A Hybrid Recommender System for Cultural Heritage Promotion

Article

Accepted version

Agate V., Concone F., Gaglio S., and Giammanco A.

In Proceedings of the 2021 IEEE International Conference on Smart Computing (SMARTCOMP)

Publisher: IEEE

A Hybrid Recommender System for Cultural Heritage Promotion

Vincenzo Agate, Federico Concone, Salvatore Gaglio, and Andrea Giammanco
Dept. of Engineering, University of Palermo, Palermo, Italy
E-mails: {vincenzo.agate, federico.concone, salvatore.gaglio, andrea.giammanco}@unipa.it

***Abstract*—Assisting users during their cultural trips is paramount in promoting the heritage of a territory. *Recommender Systems* offer the automatic tools to guide users in their decision process, by maximizing the adherence of the proposed contents with the particular preferences of every single user. However, traditional recommendation paradigms suffer from several drawbacks which are exacerbated in *Cultural Heritage* scenarios, due to the extremely wide range of users behaviors, which may also depend on their different educational backgrounds. In this paper, we propose a *Hybrid* recommender system which combines the four most common recommendation paradigms, namely *collaborative filtering*, *popularity-*, *knowledge-*, and *content-based*, according to different hybridization strategies. Experimental evaluation shows the versatility of the hybrid recommender with respect to the other paradigms adopted individually.**

***Index Terms*—Recommender systems, cultural heritage**

I. Introduction

Recommender Systems are intelligent tools capable of reasoning upon heterogeneous sources of information in order to provide tailored suggestions [1]. Such suggestions reflect the highly varied interests of the users, which can be expressed in an explicit manner, through a vote expressed in a certain scale (e.g., from 1 to 5), or in an implicit way, such as the plain information of elements' fruition.

In the context of *Cultural Heritage*, a recommender system assumes the role of an intelligent personal guide that can be queried by users [2]. Typically, museum visitors want to maximize the number of artworks to view according to their artistic backgrounds, whether related to particular cultural currents or figurative details represented. Other external factors and constraints contribute in increasing the utility of an automatic tool [3], [4] for the assistance of users in their cultural trips. These include the limited available amount of time to conduct the trip, the wide range of artworks present in a museum, and their complex set of associated features [5]. All these problems relate to the well known issue of *information overload* [6], which may be synthesized as the inherent difficulty in making effective choices when the range of options is too wide.

The issues just introduced can be exhaustively faced by recommender systems both because of the enhancement in museum visits achieved with their use, and the induced perceived presence of a personal guide with in-depth knowledge of the art collections on display [7]. The research subfield of recommender systems was widely addressed in the literature, which proposes different approaches based on the considered application scenario. In general, all the recommendation algorithms can be lead back to some general paradigms of operation, which formalize two aspects: the structure of the input received, and the practical strategy to build a set of personalized recommendations [1]. Among the main paradigms, *content-based* approaches typically suffer from an *over-specialization*, since they over fit on the single user preferences and continue to suggest items very similar to those liked in the past [1], phenomenon also known as "rabbit hole" production [8]. On the other hand, *collaborative filtering* approaches heavily suffer from the *cold-start* problem for new users, who can't be linked to any other similar users when the available expressed ratings are not sufficient, as well as the *data sparsity* problem which arise when there are no significant intersections between the elements evaluated by the community of users [1]. Approaches belonging to the *popularity* and *knowledge* paradigms realize complementary reasonings. The former is efficient in engaging new users of the system by proposing the most appreciated compositions by the community, but fails in adapting to the specific interests expressed by the visitors of the museum. The latter, conversely, can track the interests of the users, but it requires a mandatory specification of a detailed *user profile*.

To address the limitations of these approaches, in this paper we propose a *hybrid* recommender system which blends different recommendation strategies by embodying them in distinct intelligent modules, which are combined in an innovative manner. In particular, a *Knowledge* module performs a preprocessing of the items based on user profiles, which can be inferred explicitly through a survey, or implicitly from observed behaviors [9], [10]. A *Popularity* module takes in charge of dealing with new users who have never interacted with the system, so as to gain their trust in the suggestions by proposing attractive artworks, which have collected a great success among the community of users. Popular artworks are flanked by more uncommon compositions through the operation of a *Niche* module. A *Collaborative* module intercepts and exploits the similarities between user behaviors adopting a model-based approach, which opposed to a neighborhood based one possesses higher spatial efficiency, as well as lower computational complexity. A *Content* module reasons upon the contribution of the peculiar artworks' features in the formulation of the final user's preference about the composition. The interaction between these 5 modules is orchestrated through four different strategies of combination consolidated in literature [1], leading to a versatile approach capable of adapting to the highly mutable nature of *Cultural Heritage* scenarios.

