



Personalisation in cyber-physical-social systems: A multi-stakeholder aware recommendation and guidance

Bereket Abera Yilma, Yannick Naudet, Hervé Panetto

► To cite this version:

Bereket Abera Yilma, Yannick Naudet, Hervé Panetto. Personalisation in cyber-physical-social systems: A multi-stakeholder aware recommendation and guidance. 29th ACM Conference on User Modeling, Adaptation and Personalization, UMAP'21, Jun 2021, Utrecht, Netherlands. pp.251-255, 10.1145/3450613.3456847 . hal-03270669

HAL Id: hal-03270669

<https://hal.science/hal-03270669>

Submitted on 25 Jun 2021

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

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

Personalisation in Cyber-Physical-Social Systems: A Multi-stakeholder aware Recommendation and Guidance

Bereket Abera Yilma
Luxembourg Institute of Science and
Technology (LIST)
Esch-sur-Alzette, Luxembourg
Universite de Lorraine, CNRS, CRAN
Nancy, France
bereket.yilma@list.lu

Yannick Naudet
Luxembourg Institute of Science and
Technology (LIST)
Esch-sur-Alzette, Luxembourg
yannick.naudet@list.lu

Hervé Panetto
Universite de Lorraine, CNRS, CRAN
Nancy, France
herve.panetto@univ-lorraine.fr

ABSTRACT

The evolution of smart devices has led to the transformation of many physical spaces to the so-called *smart environments* collectively termed as *Cyber-Physical-Social System (CPSS)*. In CPSS users co-exist with different stakeholders influencing each other while being influenced by different environmental factors. Additionally, these environments often have their own desired goals and corresponding set of rules in place expecting people to behave in certain ways. Hence, in such settings classical approaches to personalisation which solely optimise for user satisfaction are often encumbered by competing objectives and environmental constraints which are yet to be addressed jointly. In this work we set out to (i) formalise the general problem of personalisation in CPSS from a multi-stakeholder perspective taking into account the full environmental complexity, (ii) extend the general formalisation to the case of exhibition areas and propose a personalised Multi-stakeholder aware Recommendation and Guidance method on a case study of National Gallery, London.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Personalisation, Cyber-Physical-Social System, multi-stakeholder aware recommendation

ACM Reference Format:

Bereket Abera Yilma, Yannick Naudet, and Hervé Panetto. 2021. Personalisation in Cyber-Physical-Social Systems: A Multi-stakeholder aware Recommendation and Guidance. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization (UMAP '21)*, June 21–25, 2021, Utrecht, Netherlands. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3450613.3456847>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

UMAP '21, June 21–25, 2021, Utrecht, Netherlands

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8366-0/21/06...\$15.00

<https://doi.org/10.1145/3450613.3456847>

1 INTRODUCTION

A Cyber-Physical-Social System (CPSS) characterises environments that are cohabited by humans and sensor enabled smart devices [26][27]. Smart Cities, Smart Homes, Schools, Offices, Museums, and medium to large scale industries are among the main examples where application of the CPSS notion has gained momentum [6, 16, 25, 27]. In these environments the introduction of smart devices and their tight link with daily operations has revolutionised the livelihood of people in various aspects. Nevertheless, it is also evident that the evolution towards a more and more pervasiveness of technology is gradually increasing the complexity of CPSS environments. In a CPSS people interact with each other as well as with different kinds of devices and robots that offer a variety of services. Particularly ensuring a seamless experience for people requires taking into account individual's personality [12]. This is because each person is unique and individual needs, preferences and capabilities also vary from person to person. This positions personalisation at the heart of the CPSS paradigm. However, unlike the case of virtual personalisation, in CPSS context people evolve in a physical space which by itself brings additional properties and dynamic variables to be considered [24]. Additionally multiple objectives of the co-existing stakeholders together with the complexity of human behaviour make the task of personalisation rather complex. In this work we tackle the problem of personalisation in CPSS context aiming to jointly addresses multi-stakeholder issues and environmental complexities. We first establish background on personalisation in CPSS environments covering relevant literature and present a general formalisation to the problem. Following this we instantiate the formalisation on the case of exhibition areas and propose a personalised Multi-stakeholder aware recommendation and guidance method on a case study of National gallery of London.

