014 IEEE/ACM International Conference* (pp. 540–547). IEEE/ACM. Beijing, China.
13. He, J., & Chu, W. (2010). A social network-based recommender system (SNRS). *Annals of Information Systems, 12*, 47–74.
14. Herlocker, J. L., Konstan, J. A., Terveen, L. G., & Riedl, J. T. (2004). Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems, 22*(1), 5–53.
15. Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *IEEE Computer, 42*(8), 42–49.
16. Li, L., Zheng, L., Yang, F., & Lii, T. (2014). Modelling and broadening temporal user interest in personalized news recommendation. *Expert Systems with Applications, 41*(7), 3168–3177.
17. Li, H., Wu, D., Tang, W., & Mamoulis, N. (2015). Overlapping community regularization for rating prediction in social recommender systems. In *Recommender Systems (RecSys 2015) 9th ACM International Conference* (pp. 27–34). ACM. New York, NY, USA.
18. Ma, D., Dong, L., & Li, K. (2018). Collaborative filtering recommendation algorithm based on social relation and geographic information. In *Computer Science and Application Engineering (CSAE '18) 2nd International Conference* (pp. 1–7). ACM. New York, NY, USA.
19. Margaris, D., Vassilakis, C., & Georgiadis, P. (2016). Recommendation information diffusion in social networks considering user influence and semantics. *Social Network Analysis and Mining, 6*(108), 1–22.
20. Margaris, D., Vassilakis, C., & Georgiadis, P. (2018). Query personalization using social network information and collaborative filtering techniques. *Future Generation of Computer Systems, 78*(P1), 440–450.
21. Margaris, D., Spiliotopoulos, D., & Vassilakis, C. (2019a). Social relations versus near neighbours: Reliable recommenders in limited information social network collaborative filtering for online advertising. In *Advances in Social Networks Analysis and Mining (ASONAM 2019) 2019 IEEE/ACM International Conference* (pp. 1–8). IEEE/ACM. New York, NY, USA.
22. Margaris, D., Vassilakis, C., & Spiliotopoulos, D. (2019b). Handling uncertainty in social media textual information for improving venue recommendation formulation quality in social networks. *Social Network Analysis and Mining, 9*(64), 1–19.
23. Massa, P., & Avesani, P. (2004). Trust-Aware collaborative filtering for recommender systems. In *OTM 2004: On the Move to Meaningful Internet Systems 2004: CoopIS, DOA, and ODBASE* (pp. 492–508). Springer: Agia Napa, Cyprus.
24. Mukamakuza, C., Sacharidis, & Werthner, H. (2018). Mining user behaviour in social recommender systems. In *Web Intelligence, Mining and Semantics (WIMS '18) 8th International Conference* (pp. 1–6). ACM. New York, NY, USA.
25. Pham, T., Vuong, T., Thai, T., Tran, M., & Ha, Q. (2016). Sentiment analysis and user similarity for social recommender system: An experimental study. In *Science and Applications (ICISA), Lecture Notes in Electrical Engineering* (Vol. 376). Singapore: Springer.
26. Seo, Y., Kim, Y., Lee, E., & Baik, D. (2017). Personalized recommender system based on friendship strength in social network services. *Expert Systems with Applications, 69*, 135–148.
27. Simon, M., Guillot, E., Médini, L., & Laforest, F. (2019). *RED: A rich epinions dataset for recommender systems, archives-ouvertes.*https://hal.archives-ouvertes.fr/hal-01010246. Accessed 11 Nov 2019.
28. Wang, B., Huang, J., Ou, L., & Wang, R. (2015). A collaborative filtering algorithm fusing user-based, item-based and social networks. In *2015 IEEE International Conference on Big Data* (pp. 2337–2343). IEEE. Washington, DC, USA.
29. Xing, E. P., Jordan, M. I., & Russel, S. (2002). A generalized mean field algorithm for variational inference in exponential families. In *Proceedings of the nineteenth conference on uncertainty in artificial intelligence (UAI'03)* (pp. 583–591). Morgan Kaufmann Publishers: San Francisco, CA, USA

30. Yan, S., Lin, K., Zheng, X., Zhang, W., & Feng, X. (2017). An approach for building efficient and accurate social recommender systems using individual relationship networks. *IEEE Transactions on Knowledge and Data Engineering, 29*(10), 2086–2099.
31. Yu, J., Gao, M., Rong, W., Song, Y., & Xiong, Q. (2017). A social recommender based on factorization and distance metric learning. *IEEE Access, 5*, 21557–21566.
32. Yu, K., Schwaighofer, A., Tresp, V., Xu, X., & Kriegel, H. P. (2004). Probabilistic memory-based collaborative filtering. *IEEE Transactions on Knowledge and Data Engineering, 16*(1), 56–69.
33. Zhao, G., Qian, X., & Xie, X. (2016). User-service rating prediction by exploring social users' rating behaviours. *IEEE Transactions on Multimedia, 18*, 496–506.

Safe Travelling Period Recommendation to High Attack Risk European Destinations Based on Past Attack Information

Dimitris Spiliotopoulos, Dionisis Margaris, and Costas Vassilakis

Abstract Terrorism is a significant deterrent for tourism. It affects both visitors and local citizens and personnel of a country or area. On one hand, the potential visitor will probably avoid travelling to a high attack risk country, due to safety reasons, hence will miss the opportunity to visit it, and, on the other hand, the country's tourism will decline. This work addresses the aforementioned problem by (1) showing that relatively safe visiting periods for high attack risk European countries can be predicted with high accuracy, using limited information, comprising of attack and fatality data from the past years, which are widely available, and (2) developing an algorithm that recommends relatively safe periods to potential travellers.

The results of this work will be useful for tourists, visitors, businesses and operators, as well as relevant stakeholders and actors.

Keywords Terrorist attacks · Tourism · Safety perception · Risk calculation · Safety prediction · Recommendation algorithm · Evaluation

1 Introduction

Terrorist attack reports that are on the news are taken into account by travelling agencies and individual travellers that consider their options to travel. The term 'safe destination' is used to denote countries and cities where crime is not likely to happen against visitors. Crime, in general, and terrorism, in particular, are factors that organisations devoted to global peace, such as The Organisation for Economic Co-operation and Development (OECD) that publishes the Better Life Index [56],

D. Spiliotopoulos · C. Vassilakis (✉)
Department of Management Science and Technology, University of the Peloponnese, Tripoli, Greece
e-mail: dspiliot@uop.gr; costas@uop.gr

D. Margaris
Department of Digital Systems, University of the Peloponnese, Sparti, Greece
e-mail: margaris@uop.gr

M. Çakırtaş, M. K. Ozdemir (eds.), *Big Data and Social Media Analytics*, Lecture Notes in Social Networks, https://doi.org/10.1007/978-3-030-67044-3_4

and the Institute for Economics and Peace [19] that publishes the Global Terrorism Index (GTI; [15]), use to rank countries for safety.

The country safety indices take into account several geopolitical factors as well as expert opinions based on very recent political events or even terrorist attacks on areas or nearby locations. Terrorism is one of the most impactful factors that affect the citizen perception of safety. According to Pain [42], inciting widespread fear among the global population is a key objective behind terrorist attacks. According to the Institute for Economics and Peace, global terrorism peaked in 2014, with an unprecedented increase of 80% between 2013 and 2014. While terrorist attacks had then receded to the levels of 2013, the attacks are still very widely spread between countries. This is reflected by incidents of terrorism happening in relatively low-risk countries, maintaining terrorism as a prime public fear factor for many countries.

The influence of terrorism in the decisions of potential visitors, domestic or international, is undeniable [52]. Countries and prospective visitors, being dependent on perceived safety, may utilise such information to gradually build trust between countries and visitors for economic recuperation and tourism viability, allowing for stable, unhindered economic growth [51]. On the other hand, authorities may use this information by opting to examine the unsafer time periods and prepare accordingly to shield against and ultimately prevent terrorism in tourist destinations [24]. The news may also use such information, to inform tourists, citizens and businesses on safety [22].

Perceived destination safety is dependent on the recent past events. The frequency and severity of such events is measurable. In regard to tourism, the concerned parties are the travellers and tourists, as well as businesses, such as tour operators, hotel managers and others. Risk perception and risk estimate clearly differ among these groups; hence, an aggregated and uniform risk assessment is inherently of limited utility. Since terrorism is a targeted act, past data may prove useful for the derivation of patterns that create the terrorism attack footprint for a specific country [17]. Figure 1 depicts the frequency of recorded attacks and fatalities for Turkey in 2017. As seen, the terrorist activity is not uniform over time in a year. Moreover, there are periods of time that the terrorism activity is lower, shown in blue in Fig. 1. Therefore, for a traveller, it would be useful to find out the statistically low terrorism activity season as that would be the optimal time to travel.

The above consideration is useful for travellers that plan to visit a particular European country that has experienced recent high terrorism activity. Prospective visitors would still wish to visit; however, they would also like to feel as safe as possible.

Previous work [55] utilised past terrorism information, found in *Global Terrorism Database* (GTD; [27]), and more specifically, the number of tourism-related attacks, target types and attack types, in order to estimate the number of attacks a country may suffer in the following years, targeting at upgrading a country's safety measures and vice-versa. More specifically, the two main findings of our previous work are:

1. The tourism-related attack patterns mostly followed the general attack patterns, despite the tourist and non-tourist terrorist attack ratio for a country.