The structure of the paper is as follows: Section II reviews related works on recommender systems in the *Cultural Heritage* scenario; Section III outlines the proposed hybrid recommender system; Section IV shows the experimental analysis conducted to validate the proposal; Section V discusses the conclusions of the paper.

II. Related Work

In literature, several recommendation approaches have been investigated for different application scenarios, and the limitations of single paradigms are well known by the research community. For such a reason, in recent years, recommender systems proposals focused on hybrid solutions that combine multiple heterogeneous data sources. Hybrid approaches are particularly suited for *Cultural Heritage* scenarios, given the high diversified range of visitors' behaviors which demands various considerations in order to be properly assisted in their trips [7]. In [11], it is proposed a complex hybrid system which leverages multiple kind of data for the creation of a personalized set of recommendations. Such data include user information gathered through social networks, artworks popularity scores computed with a variant of the page rank algorithm, and user opinions mined through sentiment analysis. However, the system computational effort and entry information barrier are very high for an effective utilization by the users, also due to the analysis of behaviours on online social networks. In a similar manner, the system described in [12] gathers users' opinions on museums leveraging reactions released on social networks, merging them with the behaviour captured through a pervasive network of sensors deployed inside the museum. The system also needs an in depth ontology of the museum, which has to be compiled by specialists in *Cultural Heritage*, and who may not be available in all the possible application domains. The authors of [13] propose an hybrid solution whose logical flow is similar to the system we propose here: for new users, they adopt a "statistics-based" approach which essentially suggests the most popular artworks, whereas for existing users they combine a neighborhood based collaborative approach with a content-based one. Adopting a neighborhood based approach however, contributes to dramatically reduced system scalability when data are sparse, which is a likely occurrence in *Cultural Heritage* scenarios where users may exhibit strongly different behaviors. The system in [14] incorporates user profiles, which are represented by means of 5 features, in the well known matrix factorization collaborative filtering approach. However, adopting the matrix factorization technique lets the system inherits all of its drawbacks, namely, the difficulty in treating users with few interactions with the system, or users who behave in a totally different manner from all the other individuals in the community (problem known as "gray sheep" [1]). In [15], the authors propose an approach which formulates the cultural tour planning problem as a states space search problem by means of a rigorous logic formulation. This formulation does not take into account scenarios with new users without released ratings, so that their graph-based collaborative filtering phase can not properly handle these situations. Authors of [16] propose a recommendation algorithm for generating complete itineraries, which estimates the relevance score of points of interest leveraging personal preferences, temporal constraints and social connections. However, user preferences are limited to thematic (whether if the point of interest represents a painting or a statue) and historical (such as 18th century) aspects, which have to be extended to the multiple facets of artworks, such as the cultural context in which the composition had been made.

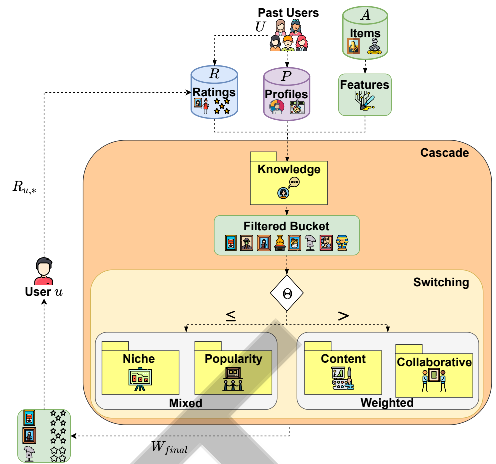

Fig. 1: The architecture of the proposed *Hybrid* system.

In light of these considerations, in this paper we propose a novel *Hybrid* recommender system which merges multiple recommendation paradigms and hybridization strategies. The goal is to offer a versatile solution capable of adapting to the highly mutable scenario of *Cultural Heritage* fruition, where both users with a strong propensity to visit the exhibitions in depth, and occasional users with little time available, all demand a personal guide to enhance their experience.