2 BACKGROUND

2.1 Personalisation in CPSS

The task of personalisation in CPSS environments such as smart homes, museums and art galleries has been an area of investigation in recent years. [4, 5, 8, 17, 21]. For instance enhancing user experience resorting game theory and optimisation methods[20], augmented reality [11], inferring art preferences from gaze [7], path recommendation in museum [18, 22], etc. Most of these works solely focus on matching preferences and interests of users in their recommendations. Thus, they often tend to neglect environmental constraints and the co-existing objectives of different stakeholders

in the same physical space influencing the user. In personalising and making CPSS environments adaptable to a particular person one should make the best possible compromise between the objectives of the user, the co-existing entities and the global objective of the CPSS while respecting environmental constraints [16, 24]. This allows to not only mitigate influential factors on the user but also maintain a desirable state for all stakeholders. Although, a multi-stakeholder approach to recommendation has gained attention in recent years [1, 2, 23, 28–30], its application is mostly limited to virtual personalisation and physical space challenges are not yet sufficiently addressed. Hence, this particular scenario leads to formulate the problem of personalisation in a CPSS context as a function of the main stakeholders. (i.e. the user u of personalisation service, the CPSS in which she evolves in $cpss$, the crowd of other people in the CPSS cr , the application device that implements the personalisation service d and the global context cx) written as:

$$Perso_u^{(CPSS)} = f(u, cpss, cr, d, cx) \quad (1)$$

Here, cx refers to the set of all other elements in the CPSS $\{x_1, x_2, \dots, x_n\}$ that have no direct or indirect influence on the user. When any of the elements in cx has an impact on the user/personalisation it will be taken as part of the formalisation f as $f(u, cpss, cr, d, x_i, cx)$; $\forall x_i \in cx$. In order to illustrate how this can be used in tackling the problem of personalisation in complex physical environments, we instantiate the formalisation on the case of exhibition areas.

2.2 Personalisation in Exhibition areas

Exhibition areas such as Museums, Galleries, temporary expose, etc. are the kind of environments that are usually composed of a large number of points of interest (POIs) to be explored by many visitors. In such environments visitors often evolve with others carrying their sensor enabled smart devices (i.e. smart phones, wearables, cameras...etc) serving corresponding individual needs. The exhibition areas themselves are also equipped with various sensors and actuating components for various purposes. In exhibition areas the physical space has rules expecting people to behave in a certain way. However, visitors are not always willing to follow museum rules and recommendations unless aligned to their state of mind and preferences. Exhibition curators also have desired goals such as making less popular items more visible, reducing congestion around popular exhibits...etc. In larger places such as the National Gallery or Louvre¹, visitors often tend to miscalculate their available time or sometimes get lost and spend wondering around the museum. Hence, offering personalisation in such systems not only improves the visitor's quality of experience but also gives meaning to the presence of people inside the museums. Nevertheless, due to the environmental complexity this task requires solutions beyond just matching user preferences and interests. This means satisfying multi-stakeholders under a constrained environment. Thus, adopting the global formalisation equation 1, the task of personalisation in exhibition areas for a user u can be formalised as a function of the user u , the CPSS which translates to the exhibition area ex , the crowd of other visitors cr , the application d and the context cx written as:

$$Perso_u^{(ex)} = f(u, ex, cr, d, cx) \quad (2)$$

This leads to formulate the task as constrained multi-objective optimisation problem. Based on this premise, in the next section we propose a personalised recommendation and guidance method on a case-study of National Gallery, London.