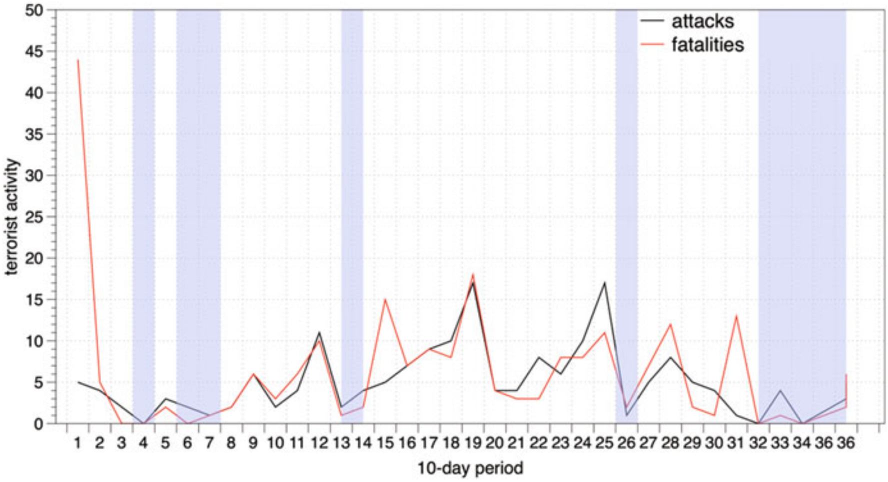

Fig. 1 Terrorist activity in Turkey in 2017 (source: GTD); the blue ribbon shows periods of lower terrorist activity

2. The number of attacks of the last three years proved to be a simple yet effective predictor for the following year's attack number, and irrelevant to the number of attacks recorded for a country.

Further analysis of these two findings is given in Sect. 3.

The present work extends the aforementioned work by proposing an algorithm that recommends optimal time slots to visit high terrorist activity European countries, providing valuable information to the query 'show me the safest time to visit country Y' a visitor will probably ask, hence elevating the state-of-the-art in this research area from safety evaluation to active safe period recommendation.

More specifically, the present work:

1. Studies how terrorist attacks disperse over the year-long periods per country
2. Proposes an algorithm that recommends the safest periods for travellers to visit a high terrorist activity European country (the previous work focused on the data analysis to understand how past information can be useful but did not recommend safe periods for travel)
3. Validates the proposed algorithm, by applying it to the European countries that have been targeted severely in the last years

The following hypotheses are examined in this work:

- H1—Is the algorithm applicable to high- and low-density data alike?
- H2—Is the algorithm accuracy retained for recommendations for earlier years?
- H3—Is the algorithm accuracy retained for non-yearly special events, such as the Olympics, which may shift tourism flow and terrorism focus in specific periods?

The experimental results show that the proposed algorithm introduces considerable prediction accuracy, achieving to predict low-risk attack visiting periods for the majority of the cases applied. It has to be mentioned that since the proposed algorithm requires only basic information (date and target type of the attack, fatalities, etc.), it is applicable to most cases.

It is also worth noting that the proposed algorithm can be enriched with additional input data that affect the terrorist attack events of a country, such as political, economic and social information.

The rest of the chapter is structured as follows: Sect. 2 overviews related work, while Sect. 3 describes the algorithm prerequisites that are used in our work. Section 4 presents the proposed recommendation algorithm. Section 5 evaluates the proposed algorithm and, finally, Sect. 6 concludes the chapter and outlines future work.

2 Related Work

In general, terrorism is negatively correlated to tourism in the literature [6, 47]. Terrorism inspires fear, which spans across citizens and visitors alike [22]. Earlier works studied the importance of safety on the attractiveness of tourist destinations, finding that safety is an important, marketable factor for tourist destinations [18]. That risk is calculated by potential visitors as a quality factor, which is as important as the natural beauty of the destination [16]. Tourism is a major source of income for tourist destination countries and terrorists know that targeting tourist destinations may force their respective governments to be involved in the international politics [2, 14, 46].

Targeting tourist destinations affects people that are directly involved in the tourist sector and businesses within the tourist industry, as well as a large portion of the economy partners that supports that industry indirectly, such as food suppliers; in total, it affects countries' economies heavily [25, 39]. Moreover, it has been observed that terrorist attacks in one tourist destination have an effect on tourism for other tourist destinations, based on correlations of places, geography, politics and other factors [5].

Studies have found that there is a difference on the level of impact of terrorism attacks between places with high and low tourist activity [7]. High tourist activity places are affected heavily in the short term but may recover over time based on factors such as media coverage, such as reports on life returning to normal, advertisements or tourist resilience to terrorism [9, 26], while low activity places are greatly affected in the long term, resulting in the tourist operations going out of business [28, 59]. The effects of terrorist attacks on an area result in tourism decline and may take a period of six months to a year for the local tourism industry to recover [44], based on how the perception of safety by potential visitors is restored [23].

Tourists, as opposed to local population, develop expectations and make decisions on whether to visit a destination based on past experiences and general perception of what they expect to experience [50]. Due to that fact, tourists are easy targets to be affected by terrorism [49]. Another business that is also involved in tourism and terrorism is the media, which reports on the impact internationally [8].

Since the recent increase in terrorist attacks, predicting future terror attacks is an important aim of the society and the individual governments. The public opinion steadily supports actions that aim to deter and shield the citizens from terrorism [45]. Understanding the dynamics of terrorism is a major step towards terrorist event prediction [21, 57]. Several works use machine learning, statistics and other big data analysis techniques to find patterns and predict terrorist events [11, 40, 41, 43, 58]. Kalaiarasi et al. [21] used the GTD and machine learning to predict terrorist threat, while Xia and Gu [60] utilised the GTD data to build a terrorism knowledge graph. Yang et al. [61] built a model to predict the lethality of terrorist attacks and tested it using the GTD data. All the aforementioned works try to predict terrorism, as in high terrorism activity or events and related parameters, such as fatalities. Recently, the work in Spiliotopoulos et al. [55] utilised limited past terrorism information, such as the number of tourism-related attacks, target types and attack types, to estimate the number of attacks a country may suffer in the following years, proving that terrorism attack estimation is possible.

None of the aforementioned works addresses the aspect of low terrorism activity prediction, and cannot be directly helpful to the prospective traveller, who plans to visit a country. The prospective visitor would benefit from knowing a predicted safe period, especially for countries that have had a significant number of terrorist attacks in the recent past.

This chapter advances the state-of-the-art on terrorism attack prediction, by analysing patterns of terrorism attacks, in yearly periods, in order to recommend relatively safe travelling periods to visit high attack risk European countries. Being able to provide specific time periods when the safety may prove to be much higher than the generic perception of safety (or non-safety), provided by international summits and politics, is a unique service to travellers and tourism businesses. The findings of this work support the fact that specific limited information from past terrorism data can be utilised to provide recommendations on timeslots that have a higher chance of being safe (higher predicted safety) for high terrorist attack risk destinations. Apart from the apparent immediate effect on visitor option for a safer travel time, such information may be used to filter time-series data such as social media data streams or surveillance data, may those be gathered through social networks [3, 34, 35, 48, 54] or other data sources, such as the IoT [33], to focus on specific points in time.

3 Algorithm Prerequisites

The analysis of the GTD by Spiliotopoulos et al. [55] was tourism centred. It utilised a subset of the GTD data that contained the tourist-related information. As an example, Fig. 2 depicts the attacks in France in the span of 18 years (2000–2017). A very high proportion of the attacks in France was tourist-related, that is, the attacks targeted locations where tourists would be, as opposed to targeting empty government facilities in remote locations. However, it is also shown that the tourism-related attack patterns mostly followed the general attack patterns. Other examples of this dispersion are Spain (geographically belonging to Western Europe), depicted in Fig. 3, and Greece (geographically belonging to Southeast Europe), depicted in Fig. 4.

France, Spain and Greece are bordering countries in Europe, but the terrorist attack pattern among them is quite opposite. One common factor, however, is that for all three countries, the patterns of the tourist and the overall attacks relatively match. For France and Spain, the number of tourist-related attacks was higher than the non-tourist-related ones. For Greece, the opposite was true. However, for all three countries the aforementioned attack matching pattern holds, albeit for Greece the non-tourist attack information is of greater value, due to the sheer difference in volume. Prediction algorithm worked well for the tourist data subset and since the subset pattern is very similar to the full data pattern, we opted to utilise the full data for the proposed algorithm, which is data driven.

The second finding of the work by Spiliotopoulos et al. [55] was the predictor data range. The results of that work showed that the data (tourist-related) of the last three years, and more specifically the accumulated attack number, are simple,