III. Hybrid Recommender System

The idea behind the *hybrid* recommender system we propose is to combine four well known recommendation paradigms, namely *Popularity*, *Knowledge*, *Content*, and *Collaborative*, and then evaluate independently the output of each paradigm in order to both take the advantages, and limit the detrimental aspects of each approach [17]. To this aim, we designed the system shown in Figure 1 in which *cascade*, *switching*, *weighted*, and *mixed* hybridization techniques are combined.

For a given user u, the algorithm processes four main parameters: the set of ratings released by the user, $R_{u,*}$; the preferences towards the features of the artworks contained in the user's profile; the number of artworks to recommend, η, also called size of the recommendation window; a discriminatory threshold, Θ, which allows to distinguish between users already recognized by the system and unknown ones. Table I introduces the notation adopted in the paper.

TABLE I: Notations used in the paper

Symbol	Description
u	An user of the system.
a	An artwork of the dataset.
U	Total number of users.
A	Total number of artworks.
R	The user-artwork ratings matrix with values $R \in \mathbb{R}^{U \times A}$.
$r_{u,a}$	Rating released by user u, for artwork a.
$\widetilde{r}_{u,a}$	Rating predicted by the system for user u w.r.t. artwork a.
$R_{u,*}$	Row of ratings released by user u.
$R_{*,a}$	Column of ratings released for artwork a.
$\overline{R}_{u,*}$	Average rating score released by user u.
$\overline{R}_{*,a}$	Average rating score released for artwork a.
$\overline{R}_{*,*}$	Average rating score released in the whole dataset.
$\lvert R_{u,*}\rvert$	Number of ratings released by the user u.
$\lvert R_{*,a}\rvert$	Number of ratings released for artwork a.
$\overline{\lvert R_{*,*}\rvert}$	Average number of ratings released by all the users.
p_u	Profile of user u.
η	Size of the final recommendations list.
π	Minimum number of ratings to label an artwork as "popular".
Θ	Threshold for distinguishing between *new* and *known* users.
α_{cb}	Weight assigned to predictions of the *Content* module.
α_{cf}	Weight assigned to predictions of the *Collaborative* module.
W_k	List generated by the *Knowledge* module.
W_p	List generated by the *Popularity* module.
W_n	List generated by the *Niche* module.
W_{cb}	List generated by the *Content* module.
W_{cf}	List generated by the *Collaborative* module.
W_{final}	The final list containing the recommended artworks.

The output of each module is a list of recommended artworks, which will be combined in an innovative manner to compose the final set of tailored suggestions, W_{final}.

A. Knowledge-based pre-filtering

At the highest level of abstraction, the *Hybrid* recommender system enables a *cascade* mechanism, whose goal is to extract a list of raw recommendations that needs to be refined in a subsequent phase. In particular, the module responsible for *knowledge-based* recommendation processes the profile of an user p_u, extracted from the *user profile* and *ratings* databases, classifies an artwork as "relevant" or "not relevant" leveraging the *Random Forest* algorithm, whose training phase works as follows. Let $\overline{R}_{u,*}$ be the average rating score released by the user u, then, the label of an artwork a in the training set is considered "relevant" if the associated rating $r_{u,a} > \overline{R}_{u,*}$; vice versa, it is considered "not relevant". The output of the model, W_k, is a *filtered bucket* of artworks in which elements labelled as "not relevant" are excluded.

Such filtered artworks are analysed through the *switching* technique that aims to activate the *weighted* or *mixed* methods based on whether the user is considered "*new*" or "*known*". Let $|R_{u,*}|$ denote the total number of ratings currently released by the user u, if $|R_{u,*}| \leq \Theta$, then u is considered to be "new". Therefore, the *switching* mechanism activates the *Popularity* and *Niche* modules, presenting their outputs in parallel. Otherwise, u is considered "known", and the *switching* technique activates the *Content* and *Collaborative* modules, presenting a weighted combination of their outputs.

B. Managing new users

If user u is considered as a "new user", the *Popularity* and *Niche* modules produce two separate lists of recommendations, namely W_p and W_n, that will be merged to form the final list of items to provide to the user. The high level reasoning behind this choice is that the system tries to engage the "new user" with the most liked artworks, flanking them with lesser-known works that might be of interest to the user.