3 PERSONALISED RECOMMENDATION AND GUIDANCE IN NATIONAL GALLERY

The National Gallery is an art museum housing a collection of over 2,300 paintings dating from the mid-13th century to 1900. This gallery located in Trafalgar Square in the City of Westminster, in Central London has a total floor area of 46,369 square meters². In 2019 it ranked 3rd nationally with 6.2 million visitors. Due to the large size and number of exhibits the gallery has been employing different techniques in its large multi-thematic venues to assist visitors and improve their quality of experience. The complex nature of the physical space together with the gradual introduction of different technologies into the gallery magnifies the CPSS nature. Hence, delivering personalised services in terms of guiding visitors to improving quality of experience while managing crowd and satisfying different goals of exhibition curators remains an open challenge. In this case-study we propose a recommendation and guidance method based on the formalisation of equation 2 which concurrently personalises for visitors and curators while respecting environmental constraints. A recommendation service in an exhibition area essentially constitutes two major components. Firstly selecting best POIs and secondly finding an optimal route to visit the selected POIs. Below we present the formulations for these two components.

3.1 POI recommendation

Let $P = \{p_1, p_2, \dots, p_y\}$ be the set of all paintings in the gallery, $P^u = \{p_1^u, p_2^u, \dots, p_n^u\}$ be the set of paintings liked or rated by a user u , and $W^u = \{w_1^u, w_2^u, \dots, w_n^u\}$ be weights representing user ratings. The central idea of POI recommendation is to suggest relevant POIs to visitors matching their preferences while at the same time reflecting the exhibition curator's goals. In this work we used explicit profiling to elicit user preferences and employed a Latent Dirichlet allocation (LDA) model trained on textual description of paintings from National gallery of London. This was previously proved to capture latent semantic similarities among the paintings[3]. As a primary stakeholder, we operationalize the notion of relevance for visitors and select a recommendation set of paintings if they closely resemble the visitor's profile. We then used this model to define a relevance score $S(p_i, u)$ for every painting p_i in our dataset P , according to the preferences P^u of user u . This is calculated based on the weighted average similarity (cosine distance) from all other paintings in P^u computed as $\frac{1}{N} \sum_{j=1}^N w_j * d(p_i, p_j)$; where w_j is weight representing the preference of user u for a painting p_j and $d(p_i, p_j)$ is the similarity between painting p_i and p_j , $p_j \in P^u$ according to the painting-LDA model. The higher the relevance score the more relevant a painting is to a user based on his preferences. In addition to unique personal preferences users also have different tendency to be interested in visiting famous POIs. Hence, we introduces a popularity score $S(p_i, Pop)$ for the paintings in the dataset.

¹<https://www.louvre.fr/>

²<https://theculturetrip.com/europe/united-kingdom/england/london/articles/history-of-the-national-gallery-london/>

This score is based on a MUST-SEE list generated according to public review from National Gallery website. By taking into account the preference of the user and also the influence of the crowd on the popularity of the paintings we introduce an aggregate score $S_{AG}(p_i, u)$ for the paintings in the dataset which is given by:

$$S_{AG}(p_i, u) = S(p_i, u) + \beta S(p_i, Pop) \quad (3)$$

where β is user provided hyperparameters determining user's interest to see popular items. The other prevalent objective of personalisation in this context comes from the exhibition curator. Regarding POI recommendation in National Gallery we identified 9 curated stories defined by artists (*Women's Lives, Contemporary Style & Fashion, Water, Women Artists & Famous Women, Warfare, Monsters & Demons, Migration & Exile, Death, Battles & Commanders*). Each story constitutes a certain number of paintings from the collection. The objective of the curators here is to increase the number of curated stories in the recommendation set. (i.e. the recommendation set contains paintings that are fairly selected from the curated story groups). Thus we define a fair story selection strategy adopted from [14]. Under the current assumption a recommendation set R is fair if it contains paintings that belong to different story groups. The fairness score for a set R that contains paintings belonging to only one or few of the story groups is lower than the one that covers all or most of the story groups. To this end we compute a fair story selection function $\psi(R)$ as:

$$\psi(R) = \sum_{i=1}^K \sqrt{\sum_{p_a \in S_i \cap R} \gamma_{p_a}} \quad (4)$$