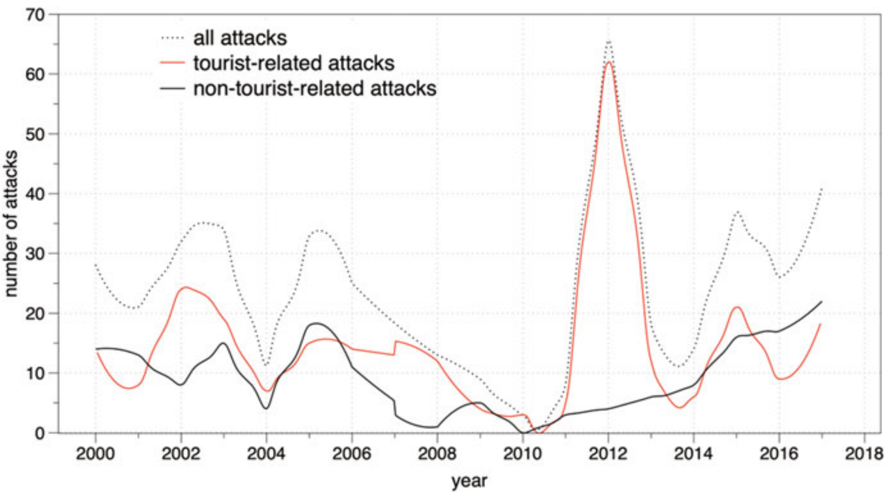

Fig. 2 Tourist- vs non-tourist-related attacks for France, 2000–2017

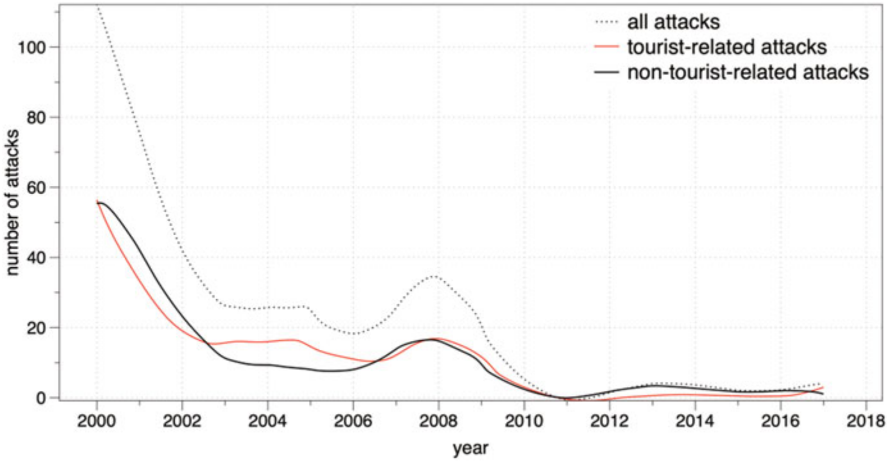

Fig. 3 Tourist- vs non-tourist-related attacks for Spain, 2000–2017

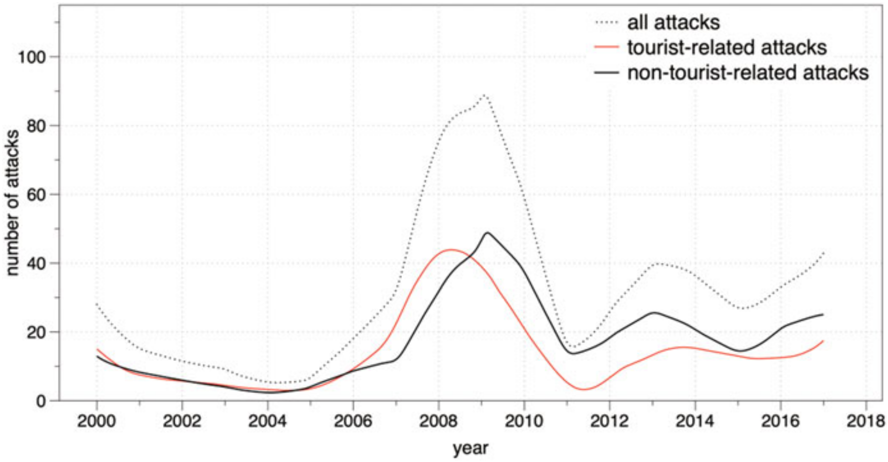

Fig. 4 Tourist- vs non-tourist-related attacks for Greece, 2000–2017

yet effective and reliable, predictors. That result was derived by using pruning techniques for enhancing prediction accuracy in recommender systems [29, 31, 32].

The proposed algorithm, therefore, is designed to use the last three years of data from the GTD. For the purpose of this work, no additional information or data source is considered. Although economic, political, social and intelligence information could result in higher accuracy or more reliable results, this work aims to develop a recommender utilising very limited information (number of attacks and time), which can be easily acquired from the GTD or similar widely available source. Therefore, the proposed algorithm is designed for standalone, broad use and is highly adaptive due to the limited information requirements. The proposed approach, however, can also be combined with other works that use multi-source data or enriched information.

4 Prediction Algorithm

The aim of this work is to design an algorithm to predict safe periods for travel to a European country, based on the limited information described in the previous sections. Tourism statistics show a variation in the number of days that visitors spend per country. Eurostat [13] reports that for the EU-28 the number of days span from 1.4 to 20.8, with an average of 6.1 days per visitor. Therefore, a ten-day window for safe timeslot prediction is a suitable baseline for our work.

Listing 1 presents the recommendation algorithm, which accepts as input the European country for which the recommendation will be generated and produces as output the recommended period of the year, computed as the period with the fewest incidents. The idea behind the proposed algorithm is to create a total of 36 ten-day slots (three per month, marking the start, middle and end of each month) and recommend the period corresponding to the slot having the smallest number of accumulated attacks over the last three years (following the work in [55]).

In the next section, we assess the performance of the aforementioned recommendation algorithm, in terms of predicting safe visiting periods for high terrorist attack risk countries.

5 Experimental Results

In this section, we report on the experiments that were designed to measure the accuracy of the proposed algorithm on recommending safe periods for visiting countries with high terrorism attack risk.

More specifically, from the GTD dataset, we selected the ten European countries that have had the most attacks over the last years, under the condition that a country had no more than two consecutive zero attacks per year, for the three years prior to the one that was predicted (since the experiments compute data from the three years prior to the target year), so that a prediction can be formulated based on the algorithm presented in Spiliotopoulos et al. [55]. Since (1) the GTD dataset contains attacks until 2017, and (2) based on the work in Spiliotopoulos et al. [55], taking into account the attacks of the last three years proved to be the optimal predictor for the volume of attacks that will happen in the next year, we store in a separate file the attacks that occurred between 2014 and 2016, for the aforementioned ten countries, targeting at recommending a relatively safe period for the year 2017, for each of the ten countries. The data for year 2017 are then used to assess the quality of the recommendation, that is, whether the algorithm achieved to recommend a period with ‘few’ incidents. This method is analogous to the ‘hide one’ technique,

```
FUNCTION populateIncidentHistogram(incidentSet): IncidentHistogram
/* Populates a histogram regarding the number of incidents that have occurred in each period of the year in a
   specific European country. Periods of years correspond to a duration equal to one third of month
   (BEGINNING, MIDDLE, END), thus totalling to 36 periods per year.
   Input: The set of incidents that have occurred in the country. Each has a "date" field, indicating when the
   incident occurred.
   Output: The histogram. The histogram elements are indexed by (month, monthPeriod) pairs, where month-
   Period∈ {BEGINNING, MIDDLE, END}. */
/* Initialize all histogram elements to 0 */
FOR month = 1 TO 12
   FOREACH monthPeriod {BEGINNING, MIDDLE, END}
      histogram[(month, monthPeriod)] = 0
   END /* FOREACH */
END /* FOR */
/* Select incidents that have occurred in the last three years. */
consideredIncidents = {i∈incidentSet: 1 ≤ CURRENT_YEAR - extractYear(i.date) ≤ 3}
FOREACH incident ∈consideredIncidents
   month = extractMonth(incident.date)
   day = extractDay(incident.date)
   /* Days 1-10 are mapped to BEGINNING; days 11-20 to MIDDLE; and all other days to END. */
   monthPeriod = mapDayToMonthPeriod(day)
   histogram[(month, monthPeriod)]++
END /* FOREACH */
RETURN histogram
END /* FUNCTION */

FUNCTION createAllHistograms(European_countries, incidentSet): IncidentHistogram[]
/* Populates a histogram array for all countries.
   Input: The set of countries, and the incident dataset. Each incident has a "date" field, indicating when the
   incident occurred and a "country" field, designating the country.
   Output: The array of histograms, indexed by the country. */
FOREACH country ∈ European_countries
   incidentsInCountry = {incident ∈incidentSet: incident.country = country}
   result[country] = populateIncidentHistogram(incidentsInCountry)
END /* FOREACH */
RETURN result
END /* FUNCTION */

FUNCTION generateRecommendation(country, histogramArray): PeriodOfYear
/* Generates a recommendation for the safest period of the year to visit a country.
   Input: The country for which the recommendation will be generated, and the array of histograms computed
   by function createAllHistograms.
   Output: Recommended period of the year, computed as the period with the fewest incidents. */
countryHistogram = histogramArray[country]
recommendation = NULL
/* Iterate over the elements of the histogram; period is the period index and value is the number of incidents in
the period. */
FOREACH (period, value) ∈countryHistogram
   IF ((recommendation == NULL) OR (value <countryHistogram[recommendation].value)) THEN
      recommendation = period
   END /* IF */
END /* FOREACH */
RETURN recommendation
END /* FUNCTION */
```

Listing 1 The proposed recommendation algorithm

Table 1 The ten European countries that suffered the highest number of terrorist attacks (2014–2016)

Country	Attacks	Fatalities
Turkey	1058	1535
United Kingdom	322	10
Russia	124	152
Germany	122	28
Greece	88	67
France	77	258
Ireland	76	2
Italy	23	1
Kosovo	14	2
Spain	8	0

commonly used in recommender systems for prediction evaluation [12, 20, 38]. The same process was repeated for safe period prediction for the year 2016, in this case by pruning the 2017 information and using the data from the three years prior to 2016 (2013–2015). Finally, the algorithm was also validated for 2015, using data from the three years prior (2012–2014) to test for H2.

Notably, the ten European countries examined in our experiment are depicted in Table 1, along with the number of attacks and the number of fatalities from those attacks between 2014 and 2016. These ten cases include EU countries with high tourism activity, such as Spain (no. 1), Italy (no. 2), the UK (no. 3) and France (no. 4) as the top European tourism destinations, according to Eurostat [13].

In order to provide a better service for the potential visitors, we adjust the algorithm to recommend a total of two periods (instead of just one), as an alternative option recommendation for the period of visit.