The *Popularity* module performs the prediction algorithm used by the popular Internet Movie Database[1]. This algorithm assigns a rating $\widetilde{r}_{*,a}$ to an artwork a according the following:

$$\widetilde{r}_{*,a} = \left(\overline{R}_{*,a} * \frac{|R_{*,a}|}{|R_{*,a}| + \pi}\right) + \left(\overline{R}_{*,*} * \frac{\pi}{|R_{*,a}| + \pi}\right), \quad (1)$$

where π is the minimum number of votes required for an artwork to appear among the most popular artworks, $|R_{*,a}|$ the number of votes released for artwork a, $\overline{R}_{*,a}$ the average vote released for artwork a, and $\overline{R}_{*,*}$ be the average rating of all elements in the dataset. The equation shows that the more $|R_{*,a}|$ is greater than π, the more burdensome will be the contribution of the average vote relative to artwork a in the final computation of its predicted vote. The predicted rating of the *Popularity* module, as explicated by the notation $\widetilde{r}_{*,a}$, is not referred to any particular user u, since this module performs a global reasoning which is equal for all the users in the dataset.

The *Niche* module realizes the specular reasoning as it considers those elements whose number of evaluations present in the dataset remains lower than π. The formula for calculating the predicted vote remains Eq. 1, but in this case we will have $|R_{*,a}| < \pi$, and this results in a higher contribution of the average vote $\overline{R}_{*,*}$ of all the works contained in the dataset.

The outputs of the *Population* and *Niche* modules are finally presented in parallel to the "new user" u, according to the *mixed* hybridization paradigm [17]. In other words, the two lists W_p and W_n are concatenated together to obtain the final recommendation list W_{final}.

C. Managing known users

If user u is considered as a "known user", the system performs a *weighted* combination of the predictions coming from the *Content* and *Collaborative* modules.

The *Content* module focuses its analysis on the features characterizing the artworks, mapping the ratings released by the users into preferences towards the single features. To this aim, since a numerical encoding is required, we adopted a *one-hot encoding*, which allows to represent the categorical features of the artworks through binary vectors: the value 1 signals the presence of a certain feature, 0 otherwise. In the case of a recommender system set in *Cultural Heritage* scenario, one of the most relevant and critical features is the caption associated with the artwork, that is a verbose description. In order to use the *one-hot encoding* also for this kind of feature, the system processes the caption as set of

[1]https://help.imdb.com/article/imdb/track-movies-tv/ratings-faq/G67Y87TFYYP6TWAV#

TABLE II: Example of user profile

Demographics		**Generic Features Weights**						**Specific Features Weights**					
Age	*Sex*	*Artist Fullname*	*Artist Nationality*	*Object Type*	*Period*	*Medium*	*Other*	*Gilles Peress*	...	*French*	...	*Sake Cup*	...
32	F	0.299	0.237	0.079	0.156	0.194	0.035	0.717	...	0.587	...	0.653	...

words on which a *TF-IDF* score [18] can be computed. Once the feature vectors are generated, a linear regressor can be trained using the released votes by the current user $R_{u,*}$ as labels. By solving the regression problem for all the artworks the user u has not seen yet, the *Content* module produces a list of predicted ratings for each of these artworks, W_{cb}, which will then be combined with the prediction coming from the *Collaborative* module.

The *Collaborative* module relies on the *AutoRec* approach [19], which essentially uses an autoencoder with a single hidden layer, and an output layer of the same cardinality as the input. In this approach, predicted ratings are obtained by solving a regression problem in order to forecast the missing rating values, based on the analysis of ratings released by other users. In particular, we implemented the "item-based" variant of AutoRec, feeding the neural network with the user-artwork matrix R grouped by columns, finally producing the list of recommendations W_{cf}.

The predicted ratings of both the *Content* and *Collaborative* modules are combined through a *weighted* hybridization technique [17] as follows. Let α_{cb} be the weight for the *Content* module predictions, and let α_{cf} be the weight of the *Collaborative* module predictions. These weights are calculated adaptively with respect to the number of ratings currently available to the active user, $|R_{u,*}|$. Then, considering $\overline{|R_{*,*}|}$ to be the average number of ratings for all the users in the dataset, we compute α_{cb} as follows:

$$\alpha_{cb} = 1 - \frac{|R_{u,*}|}{|R_{u,*}| + \overline{|R_{*,*}|}} \tag{2}$$

and $\alpha_{cf} = 1 - \alpha_{cb}$. This relation expresses that the more ratings there are for the current user, the more reliance can be placed on collaborative prediction for ratings. Conversely, when the current user has few ratings, the system will place more emphasis on content-based suggestions instead. In this way, it is possible to circumvent both the over-specialization that typically affects content-based systems and the *gray sheep* problem [1] that plagues collaborative filtering.