where K is the total number of story groups, S_i is the i^{th} story group and γ_{p_a} is a count for every painting p_a selected from a story group i . $\psi(R)$ rewards recommendation sets that are diverse in terms of the different story groups covered. Given the nature of the function, there is more benefit to selecting a painting from a story group not yet having one of its paintings already chosen. As soon as a painting is selected from a story group, other paintings from the same story group start having diminishing gain owing to the square root function (e.g. For $K = 3$ a recommendation set that chooses 2 paintings from S_1 and 1 Painting from each of S_2 and S_3 gets a higher score of $\psi(R)$ compared to a recommendation set that chooses 4 paintings from just S_1 . i.e. $\sqrt{2} + \sqrt{1} + \sqrt{1} > \sqrt{4} + \sqrt{0} + \sqrt{0}$). The discussion above yields two different policies to investigate. The first policy is matching user preferences. (i.e. Given a set of paintings P and a user u we select the most relevant set R to recommend which maximizes the following:

$$\text{Policy1} : \left\{ \operatorname{argmax} \sum_{i=1}^R S_{AG}(p_i, u) \right\} \quad (5)$$

POSTULATE 1. *It is well known that optimizing for user preferences has a positive impact on user satisfaction [9, 10, 13]. Thus, we expect to achieve higher user satisfaction for this policy as a baseline.*

The second Policy we investigate is matching the curator's goal. This aims at recommending POIs that maximally cover the story groups in the dataset given by:

$$\text{Policy2} : \left\{ \operatorname{argmax} \psi(R) = \operatorname{argmax} \sum_{i=1}^K \sqrt{\sum_{p_a \in S_i \cap R} \gamma_{p_a}} \right\} \quad (6)$$

(i.e. maximising over $\psi(R)$ ensures the diversity of the recommendation set in terms of the story groups covered. However, this could

also be a minimisation depending on the curator's goal (i.e. minimising over $\psi(R)$ maintains the consistency of the recommendation set regarding the curated stories presented). The task of POI recommendation here is to concurrently satisfy the two policies. Hence, in this approach we depart from solely optimising for user preferences unlike the classical cases instead we combine the two objectives. Additionally we take into account the available time of visitors in order to limit the size of recommendation set. This is a soft constraint introduced by estimating visiting times per POI depending on the type of visitors. For this we adopted the analogy used in the work of *Najbri et al.*[15] to classify museum visitors in to four visiting style metaphors (Ant, butterfly, fish and grasshopper). Thus, the task is to suggest a recommendation set $R(u)$ for a user u by solving the following MIP problem:

$$\operatorname{argmax} \left(1 - \xi \sum_{a=1}^R p_a * S_{AG}(p_a, u) + \xi \sum_{i=1}^K \sqrt{p_a \sum_{p_a \in S_i \cap R} \gamma_{p_a}} \right) \quad (7)$$

$$\text{s.t.} \sum_{a=1}^R T_v(p_a) \leq T_{ava} \quad (8)$$

Where ξ is a user provided hyperparameter determining user's tolerance to diversity, $T_v(p_a)$ is the estimated visiting time for the painting p_a according to the visiting style of the user and T_{ava} is the total available time of the user.

3.2 Path recommendation

The second component of personalisation in exhibition areas is that of optimal path recommendation. Once a recommendation set R is generated. We then map the paintings in R to their corresponding venues. Thus, the task of Path recommendation is to find an optimal traversing path between the venues containing recommended paintings. Let $V = \{v_1, v_2, \dots, v_v\}$ be a set of all venues in the gallery and $V_R = \{v_1, v_2, \dots, v_q\}$ be a set of venues containing R . The optimal path is defined as the one that is of high relevance and the least expensive (shortest) which is subject to crowd size and time constraints. Since a path is a combination of venues, the relevance of a path is the total sum of the relevance of each venue on a path. The relevance score $S(v_i)$ of a venue $v_i \in V$ is defined as the total sum of the relevance scores of the paintings from the recommendation set R that are found in v_i . Depending on the interest of a visitor, $S(v_i)$ could be defined in two different ways as (Quality or Quantity). The deadlock here is that two venues say v_i and v_j could have the same relevance score $S(v_i) = S(v_j)$ but one might be composed of a single painting of very high score while the other is composed of many paintings with lower scores. Thus, in order to prioritise between venues containing recommendations we define a Quality relevance score $\Theta(v_i)$ and Quantity relevance score $\delta(v_i) \forall v_i \in V_R$ as:

$$\Theta(v_i) = \sum_{a=1}^h S_{AG}(p_a, u) + \frac{x}{K}; \quad \delta(v_i) = |h_i| + \frac{x}{K} \quad (9)$$

where h is the total number of paintings that are in $v_i \cap R$, x is the number of story groups covered by the recommended paintings in v_i . $\frac{x}{K}$ contributes diversity to the quality score. $\Theta(v_i)$ is taken as relevance for users interested in visiting fewer but most relevant paintings (i.e. paintings with high score $S_{AG}(p, u)$). In the contrary,

for those users that are interested in covering as many relevant paintings as possible we use $\delta(v_i)$.

The second objective of the optimal path recommender is to find a path with minimal cost that navigates the user from the current venue to all relevant venues and back. This can easily be solved as the travelling salesman(TSP) problem. However, the physical lay out of the museum forces visitors to traverse more than once in some venues. Hence, we rather define the cost of a path as the total sum of the travel distance between consecutive venues ($\sum_{i=1}^M \text{dist}(v_i, v_{i+1})$). Following the above discussion we get two policies to investigate in order to find an optimal path. The first one is to maximize the relevance of the venues that make up the path and the second one is minimizing the cost of the path given by:

$$\text{Policy1} : S(v_i) = \begin{cases} \arg\max \sum_{i=1}^M \Theta(v_i) & \text{Quality} \\ \arg\max \sum_{i=1}^M \delta(v_i) & \text{Quantity} \end{cases} \quad (10)$$

$$\text{Policy2} : \left\{ \arg\min \sum_{i=1}^M \text{dist}(v_i, v_{i+1}) \right\} \quad (11)$$

Combining this two policies, the optimal path $PT(u)$ recommended for a user u is computed by solving the following MIP problem:

$$\arg\max \left(1 - \lambda \sum_{i=1}^M v_i * S(v_i) + \lambda v_i * \frac{1}{\sum_{i=1}^M \text{dist}(v_i, v_{i+1})} \right) \quad (12)$$

$$s.t \sum_{i=1}^R T v(p_i) + T t(v) \leq T_{ava}; \quad (13)$$

$$\begin{aligned} T t(v) &= \sum_{i=1}^{M-1} T t(v_i, v_{i+1}) \\ C r_s(V_i) &\leq C r_t(u); \\ \forall V_i, 1 \leq i \leq M \end{aligned} \quad (14)$$

where λ is a user provided hyperparameter that entails user's tolerance to fatigue (i.e. walking). Constraint (13) entails that the total estimated time for visiting (Tv) and traveling (Tt) should not exceed the available time(T_{ava}) of the visitor. Constraint (14) entails that the crowd size $Cr_s(v_i)$ in every venue v_i on the optimal path should not exceed the Crowd tolerance $Cr_t(u)$ of the user.

4 PRELIMINARY EXPERIMENT & RESULTS

4.1 Experimental setup & Dataset

We studied the performance of our method in a controlled experiment conducted with 40 real users. Our experimental set up follows closely the one used in Yilma et al. [3]. We used a similar methodology and dataset of painting descriptions to train our painting-LDA model. This was used as a baseline recommender along with ResNet-50 (a 50-layer deep convolutional neural network trained on more than a million images from the ImageNet database³) to compare against our method. The POI dataset used in our experiment contains 2368 paintings from the National Gallery of London which were categorised into 9 curated story groups. Each user provided a rating for 80 paintings from the dataset to be used for profiling. All algorithms were implemented using python leveraging Gurobi 9.0⁴ to solve the core MIP problems. All experiments were run on a 1,4 GHz Intel CPU with 5 cores and 16 GiB of RAM.