In this experiment, we use the following evaluation metrics for the periods recommended by our algorithm:

1. The number of the incidents (attacks or fatalities) that took place in the proposed period, when compared to the average attack number per period per country (#attacks in 2017 over 36 ten-day periods). This metric will be denoted as RNoI—Relative Number of Incidents. A recommendation is considered successful when its RNoI value is less than 100.0% (i.e. less than the average number of attacks per ten-day period).
2. The quartile that the recommended period belongs to, for each country. A recommendation is considered successful when a period belongs to either Q1 (very successful) or Q2 (successful).

In the remainder of this section, we present and discuss the results obtained from applying the algorithm presented above to the ten countries listed in Table 1, using the two aforementioned metrics, as well as two evaluation parameters, the number of attacks (Sect. 5.1) and the number of fatalities (Sect. 5.2). The reasons behind the evaluation of the proposed algorithm with two different parameters are (1) because both of these parameters are extremely important for the safety of tourists and (2)

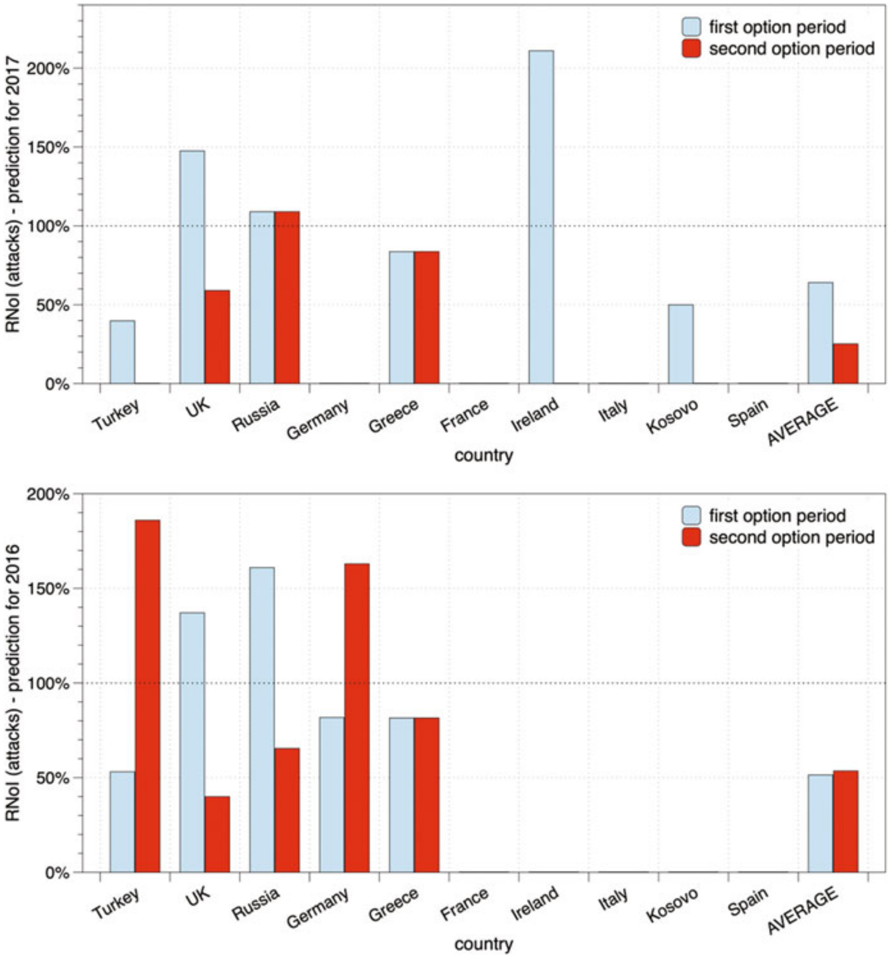

Fig. 5 RNoI results when taking the number of attacks as the evaluation parameter

to prove that the algorithm is parameter-independent, hence it can be applied in many-use cases.

5.1 Number of Attacks as the Evaluation Parameter

Figure 5 illustrates the results of the experiments, regarding the RNoI metric for the ten countries, when using the number of attacks as the evaluation parameter, for the 2017 prediction (top) and for the 2016 prediction (bottom). The average value for all the countries tested in our experiment is also included.

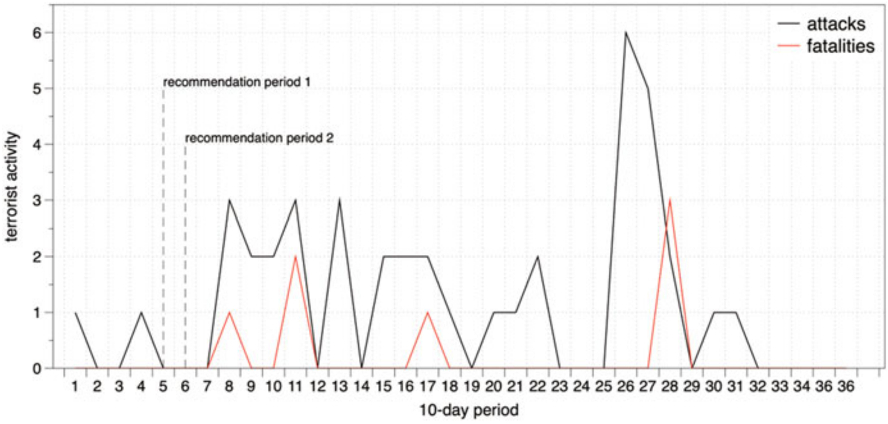

Fig. 6 Recommended periods for France (prediction year: 2017; predictor: *attacks*)

Regarding the predictions for year 2017, we can observe that the average RNoI value for the 10 countries is 45% (64% for the first recommended period and 25% for the second), while for 80% of the 20 recommendations produced (10 countries * 2 recommendations per country), the RNoI values were found to be less than 100% (i.e. less than the average number of attacks over the 36 ten-day periods in a year), indicating successful recommendations. Furthermore, we can clearly see that for the cases of Germany, France, Italy and Spain, the proposed algorithm achieved to recommend two periods with zero attacks for the four countries.

We repeated the same experiment for prediction of 2016 (after pruning the attacks of 2017 and using the attacks that occurred between 2013 and 2015 as predictor input) and the results are shown in the lower part of Fig. 5. More specifically, the average number of the RNoI value of the ten countries is 53% (51% for the first recommended period and 54% for the second). As before, for 80% of the 20 total recommendations produced, the RNoI values were found to be less than 100% (threshold), indicating successful recommendations. In this case, the algorithm achieved to recommend two periods with zero attacks for five countries (France, Ireland, Italy, Kosovo and Spain).

Similarly, for the 2015 predictions, the average was 47% for the primary and 50% for the secondary prediction. This shows that the algorithm worked successfully for past target years, using data from the three years prior to the target year.

More specifically, in the case of France, which was referred to in the introduction of this chapter, the proposed algorithm recommended, as safe visiting periods, the middle and the end of February, that is, the fifth and sixth ten-day period of the year (as indicated in Fig. 6).

Moreover, as far as Turkey is concerned, a country where 181 attacks occurred in 2017, the proposed algorithm successfully recommended periods with only 1 attack and 0 attacks, for the first and second recommendation, respectively. Considering that the average number of attacks in a 10-day period is 29.39 (i.e. 1058/36),

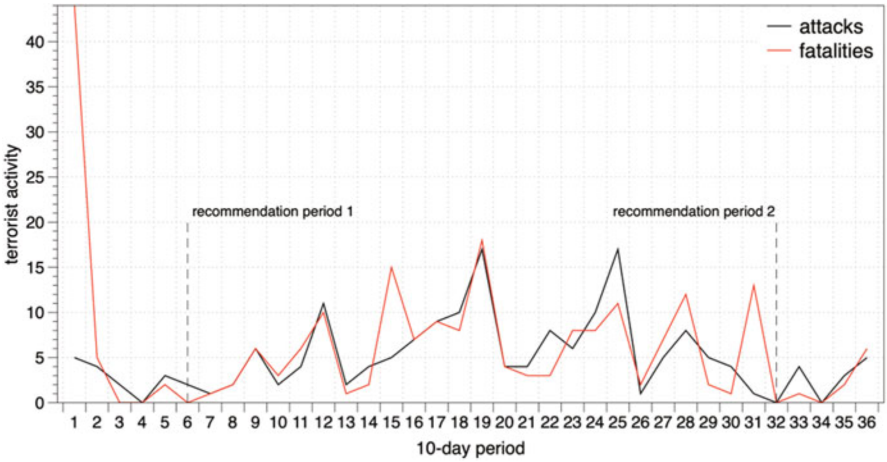

Fig. 7 Recommended periods for Turkey (prediction year: 2017; predictor: *attacks*)

these recommendations are deemed as very successful. Figure 7 illustrates the two periods proposed for visiting Turkey (end of February and middle of November, corresponding to the 6th and 32nd 10-day period of the year), along with a graph depicting the number of attacks and fatalities throughout the year.

Figure 8 illustrates the results of the evaluation under the second metric, that is, the quartile that each recommended period belongs to, for the ten countries. The average value for all countries considered in our experiment is also calculated and shown to the far right of the figure.

Figure 9 presents the aggregate results for all recommendations generated (20 in total = 10 countries * 2 recommendations per country). For each aggregated quartile value, the left column represents the algorithm prediction and the right column represents the aggregated values from ten randomly chosen ten-day periods averaged for each country.

Regarding the 2017 prediction (top of Fig. 9), we can observe that the vast majority of the recommendations (85%) are considered successful (Q1 and Q2), while only the 10% belongs to the Q4 (meaning that the recommendation is considered very unsuccessful, actually recommending a period with high attack risk to the visitors). The investigation and handling of this phenomenon will be part of our future work. At country level, we can see that only 3 out of the 20 cases that Fig. 9 presents are categorised in Q3 and Q4; however, 12 of them are categorised in Q1, while all the others in Q2. In the same figure, the results from the randomly chosen periods are significantly worse, showing safe periods (Q1 and Q2) for only 60% of the cases.