IV. Experimental Evaluation

In order to validate the proposed *Hybrid* recommender system, the first step consists in selecting an appropriate dataset with explicit preferences released by users. To the best of our knowledge, there is not a real world dataset that provides such information in a *Cultural Heritage* scenario. For this reason, we started from a tool for the generation of synthetic datasets of ratings [20]. Then, we present an exhaustive experimental evaluation of the proposed system.

A. Dataset Generation

In order to better simulate a real scenario, we based the synthetic ratings generation on a real world dataset of artworks, namely the dataset of the Minneapolis Institute of Arts (MIA)[2]. The synthetic ratings dataset generation tool [20] allows to model artificial users with a particular profile, and generate coherent explicit preferences accordingly.

An *user profile* is fully described by the demographics info, i.e., sex and age, a set of weights for the *generic features*, and a set of weights for the *specific features* of the dataset. The *generic features* represent the high level characteristics of the artworks. In our simulation, we considered a set made up of five elements as *generic features*: the artist's full name and nationality; the object type of the artwork (e.g., sculpture or print); the period of composition; the medium in which the artwork had been realized. A *generic feature* called "other" is added in order to model unpredictable factors in the simulation of a rating. For a given user, the set of weights for the *generic features* has unitary sum, so that each single weight may easily be interpreted as the degree of importance that the user gives to the feature. The last part of the *user profile* is made up of a set of weights for the *specific features* of the dataset, which represent the particular values that a *generic feature* may assume. Unlike the set of weights associated with the *generic features*, the set of weights for the *specific features* does not sum to 1. For sake of clarity, Table II shows one example of an *user profile* generated with the described tool. This profile represents the input to the *Knowledge* module.

Given the vast amount of artworks present in the MIA dataset, it would be onerous to keep track of all the weights associated with all the possible *specific features*. For this reason, as input to our implementation of the generator, we provide a parameter ξ (which we set to $\xi = 5$) which represents the number of *specific features* to store. A group of users with the same artistic preferences share the identical subset of *specific features*. This subset is leveraged by the system for extracting artworks that will be proposed to the user for evaluation. The final synthetic rating that the simulated user releases to the proposed artwork is a weighted sum of the *specific features* times the *generic features* weights, scaled in the range of $[1, 5]$, rounding to the nearest multiple of 0.5.

We generated a dataset containing 2000 users belonging to 10 groups of similar users. We restricted the analysis to 1000 artworks of the MIA dataset. In our simulation, each user expresses a total number of ratings which lies in the range $[10, 100]$. In conclusion, we obtain 57682 total ratings, with an average number of $\overline{|R_{*,*}|} = 29$ ratings per user. The resulting user-artwork heatmap is shown in Figure 2.

[2]https://github.com/artsmia/collection

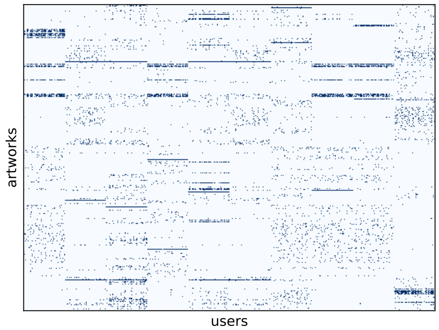

Fig. 2: The heatmap for the generated dataset of ratings with 1000 artworks and 2000 users.

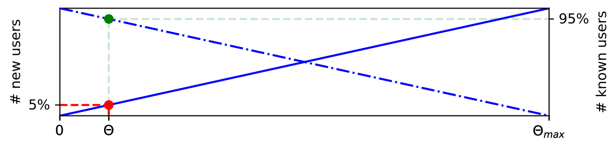

Fig. 3: The amount of "new/known users" against Θ.