4.2 Results

In order to compare the performance of our method with LDA and ResNet-50 we adopted a user-centric evaluation framework [19]. This was done through a questionnaire, where participants had to express their opinion in a five points Likert scale to four of the following statements. “The recommendations match my personal preferences and interests” (Predictive accuracy); “The recommender helped me discover paintings I did not know before” (Novelty); “The recommender helped me discover surprisingly interesting paintings I might not know otherwise.” (Serendipity); “The recommended paintings are diverse” (Diversity). Figure 1 summarizes the results in terms of average values of user ratings for the three recommendation pipelines. As reported on the figure, our method (CPSS) achieved significantly higher diversity values (3.9/5) compared to LDA (3.6/5) and ResNet-50 (2.3/5). This can be attributed to the introduction of popularity bias and fair story selection strategy employed in our algorithm. LDA performs slightly lower than CPSS but shows significantly better performance compared to ResNet-50 in terms of diversity. This is mainly due to the fact that the similarity scoring mechanism in LDA is based on semantically dominant concepts shared among the paintings. This was proved to uncover semantic relationships of paintings that do not necessarily have resembling visual features[3]. In particular LDA based recommendation contain visually diverse but semantically related paintings. Since we also used LDA in our method as a first layer scoring function it contributes to the higher diversity in the CPSS together with the popularity and fairness biases. This can also justify the higher values of serendipity and Novelty. In terms of prediction accuracy (i.e matching user's preferences) LDA performed better than CPSS and ResNet-50. The slightly lower accuracy of CPSS can also be attributed to the popularity and fairness bias compared to LDA and ResNet-50 whose objective solely optimises for user preferences. This result validates Postulate 1. Although in CPSS we considered the influence of the other stakeholders in the system (i.e crowd and curator), optimising for multiple objectives did not cause a significant worsening of performance accuracy (4.5/5) to (4.2/5). We also observed that this is significantly better than ResNet-50 (3.1/5). This supports the claim in Yilma et al. [3] which deduced that the exploration of textual descriptions of paintings with LDA empowers visual art recommenders to reveal hidden semantic relationships among paintings. In all experiments computational time was measured to an average of 2 seconds which is fast enough given the size of the search space. Furthermore in the survey participants were asked if the CPSS recommendations were acceptable relative to LDA and ResNet-50. Interestingly 95% of the participants responded “Yes”. This shows that the compromise to concurrently satisfy the goals of the curator and visitors was reasonable. In general the reported results in our experiment illustrate that the multi-stakeholder approach is a promising direction to pursue for personalisation in such CPSS settings. Nevertheless, these results are preliminary and conducted under a controlled experiment. Thus, further experimental validation in a non-controlled setting is viable to further justify the results of this work.

³ImageNet. <http://www.image-net.org>

⁴<https://www.gurobi.com/>

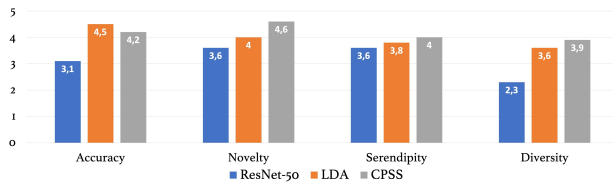


Fig. 1. Comparison of ResNet-50, LDA and CPSS interms of Accuracy, Novelty, Serendipity and Diversity

5 CONCLUSION AND FUTURE WORK

In this work we tackled the problem of personalisation in CPSS aiming to jointly address multi-stakeholder issues and environmental complexities. In particular we proposed a multi-stakeholder aware recommendation and guidance method for exhibition areas formulated as a constrained multi-objective optimisation problem. Results obtained from a preliminary experiment indicate that a multi-stakeholder approach to personalisation in CPSS is a promising direction. It not only compares favorably against baselines but also allows to better manage the environment as it makes best possible compromises between coexisting objectives of stakeholders. Although our method offers POIs recommendations and optimal traversal, the dynamic nature of the environment imposes more challenges as the visit commences. For instance variations in crowd size, closure of certain venues or the deviation of the user from the recommended path due to unknown reasons, etc. In future work we plan to integrate a real time tracking and optimising strategy to our method.