For the 2016 prediction (bottom of Fig. 9), 80% of the recommendations are successful (Q1 and Q2). As for 2017, the results from the randomly chosen periods are successful for 60% of the cases.

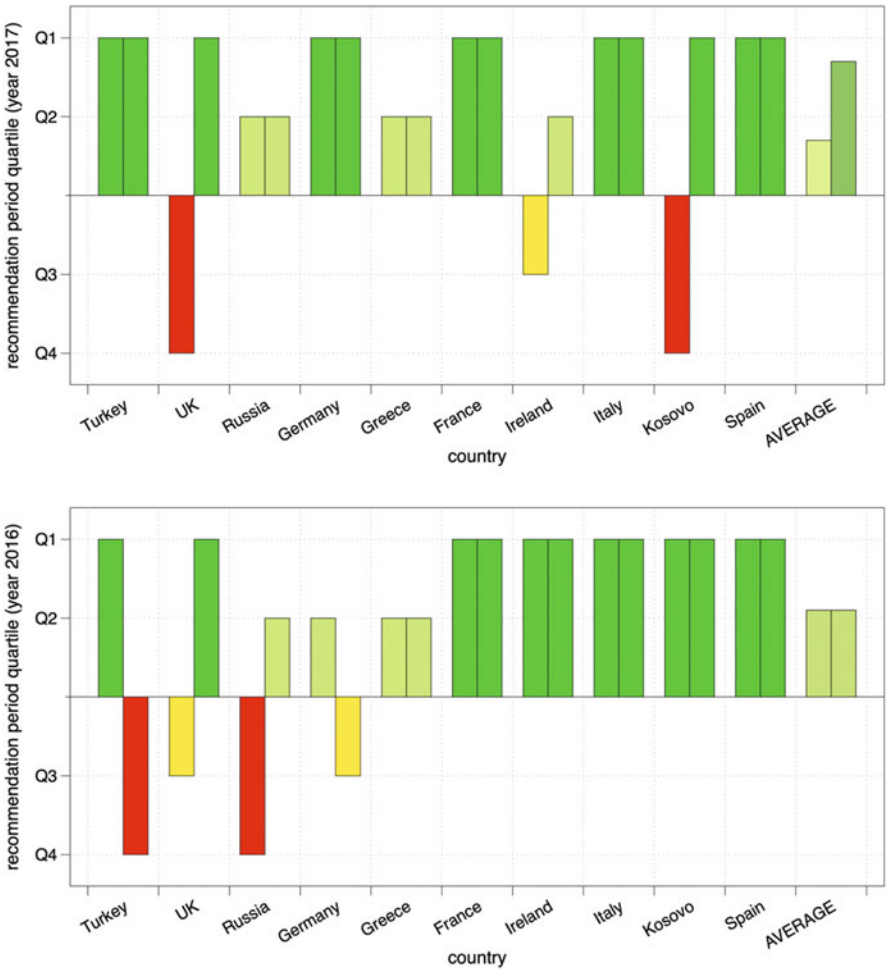

Fig. 8 Recommended period success for various countries when classifying the number of attacks into quartiles. For each country, the left value is the first recommended period and the right value is the second recommended period

Similarly, for the 2015 (two years back) recommendation experiments, Q1 was achieved for 60% of the cases, Q2 for 20%, Q3 for 15% and Q4 for 5%. The results from the randomly selected periods, Q1 and Q2, were achieved for only 60% of the cases.

To examine the cases when non-yearly periodic events (see H3), such as the Summer Olympics, occur, we evaluated the algorithm for 2012 UK and 2004 Greece. Both were Summer Olympics years for the respective countries. The quartiles for the predicted values were Q2 (primary) and Q1 (secondary) for UK and Q1 for both primary and secondary for Greece (Fig. 10). This shows that there

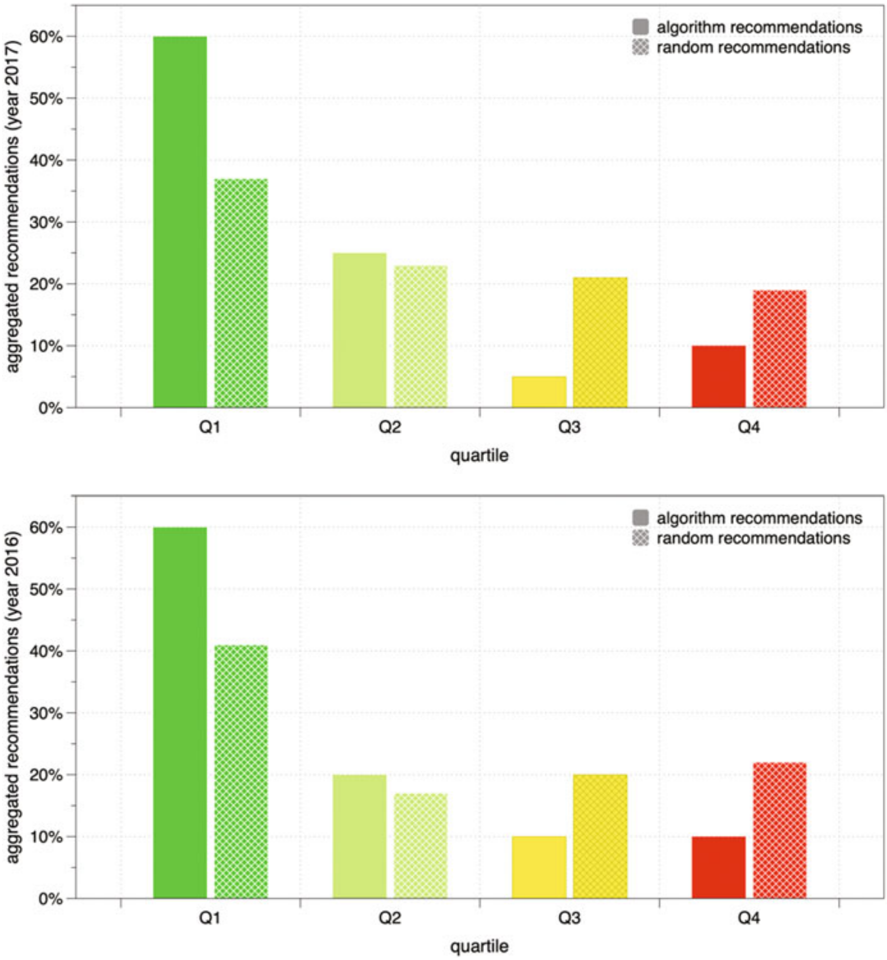

Fig. 9 Aggregated recommendations into quartiles for all countries (attacks)

are no exceptions for the algorithm applicability for special event years, since it resulted in recommendations with high accuracy.

5.2 *Number of Fatalities as the Evaluation Parameter*

This subsection analyses the results regarding the formulated recommendations' success when the number of fatalities is used as the evaluation parameter.

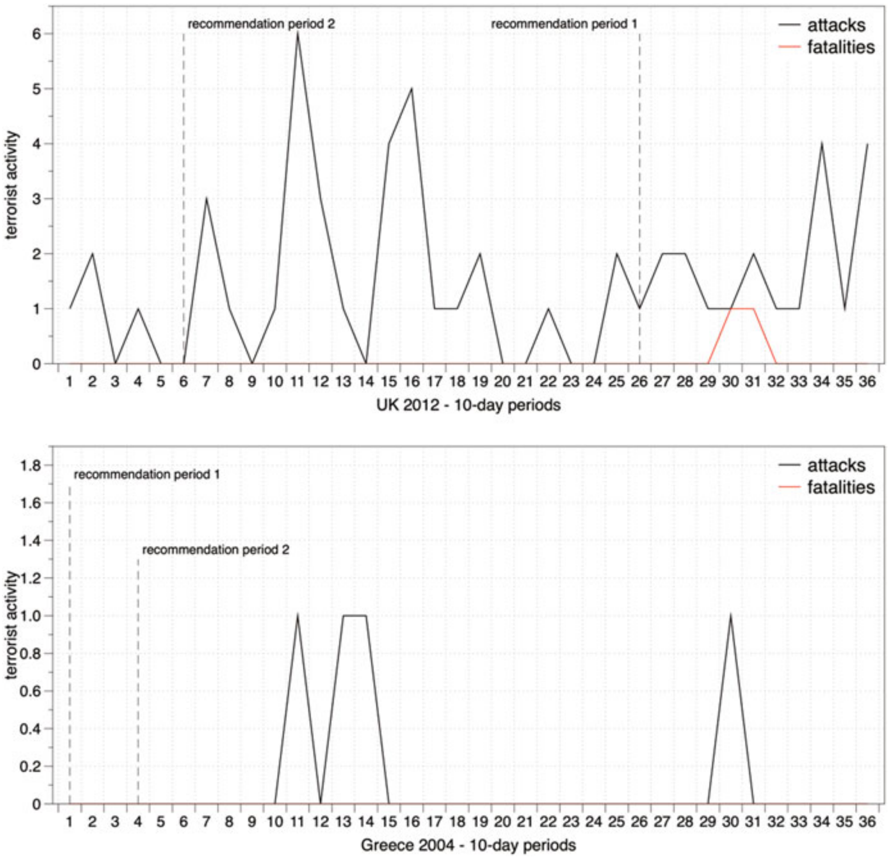

Fig. 10 Recommended periods for UK (prediction year: 2012; predictor: *attacks*) (top) and Greece (prediction year: 2004; predictor: *attacks*) (bottom)

Figure 11 illustrates the results of the experiments regarding the RNoI metric for the same ten countries. The average value of all the countries tested in our experiment is also depicted in Fig. 11.