B. Evaluation Metrics

The performance of the system was evaluated by adopting standard classification metrics aimed to capture the capability of the recommender of separating *relevant* and *irrelevant* artworks [21]. In particular, an artwork a is considered *relevant* for a given user u, if the associated rating is above the average score released by u, i.e., $r_{u,a} > \overline{R}_{u,*}$. As a consequence, each recommendation can be: a True Positive (TP), when a relevant artwork is recommended to the user; a True Negative (TN), when an irrelevant artwork is not recommended to the user; a False Negative (FN), if a relevant artwork is not recommended to the user; a False Positive (FP), if an irrelevant artwork is recommended to the user. Given these values, three classification metrics are used by adopting the notation $@\eta$, which indicates that the evaluation is restricted to the recommendation window η. When an artwork in the recommendation window does not have a correspondence among the test set for the current user, we neglect the given artwork in the evaluation. The three metrics are as follows.

$\boldsymbol{precision@\eta}$ is defined as the ratio between the number of relevant artwork recommended and the total number of recommended artworks, and represents the probability that a recommended artwork matches the user preferences [21]: $precision@\eta = \frac{TP}{TP+FP}$.

$\boldsymbol{recall@\eta}$ is defined as the ratio between the number of relevant artworks recommended and the total number of relevant artworks, and represents the probability that a relevant artwork is recommended within the η elements in the recommendation window [21]: $recall@\eta = \frac{TP}{TP+FN}$.

$\boldsymbol{fscore@\eta}$ is defined as the harmonic mean between $precision@\eta$ and $recall@\eta$.

C. Performance Evaluation

Before describing its performances, we briefly introduce the hyperparameters adopted in the proposed system. The autoencoder of the *Collaborative* module is composed of 500 units in the hidden layer, whereas the number of input and output neurons is equal to the number of users in the dataset, the activation functions are the identity function for the hidden layer, and the sigmoid function for the output layer, and the whole network is trained with 20 epochs using the Adam optimizer with learning rate set to 0.001 and $\lambda = 1e-8$.

In order to tune the parameter Θ we can refer to Figure 3, in which the relationship between "new users" and "known users" is shown. In particular, it is described the percentage of the two categories of users recognized by the *switching* mechanism for a dataset, against the range of possible values of $\Theta \in [0, \Theta_{max}]$, where Θ_{max} is equal to the maximum amount of ratings released by any user. In our simulation, the value of Θ has been set in order to have an amount of 5% "new users", and 95% "known users" in each of the considered subsets. In a similar way, the value of π is set in order to classify as "popular" an amount of 5% of the artworks in each of the considered subsets. In both cases, these parameters have been chosen to fit the addressed scenario; in general, they can be tuned according to the application specific goals.

We performed several experiments, by considering different portions of the generated dataset of ratings. In particular, the portions contain an amount of $|U|$ users which are selected randomly for each run. We executed 50 runs for each amount $|U| \in \{50, 100, 150, 200, 250\}$. During a given run, the subset of ratings is split in 80% train test and 20% test set. We investigated the size of the recommendation window $\eta \in \{5, 10, 15\}$.

Figure 4 shows the experimental results, averaged across all the 50 runs, and plotted against all the considered subset of users, for the *Popularity*, *Content* and *Collaborative* modules, together with the *Hybrid* system on the whole. The performances of the *Hybrid* system express a trade-off between the accuracy of the pure *Popularity*, *Content* and the *Collaborative* approaches. Analyzing scenarios in which the number of users is quite low, i.e. $\{50, 100\}$, it is possible to notice how the best performances are obtained by the *Content* module, followed by the proposed *Hybrid* approach. Conversely, as the number of users increases, i.e. $\{150, 200, 250\}$, the highest performances are obtained by the *Collaborative* module, followed by the *Hybrid* system. Indeed, increasing $|U|$, the ability of the *Collaborative* module in finding correlations between user behaviors grows accordingly. The experiments show that the recommendation window does not affect the fscore. However, an opportune sized and well tailored set of suggestions play a fundamental role in enhancing the user engagement with the system. We now analyze Fig. 4 averaging the results obtained by both varying η and $|U|$. Compared to the single modules, the *Hybrid* solution achieves the following improvements in $precision$ and $recall$ respectively: 91.63% and 87.44% w.r.t. the *Popularity* module; 6.60% and 6.18% considering the *Content* module; 0.43% and −0.52% examining the *Collaborative* module. We finally notice how the proposed *Hybrid* system can manage both situations with low and high number of users, achieving the best trade-off in the performance metrics. Being able to efficiently handle these mutable situations is paramount in *Cultural Heritage*, especially when the number of visitors can be extremely varied.