REFERENCES

- [1] Himan Abdollahpour, Gediminas Adomavicius, Robin Burke, Ido Guy, Dietmar Jannach, Toshihiro Kamishima, Jan Krasnodebski, and Luiz Pizzato. 2020. Multistakeholder recommendation: Survey and research directions. *User Modeling and User-Adapted Interaction* 30, 1 (2020), 127–158.
- [2] Himan Abdollahpour and Robin Burke. 2019. Multi-stakeholder recommendation and its connection to multi-sided fairness. *arXiv preprint arXiv:1907.13158* (2019).
- [3] Bereket Abera Yilma, Najib Aghenda, Marcelo Romero, Yannick Naudet, and Herve Panetto. 2020. Personalised Visual Art Recommendation by Learning Latent Semantic Representations. *arXiv e-prints* (2020), arXiv–2008.
- [4] Lyuba Alboul, Martin Beer, and Louis Nisiotis. 2019. Robotics and virtual reality gaming for cultural heritage preservation. (2019).
- [5] Saeed Amal, Mustafa Adam, Peter Brusilovsky, Einat Minkov, Zef Segal, and Ts vi Kuflik. 2020. Demonstrating Personalized Multifaceted Visualization of People Recommendation to Conference Participants. In *Proceedings of the 25th International Conference on Intelligent User Interfaces Companion*. 49–50.
- [6] Christos G Cassandras. 2016. Smart cities as cyber-physical social systems. *Engineering* 2, 2 (2016), 156–158.
- [7] Sylvain Castagnos, Florian Marchal, Alexandre Bertrand, Morgane Colle, and Djailila Mahmoudi. 2019. Inferring Art Preferences from Gaze Exploration in a Museum. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*. 425–430.
- [8] María del Carmen Rodríguez-Hernández, Sergio Ilarri, Ramón Hermoso, and Raquel Trillo-Lado. 2017. Towards trajectory-based recommendations in museums: evaluation of strategies using mixed synthetic and real data. *Procedia computer science* 113 (2017), 234–239.
- [9] SeoYoung Lee and Junho Choi. 2017. Enhancing user experience with conversational agent for movie recommendation: Effects of self-disclosure and reciprocity. *International Journal of Human-Computer Studies* 103 (2017), 95–105.
- [10] Ting-Peng Liang, Hung-Jen Lai, and Yi-Cheng Ku. 2006. Personalized content recommendation and user satisfaction: Theoretical synthesis and empirical findings. *Journal of Management Information Systems* 23, 3 (2006), 45–70.
- [11] Eran Litvak and Tsvi Kuflik. 2020. Enhancing cultural heritage outdoor experience with augmented-reality smart glasses. *Personal and Ubiquitous Computing* (2020), 1–14.
- [12] Ioanna Lykourantzou, Angeliki Antoniou, Yannick Naudet, and Steven P Dow. 2016. Personality matters: Balancing for personality types leads to better outcomes for crowd teams. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*. 260–273.
- [13] Rishabh Mehrotra, James McInerney, Hugues Bouchard, Mounia Lalmas, and Fernando Diaz. 2018. Towards a fair marketplace: Counterfactual evaluation of the trade-off between relevance, fairness & satisfaction in recommendation systems. In *Proceedings of the 27th acm international conference on information and knowledge management*. 2243–2251.
- [14] Rishabh Mehrotra and Emine Yilmaz. 2015. Representative & informative query selection for learning to rank using submodular functions. In *Proceedings of the 38th international ACM sigir conference on research and development in information retrieval*. 545–554.
- [15] Lukáš Najbrt and Jana Kapounová. 2014. Categorization of museum visitors as part of system for personalized museum tour. *International Journal of Information and Communication Technologies in Education* 3, 1 (2014), 17–27.