Regarding the 2017 prediction (top of Fig. 11), we can observe that the average RNoI value of the ten countries is 17% (0% for the first recommended period and 34% for the second), while for the 90% of the recommendations generated, the relative number of fatalities is zero (indicating a successful recommendation). Similarly, for the 2016 prediction (bottom of Fig. 11), the average RNoI value of the ten countries is 16% (2% for the first recommended period and 29% for the second). For the 2015 prediction, the average of RNoI values was 22.5% for the primary prediction and 14% for the secondary. This clearly shows that the algorithm accuracy is retained for the earlier years.

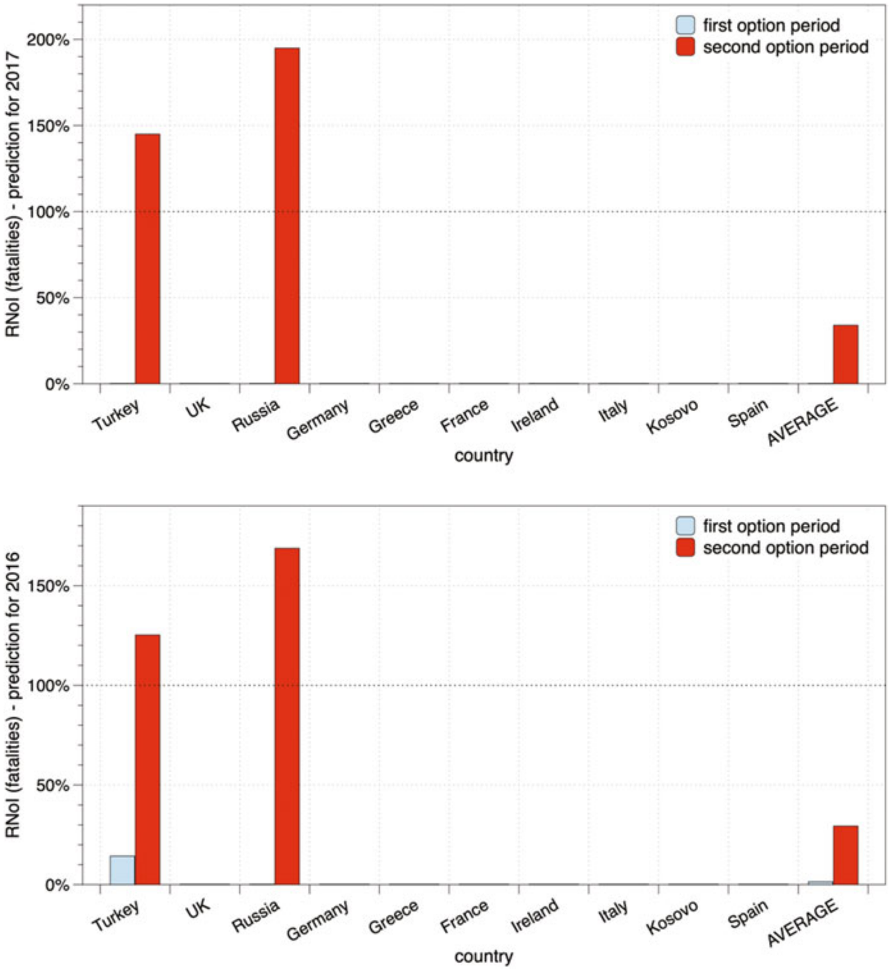

Fig. 11 RNoI results when taking the number of fatalities as the evaluation parameter

More specifically, in the case of France (Fig. 12), which was referred to in the introduction of this chapter, the proposed algorithm recommended, as safe visiting periods, the middle and the end of January, that is, the second and third ten-day period of the year (indicated by the dashed lines).

Figure 13 illustrates the results of the evaluation when considering the quartile that the recommended period belongs to, for the same ten cases. The average value of all the ten countries tested in our experiment is also included in the graph. Figure 14 presents the aggregate results for all the ten countries considered in our experiment for both 2017 and 2016 prediction.

For the 2017 prediction (top of Fig. 14), we can observe that the number of recommendations that are considered successful (Q1) is 90% (18 out of 20

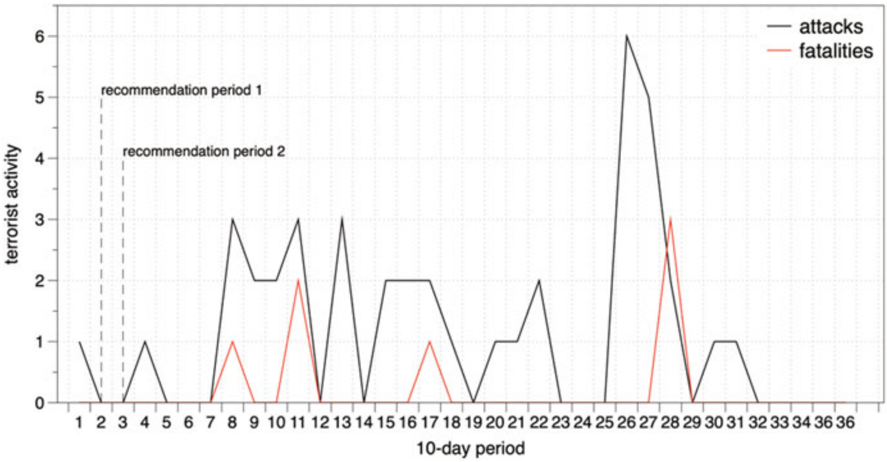

Fig. 12 Recommended periods for France (predictor: *fatalities*)

recommendations), while only the 10% of the predictions (2 out of 20) belongs to Q4 (meaning that the recommendation is considered very unsuccessful, actually recommending a period with high fatality numbers from attacks to the visitor). These cases will be part of our future work. On the same graph, the results from the randomly selected ten-day periods are worse, since less than 75% of the cases achieve a Q1 or Q2 recommendation. The reason for the relatively high accuracy for both the predicted and the randomly selected periods, when compared to the respective accuracy when the 'number of attacks' metric was used as the predictor (Sect. 5.1), is that several countries had registered minimal fatalities; therefore, the random or predicted choice had a higher chance to land on a zero-fatality period, increasing the number of Q1 predictions. In comparison, however, the algorithm prediction was much more accurate than the randomly selected period averages, proving algorithm effectiveness even in low data point situations.

For the 2016 prediction (bottom of Fig. 14), we can observe a similar dispersion. The number of very successful recommendations (Q1) is 90%, while the one from the randomly selected periods is less than 70%.

Similarly, for the 2015 (two years back) recommendation experiments, the percentage of the recommendations belonging to Q1 was 85%, to Q2 was 10% and to Q3 was 5%, proving the algorithm effectiveness when applied to past years.

6 Conclusion and Future Work

Potential visitors to destinations are discouraged by reports on terrorism or fear of terrorist attacks. Destination safety is quite hard to assess since past attacks may continue to be mentioned and still distil fear to local population, businesses and

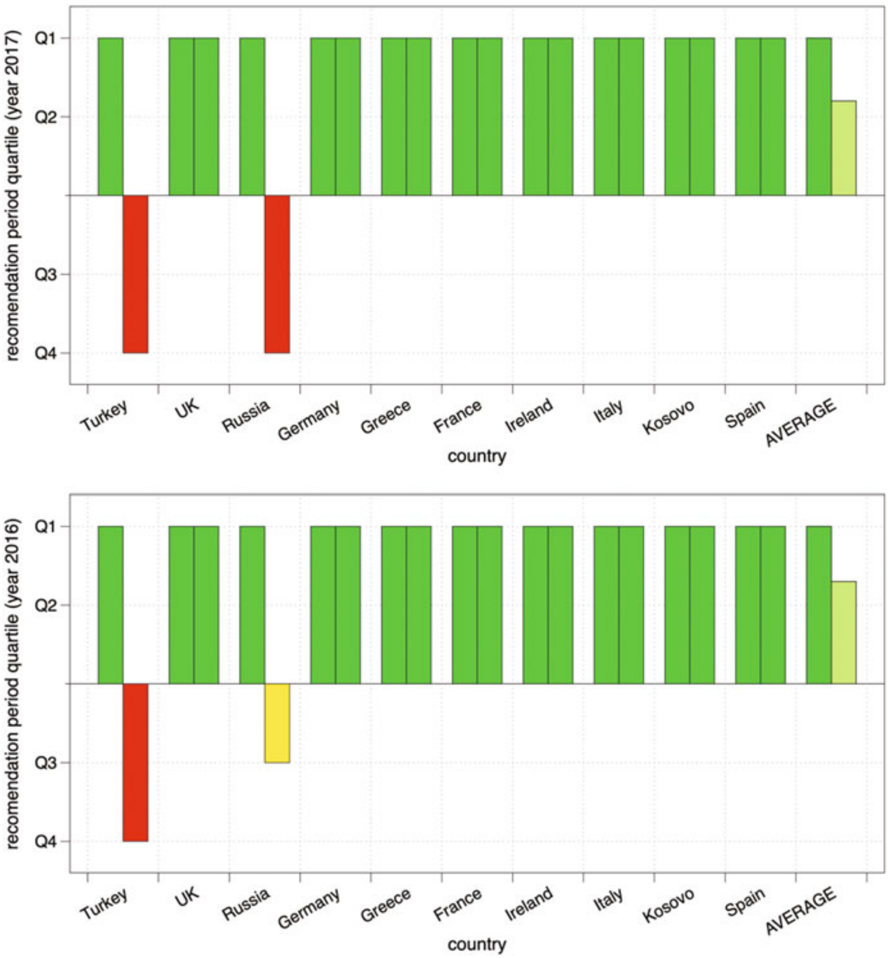

Fig. 13 Recommended period success for different countries when classifying into quartiles the number of fatalities. For each country, the left value is the first recommended period and the right value is the second recommended period

visitors alike, for several years after. A terrorist event results in a state of flux. On the one hand, tour operators and local businesses aim to restore normality and welcome new visitors. On the other hand, visitors, businesses and local governments are uncertain about safety and are looking for ways to assess the situation in the aftermath. The same is true for neighbouring locations, as well.