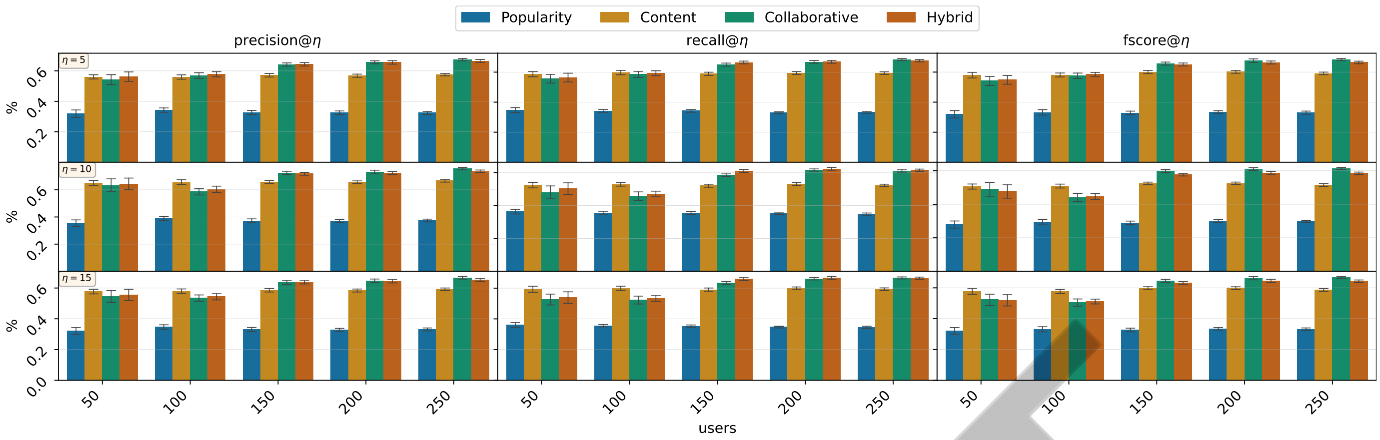

Fig. 4: The classification metrics for the *Hybrid Recommender System* and the single modules.

V. Conclusions

In this paper we proposed a novel *Hybrid Recommender System* for *Cultural Heritage* scenarios. Our system merges four different recommendation paradigms, and four different hybridization strategies. The experimental results show how this *Hybrid* solution represents a versatile trade-off capable of adapting to situations with mutable number of users. As part of our future work, we plan to define a smarter switching mechanism which allows to better decide the situations in which the *Popularity* and *Niche* modules would be more effective in catching the user's interests. We also plan to increase the security of the system against shilling attacks with the support of a tailored reputation system [22]. Finally, we want to extend the experimental evaluation exploiting real users data that will be available as part of the italian VASARI research project.

Acknowledgment

This research is partially funded by the Project VASARI of Italian MIUR (PNR 2015-2020, DD MIUR n. 2511).

References

[1] C. C. Aggarwal, *Recommender systems - The textbook*. Springer, 2016, vol. 1.

[2] P. Ferraro and G. Lo Re, *Designing Ontology-Driven Recommender Systems for Tourism*. Cham: Springer International Publishing, 2014, pp. 339–352.

[3] T. Catarci, F. Leotta, A. Marrella, M. Mecella, D. Sora, P. Cottone, G. Lo Re, M. Morana, M. Ortolani, V. Agate, G. R. Meschino, G. Pecoraro, and G. Pergola, "Your friends mention it. what about visiting it? a mobile social-based sightseeing application," in *Proceedings of the International Working Conference on Advanced Visual Interfaces*, ser. AVI '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 300–301.

[4] V. Agate and S. Gaglio, "A gesture recognition framework for exploring museum exhibitions," in *Proceedings of the 2018 International Conference on Advanced Visual Interfaces*, ser. AVI '18. New York, NY, USA: Association for Computing Machinery, 2018.