- [16] Yannick Naudet, Bereket Abera Yilma, and Hervé Panetto. 2018. Personalisation in cyber physical and social systems: the case of recommendations in cultural heritage spaces. In *2018 13th International Workshop on Semantic and Social Media Adaptation and Personalization (SMAP)*. IEEE, 75–79.
- [17] Louis Nisiotis, Lyuba Alboul, and Martin Beer. 2020. A Prototype that Fuses Virtual Reality, Robots, and Social Networks to Create a New Cyber-Physical-Social Eco-Society System for Cultural Heritage. *Sustainability* 12, 2 (2020), 645.
- [18] Pierre-Edouard Osche, Sylvain Castagnos, Amedeo Napoli, and Yannick Naudet. 2016. Walk the line: Toward an efficient user model for recommendations in museums. In *2016 11th International Workshop on Semantic and Social Media Adaptation and Personalization (SMAP)*. IEEE, 83–88.
- [19] Pearl Pu, Li Chen, and Rong Hu. 2011. A user-centric evaluation framework for recommender systems. In *Proceedings of the fifth ACM conference on Recommender systems*. 157–164.
- [20] Eirini Eleni Tsiropoulou, George Kousis, Athina Thanou, Ioanna Lykourantzou, and Symeon Papavassiliou. 2018. Quality of experience in cyber-physical social systems based on reinforcement learning and game theory. *Future Internet* 10, 11 (2018), 108.
- [21] Eirini Eleni Tsiropoulou, Athina Thanou, and Symeon Papavassiliou. 2016. Modelling museum visitors' Quality of Experience. In *2016 11th International Workshop on Semantic and Social Media Adaptation and Personalization (SMAP)*. IEEE, 77–82.
- [22] Eirini Eleni Tsiropoulou, Athina Thanou, and Symeon Papavassiliou. 2017. Quality of Experience-based museum touring: A human in the loop approach. *Social Network Analysis and Mining* 7, 1 (2017), 33.
- [23] Willème Verdeaux, Clément Moreau, Nicolas Labroche, and Patrick Marcel. 2020. Causality based explanations in multi-stakeholder recommendations. In *EDBT/ICDT Workshops*.
- [24] Bereket Abera Yilma, Yannick Naudet, and Hervé Panetto. 2018. Introduction to personalisation in cyber-physical-social systems. In *OTM Confederated International Conferences "On the Move to Meaningful Internet Systems"*. Springer, 25–35.
- [25] Bereket Abera Yilma, Yannick Naudet, and Hervé Panetto. 2020. A new paradigm and meta-model for cyber-physical-social systems. In *21st IFAC World Congress, IFAC 2020*. Elsevier.
- [26] Bereket Abera Yilma, Hervé Panetto, and Yannick Naudet. 2019. A Meta-Model of Cyber-Physical-Social System: The CPSS paradigm to support Human-Machine collaboration in Industry 4.0. In *Working Conference on Virtual Enterprises*. Springer, 11–20.
- [27] Jun Jason Zhang, Fei-Yue Wang, Xiao Wang, Gang Xiong, Fenghua Zhu, Yisheng Lv, Jiachen Hou, Shuangshuang Han, Yong Yuan, Qingchun Lu, et al. 2018. Cyber-physical-social systems: The state of the art and perspectives. *IEEE Transactions on Computational Social Systems* 5, 3 (2018), 829–840.
- [28] Yong Zheng. 2019. Multi-stakeholder personalized learning with preference corrections. In *2019 IEEE 19th International Conference on Advanced Learning Technologies (ICALT)*, Vol. 2161. IEEE, 66–70.
- [29] Yong Zheng. 2019. Multi-stakeholder recommendations: case studies, methods and challenges. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 578–579.
- [30] Yong Zheng, Nastaran Ghane, and Milad Sabouri. 2019. Personalized educational learning with multi-stakeholder optimizations. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*. 283–289.