Perceived safety about a location or a country can be computed through user feedback on social media and is often reflected in the visitor count. Finding ways to measure and predict likeness of terrorist activity in specific time periods would enhance the perceived safety for those periods. Terrorist activity is not uniform over

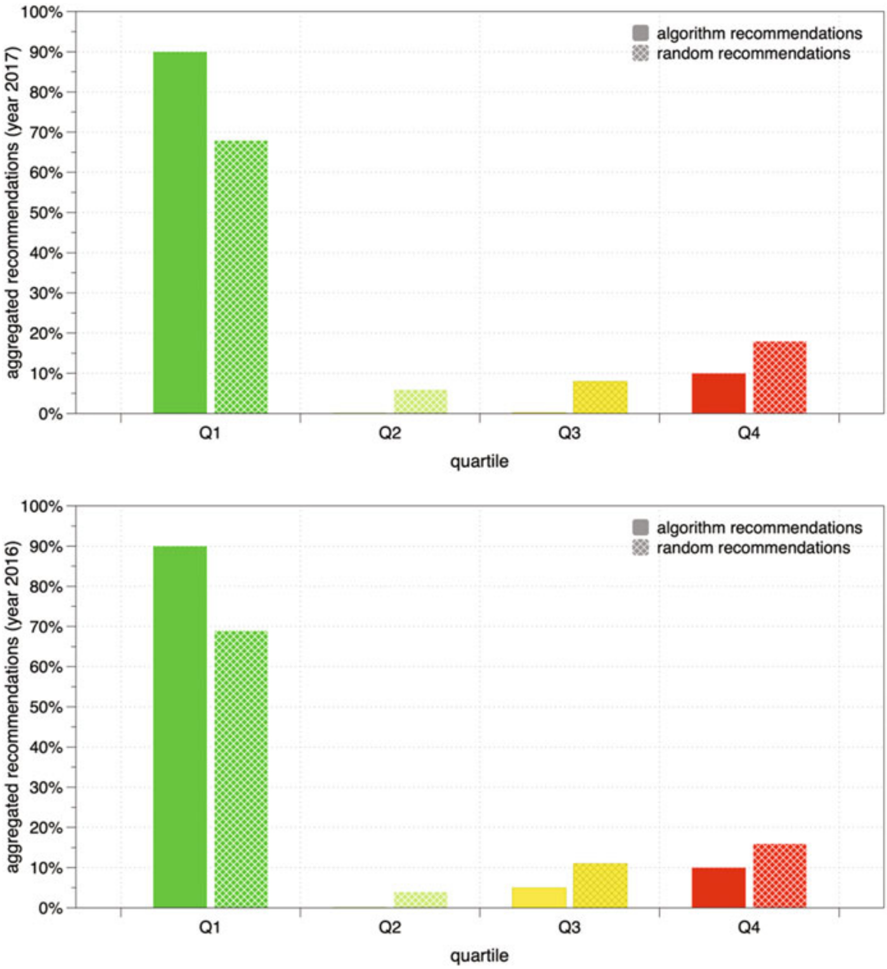

Fig. 14 Aggregated recommendation into quartiles for all countries (fatalities)

time in the span of year. Past data may be used to investigate patterns of activity and predict the safest periods to travel to locations that suffered recent terrorist attacks.

In this chapter, we have used the GTD data from past years and explored through the prism of *attacks* and *fatalities* as the means to understand and predict near-future terrorist activity for European countries. This activity can be approached through these two metrics for recommended periods of relative safety for travellers. Such prediction can help visitors and citizens get an up-to-date, real view of risk and safety, as well as help businesses and authorities plan on rebuilding safety and the perception of the people through resilience [10, 62].

Furthermore, an algorithm that recommends the relatively safe timeslots for travel over as ten-day periods, without season limitations, was proposed. The proposed algorithm is capable of working with limited information, in our case the number of attacks and fatalities of the past three years, which has been proposed in Spiliotopoulos et al. [55], as a simple yet effective predictor, along with their timestamps. Hence, it can be easily applied to many domains, such as travellers taking pilgrimage to a holy place or a football team travelling to an away game.

The presented algorithm was experimentally validated using ten European countries that had suffered from a large number of attacks in the 2014–2016 timeframe. The evaluation results showed that, for the majority of the cases, the ten-day period recommendations that were produced were considered successful. More specifically, the algorithm was tested using two evaluation parameters, the number of attacks and the number of fatalities, as well as two evaluation metrics, the relative number of incidents and the quartile each recommendation belongs to. In all cases, at least 80% of the recommendations were considered successful, by both metrics, at the same time. Furthermore, in 60% of the cases where the number of attacks was used as evaluation parameters and 90% of the cases where the number of fatalities were used, the recommended ten-day visiting period was considered very successful, since the number of attacks/fatalities was extremely low, for that country, or even zero (as in the case of France).

Through the validation, the algorithm was found to satisfy the original hypotheses set for this work. In regard to H1, the algorithm was found to accurately predict safe periods for travel using both number of attacks and number of fatalities. Especially for the latter, based on the experimental data, the algorithm was able to recommend Q1 safety periods even for countries that had suffered less than a total of 150 number of fatalities for the 2014–2016 period. In regard to H2, the algorithm was evaluated for years 2017, 2016 and 2015, and was found to exhibit the same high accuracy for all three predictions. In regard to H3, the algorithm was evaluated for year–country combinations that included Summer Olympics as a special event and was found to be equally accurate and applicable as it was for the other tested years.

One limitation of this work is on the algorithm applicability to very large countries. Such destinations may suffer terrorist attacks in specific regions, while other areas may be relatively safe. City-level safety prediction requires data on cities and countries to determine relative safety. Another limitation is that opportunistic terrorism targets that do not periodically occur cannot be predicted based on past information. However, identifying similarities with past opportunistic events and clustering them as special events may provide enough information (albeit lacking periodicity) for prediction shifts.

Our future work will focus on utilising social network streams, location features, transport and geographical data for real-time predictions [1, 4, 30, 36]. Geolocation, multimedia and text content [37, 53] can also be analysed for the sentiment of local citizens and visitors for modelling their perception of safety, worry or fear, towards a prospective visit to a country.

References

1. Aivazoglou, M., Roussos, A., Margaris, D., Vassilakis, C., Ioannidis, S., Polakis, J., & Spilitopoulos, D. (2020). A fine-grained social network recommender system. *Social Network Analysis and Mining, 10*(1), 8.
2. Albu, C. E. (2016). Tourism and terrorism: A worldwide perspective. *CES Working Papers, 8*(1), 1–19.
3. Antonakaki, D., Spiliotopoulos, D., Samaras, C. V., Ioannidis, S., & Fragopoulou, P. (2016). Investigating the complete corpus of referendum and elections tweets. In *Advances in social networks analysis and mining 2016 IEEE/ACM conference* (pp. 100–105). IEEE/ACM. San Francisco, CA, USA.
4. Antonakaki, D., Spiliotopoulos, D., Samaras, C. V., Pratikakis, P., Ioannidis, S., & Fragopoulou, P. (2017). Social media analysis during political turbulence. *PLoS ONE, 12*(10), 1–23.
5. Araña, J., & León, C. (2008). The impact of terrorism on tourism demand. *Annals of Tourism Research, 35*(2), 299–315.
6. Asongu, S. A., Nnanna, J., Biekpe, N., & Acha-Anyi, P. (2019a). Contemporary drivers of global tourism: Evidence from terrorism and peace factors. *Journal of Travel & Tourism Marketing, 36*(3), 345–357.
7. Asongu, S. A., Uduji, J. I., & Okolo-Obasi, E. N. (2019b). Tourism and insecurity in the world. *International Review of Economics, 66*(4), 453–472.
8. Burns, P., Lester, J., & Bibbings, L. (2010). Tourism and visual culture. *Methods and Cases., 2*, 1–234.
9. Cavlek, N. (2002). Tour operators and destination safety. *Annals of Tourism Research, 29*(2), 478–496.
10. Cox, A., Prager, F., & Rose, A. (2011). Transportation security and the role of resilience: A foundation for operational metrics. *Transport Policy, 18*(2), 307–317.
11. Ding, F., Ge, Q., Jiang, D., Fu, J., & Hao, M. (2017). Understanding the dynamics of terrorism events with multiple-discipline datasets and machine learning approach. *PLoS ONE, 12*(6), e0179057.
12. Ekstrand, M., Riedl, R., & Konstan, J. (2011). Collaborative filtering recommender systems. *Foundations and Trends in Human-Computer Interaction, 4*(2), 81–173.
13. Eurostat. (2017). *Tourism statistics*. https://ec.europa.eu/eurostat/statistics-explained/index.php/Tourism_statistics. Accessed 21 June 2020.
14. Fuchs, G., Uriel, N., Reichel, A., & Maoz, D. (2012). Vacationing in a terror-stricken destination. Tourists' risk perceptions and rationalizations. *Journal of Travel Research, 52*(2), 182–191.
15. Global Terrorism Index. (2019). http://globalterrorismindex.org/. Accessed 12 Dec 2019.
16. Goldman, O. S., & Neubauer-Shani, M. (2016). Does international tourism affect transnational terrorism? *Journal of Travel Research, 56*(4), 451–467.
17. Guo, W. (2019). Common statistical patterns in urban terrorism. *Royal Society Open Science, 6*(9), 2–13.
18. Gupta, A. (2011). Terrorism and its impact on financial performance: A case of tourism industry. *International Journal of Financial Management, 1*(4), 46–52.
19. Institute for Economics and Peace. (2019). http://economicsandpeace.org/. Accessed 12 Dec 2019.
20. Herlocker, J., Konstan, J., Terveen, L., & Riedl, J. (2004). Evaluating collaborative filtering recommender systems. *ACM Transactions in Information Systems, 22*(1), 5–53.
21. Kalaiarasi, S., Mehta, A., Bordia, D., & Sanskar, D. (2019). Using Global Terrorism Database (GTD) and machine learning algorithms to predict terrorism and threat. *International Journal of Engineering and Advanced Technology (IJEAT), 9*(1), 1–6.
22. Kapuściński, G., & Richards, B. (2016). News framing effects on destination risk perception. *Tourism Management, 57*, 234–244.