[5] F. Concone, R. Giaconia, G. Lo Re, and M. Morana, "A smart assistant for visual recognition of painted scenes," in *26th International on Intelligent User Interfaces (IUI), April 13-17, 2021, College Station, USA*. CEUR-WS, 2021.

[6] S. S. Iyengar and M. R. Lepper, "When choice is demotivating: Can one desire too much of a good thing?" *Journal of personality and social psychology*, vol. 79, no. 6, p. 995, 2000.

[7] G. Pavlidis, "Recommender systems, cultural heritage applications, and the way forward," *Journal of Cultural Heritage*, vol. 35, 2019.

[8] L. Tang, K. Fujimoto, M. T. Amith, R. Cunningham, R. A. Costantini, F. York, G. Xiong, J. A. Boom, and C. Tao, ""down the rabbit hole" of vaccine misinformation on youtube: Network exposure study," *Journal of Medical Internet Research*, vol. 23, no. 1, p. e23262, 2021.

[9] F. Concone, G. L. Re, and M. Morana, "A fog-based application for human activity recognition using personal smart devices," *ACM Trans. Internet Technol.*, vol. 19, no. 2, Mar. 2019.

[10] F. Concone, S. Gaglio, G. Lo Re, and M. Morana, "Smartphone data analysis for human activity recognition," in *AI*IA 2017 Advances in Artificial Intelligence*. Cham: Springer International Publishing, 2017, pp. 58–71.

[11] X. Su, G. Sperlì, V. Moscato, A. Picariello, C. Esposito, and C. Choi, "An edge intelligence empowered recommender system enabling cultural heritage applications," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 7, pp. 4266–4275, 2019.

[12] M. Torres-Ruiz, F. Mata, R. Zagal, G. Guzmán, R. Quintero, and M. Moreno-Ibarra, "A recommender system to generate museum itineraries applying augmented reality and social-sensor mining techniques," *Virtual Reality*, vol. 24, no. 1, pp. 175–189, 2020.

[13] L. Kovavisaruch, T. Sanpechuda, K. Chinda, T. Wongsatho, S. Wisadsud, and A. Chaiwongyen, "Incorporating time constraints into a recommender system for museum visitors," *Journal of information and communication convergence engineering*, vol. 18, no. 2, pp. 123–131, 2020.

[14] M. Konstantakis, G. Alexandridis, and G. Caridakis, "A personalized heritage-oriented recommender system based on extended cultural tourist typologies," *Big Data and Cognitive Computing*, vol. 4, no. 2, 2020.

[15] F. Amato, F. Moscato, V. Moscato, F. Pascale, and A. Picariello, "An agent-based approach for recommending cultural tours," *Pattern Recognition Letters*, vol. 131, pp. 341–347, 2020.

[16] L. Rajaonarivo, A. Fonteles, C. Sallaberry, M.-N. Bessagnet, P. Roose, P. Etcheverry, C. Marquesuzaà, A. L. P. Lacayrelle, C. Cayèré, and Q. Coudert, "Recommendation of heterogeneous cultural heritage objects for the promotion of tourism," *ISPRS International Journal of Geo-Information*, vol. 8, no. 5, 2019.

[17] R. Burke, "Hybrid recommender systems: Survey and experiments," *User modeling and user-adapted interaction*, vol. 12, no. 4, 2002.

[18] G. Salton, *Automatic Text Processing*. USA: Addison-Wesley Longman Publishing Co., Inc., 1988.

[19] S. Sedhain, A. K. Menon, S. Sanner, and L. Xie, "Autorec: Autoencoders meet collaborative filtering," in *Proceedings of the 24th international conference on World Wide Web*, 2015, pp. 111–112.

[20] M. del Carmen Rodríguez-Hernández, S. Ilarri, R. Hermoso, and R. Trillo-Lado, "Datagencars: A generator of synthetic data for the evaluation of context-aware recommendation systems," *Pervasive and Mobile Computing*, vol. 38, pp. 516–541, 2017.

[21] A. Gunawardana and G. Shani, "A survey of accuracy evaluation metrics of recommendation tasks." *Journal of Machine Learning Research*, vol. 10, no. 12, 2009.

[22] V. Agate, A. D. Paola, G. L. Re, and M. Morana, "A simulation software for the evaluation of vulnerabilities in reputation management systems," *ACM Trans. Comput. Syst.*, vol. 37, no. 1–4, Jun. 2021.