23. Khan, S. (2011). Gendered leisure: Are women more constrained in travel for leisure? *Tourismos: An International Multidisciplinary Journal of Tourism, 6*(1), 105–121.
24. Kılıçlar, A., Uşaklı, A., & Tayfun, A. (2018). Terrorism prevention in tourism destinations: Security forces vs. civil authority perspectives. *Journal of Destination Marketing and Management, 8*, 232–246.
25. Korstanje, M. E. (2013). Preemption and terrorism. When the future governs. *Cultura, 10*(1), 167–184.
26. Liu, A., & Pratt, S. (2017). Tourism's vulnerability and resilience to terrorism. *Tourism Management, 60*, 404–417.
27. LaFree, G., & Dugan, L. (2007). Introducing the global terrorism database. *Political Violence and Terrorism, 19*, 181–204.
28. Larsen, S., Brun, W., & Ogaard, T. (2009). What tourists worry about. Construction of a scale measuring tourist worries. *Tourism Management, 30*(2), 260–265.
29. Margaris, D., & Vassilakis, C. (2016). Pruning and aging for user histories in collaborative filtering. In *7th Symposium Series on Computational Intelligence (SSCI) 2016 International Symposium* (pp. 1–8). IEEE. Athens, Greece.
30. Margaris, D., Vassilakis, C., & Georgiadis, P. (2016). Recommendation information diffusion in social networks considering user influence and semantics. *Social Network Analysis and Mining, 6*(1), 108.
31. Margaris, D., & Vassilakis, C. (2017a). Enhancing user rating database consistency through pruning. In *Transactions on large-scale data- and knowledge-centered systems* (Vol. XXXIV, pp. 33–64). Berlin: Springer.
32. Margaris, D., & Vassilakis, C. (2017b). Improving collaborative filtering's rating prediction quality in dense datasets, by pruning old ratings. In *Symposium on Computers and Communications (ISCC) 2017 IEEE* (pp. 1168–1174). IEEE. Heraklion, Greece.
33. Margaris, D., & Vassilakis, C. (2017c). Exploiting internet of things information to enhance venues' recommendation accuracy. *Service Oriented Computing & Applications, 11*(4), 393–409.
34. Margaris, D., & Vassilakis, C. (2018). Exploiting rating abstention intervals for addressing concept drift in social network recommender systems. *Informatics, 5*(2), 21.
35. Margaris, D., Vassilakis, C., & Georgiadis, P. (2018). Query personalization using social network information and collaborative filtering techniques. *Future Generation of Computer Systems, 78*, 440–450.
36. Margaris, D., Spiliotopoulos, D., & Vassilakis, C. (2019a). Social relations versus near neighbours: Reliable recommenders in limited information social network collaborative filtering for online advertising. In *2019 Advances in Social Networks Analysis and Mining (ASONAM 2019) IEEE/ACM International Conference* (pp. 1160–1167). IEEE/ACM. New York, NY, USA.
37. Margaris, D., Vassilakis, C., & Spiliotopoulos, D. (2019b). Handling uncertainty in social media textual information for improving venue recommendation formulation quality in social networks. *Social Network Analysis and Mining, 9*(1), 64.
38. Margaris, D., Kobusinska, A., Spiliotopoulos, D., & Vassilakis, C. (2020). An adaptive social network-aware collaborative filtering algorithm for improved rating prediction accuracy. *IEEE Access, 8*(1), 68301–68310.
39. Meierrieks, D., & Gries, T. (2013). Causality between terrorism and economic growth. *Journal of Peace Research, 50*(1), 91–104.
40. Meng, X., Nie, L., & Song, J. (2019). Big data-based prediction of terrorist attacks. *Computer Electric Engineering, 77*, 120–127.
41. Mo, H., Meng, X., Li, J., & Zhao, S. (2017). Terrorist event prediction based on revealing data. In *Big data 2017 international conference* (pp. 239–244). IEEE. Beijing, China.
42. Pain, R. (2014). Everyday terrorism: Connecting domestic violence and global terrorism. *Progress in Human Geography, 38*(4), 531–550.
43. Palak, A., Mahak, S., & Satish, C. (2019). Comparison of machine learning approaches in the prediction of terrorist attacks. In *Contemporary Computing (IC3) 2019 Twelfth International Conference* (pp. 1–7). IEEE. Noida, India.

44. Pizam, A., & Fleischer, A. (2002). Severity versus frequency of acts of terrorism: Which has a larger impact on tourism demand? *Journal of Travel Research, 40*(3), 337–339.
45. Rai, T. S. (2019). Predicting future terror attacks. *Science, 366*(6467), 834–835.
46. Ranga, M., & Pradhan, P. (2014). Terrorism terrorizes tourism: Indian Tourism effacing myths? *International Journal of Safety and Security in Tourism, 1*(5), 26–39.
47. Samitas, A., Asteriou, D., Polyzos, S., & Kenourgios, D. (2018). Terrorist incidents and tourism demand: Evidence from Greece. *Tourism Management Perspectives, 25*, 23–28.
48. Schefbeck, G., Spiliotopoulos, D., & Risse, T. (2012). The recent challenge in web archiving: Archiving the social web. In *Archives Congress 2012 International Council* (pp. 20–24).
49. Seabra, C., Dolnicar, S., Abrantes, J., & Kastenholz, E. (2013). Heterogeneity in risk and safety perceptions of international tourists. *Tourism Management, 36*, 502–510.
50. Seddighi, H., & Theocharous, A. (2002). A model of tourism destination choice: a theoretical and empirical analysis. *Tourism Management, 23*(5), 475–487.
51. Sönmez, S. (1998). Tourism, terrorism, and political instability. *Annals of Tourism Research, 25*(2), 416–456.
52. Sönmez, S., & Graefe, A. (1998). Influence of terrorism risk on foreign tourism decisions. *Annals of Tourism Research, 25*(1), 112–144.
53. Spiliotopoulos, D., Margaris, D., Vassilakis, C., Petukhova, V., & Kotis, K. (2019a). A methodology for generated text annotation for high quality speech synthesis. In *Information, Intelligence, Systems and Applications (IEEE IISA 2019) 10th IEEE International Conference* (pp. 1–8). IEEE. Patras, Greece.
54. Spiliotopoulos, D., Tzoannos, E., Stavropoulou, P., Kouroupetroglou, G., & Pino, A. (2012). Designing user interfaces for social media driven digital preservation and information retrieval. In: Miesenberger K., Karshmer A., Penaz P., Zagler W. (eds). Computers Helping People with Special Needs. ICCHP 2012. Lecture Notes in Computer Science, vol 7382. Springer, Berlin, Heidelberg. https://doi.org/10.1007/978-3-642-31522-0_87.
55. Spiliotopoulos, D., Vassilakis, C., & Margaris, D. (2019b). Data-driven country safety monitoring terrorist attack prediction. In *2019 Advances in Social Networks Analysis and Mining (ASONAM 2019) IEEE/ACM International Conference* (pp. 1128–1135). IEEE/ACM. New York, NY, USA.
56. The Organisation for Economic Co-operation and Development (2019). http://www.oecdbetterlifeindex.org/topics/safety/. Accessed 12 Dec 2019.
57. Feng, Y., Wang, D., & Yin, Y et al. (2020). An XGBoost-based casualty prediction method for terrorist attacks. Complex Intell. Syst. *6*, 721–740. https://doi.org/10.1007/s40747-020-00173-0.
58. Troian, J., Arciszewski, T., & Apostolidis, T. (2019). The dynamics of public opinion following terror attacks: Evidence for a decrease in equalitarian values from Internet Search Volume Indices. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace, 13*(3), 4.
59. Wolff, K., & Larsen, S. (2014). Can terrorism make us feel safer? Risk perceptions and worries before and after the July 22nd attacks. *Annals of Tourism Research, 44*, 200–209.
60. Xia, T., & Gu, Y. (2019). Building terrorist knowledge graph from global terrorism database and Wikipedia. In *Intelligence and Security Informatics (ISI) 2019 IEEE International Conference* (pp. 194–196). IEEE. Shenzhen, China.
61. Yang, Y., Pah, A. R., & Uzzi, B. (2019). Quantifying the future lethality of terror organizations. *Proceedings of the National Academy of Sciences, 116*(43), 21463–21468.
62. Zemishlany, Z. (2012). Resilience and vulnerability in coping with stress and terrorism. *Israel Medical Association Journal, 14*(5), 307–309.