



HAL
open science

Personalised visual art recommendation by learning latent semantic representations

Bereket Abera Yilma, Najib Aghenda, Marcelo Romero, Yannick Naudet,
Hervé Panetto

► **To cite this version:**

Bereket Abera Yilma, Najib Aghenda, Marcelo Romero, Yannick Naudet, Hervé Panetto. Personalised visual art recommendation by learning latent semantic representations. 15th IEEE International Workshop on Semantic and Social Media Adaptation and Personalization, SMAP 2020, Oct 2020, Zakynthos, Greece. pp.1-6, 10.1109/SMAP49528.2020.9248448 . hal-02997497

HAL Id: hal-02997497

<https://hal.science/hal-02997497>

Submitted on 10 Nov 2020

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.

Personalised Visual Art Recommendation by Learning Latent Semantic Representations

Bereket Abera Yilma^{*†}, Najib Aghenda^{*}, Marcelo Romero^{*†} Yannick Naudet^{*} and Hervé Panetto[†]

^{*}Luxembourg Institute of Science and Technology (LIST)

[†]Université de Lorraine, CNRS, CRAN

{bereket.yilma, Najib.aghenda, marcelo.romero, yannick.naudet}@list.lu, herve.panetto@univ-lorraine.fr

Abstract—In Recommender systems, data representation techniques play a great role as they have the power to entangle, hide and reveal explanatory factors embedded within datasets. Hence, they influence the quality of recommendations. Specifically, in Visual Art (VA) recommendations the complexity of the concepts embodied within paintings, makes the task of capturing semantics by machines far from trivial. In VA recommendation, prominent works commonly use manually curated metadata to drive recommendations. Recent works in this domain aim at leveraging visual features extracted using Deep Neural Networks (DNN). However, such data representation approaches are resource demanding and do not have a direct interpretation, hindering user acceptance. To address these limitations, this work introduces an approach for Personalised Recommendation of Visual arts based on learning latent semantic representation of paintings. Specifically, we trained a Latent Dirichlet Allocation (LDA) model on textual descriptions of paintings. Our LDA model manages to successfully uncover non-obvious semantic relationships between paintings whilst being able to offer explainable recommendations. Experimental evaluations demonstrate that our method performs better than other approaches based on the use of visual features extracted using pre-trained Deep Neural Networks.

I. INTRODUCTION

Paintings are important pieces in visual art that bring together complex elements such as drawings, gestures, narration, composition, abstraction, ...etc[2]. These elements carry deeper semantics beyond their usual categorizations (i.e. time period, material, size, color, ...etc.) Paintings are also perceived differently by people as they trigger different emotional and cognitive reflections depending on the background and personality of a person [2]. Personalised recommendation of visual arts often suggest similar contents to those users have already seen or previously indicated that they liked. On the contrary, a recent work proposed by [17] presents an Anti-Recommendation approach called "Art I don't like". This anti-recommender exposes users to a variety of content and suggests artworks that are dissimilar to the ones users selected aiming to allow serendipity and exploration. In comparison to other areas such as movie and music recommendation, Visual art has received little attention [14]. Nevertheless, the huge potential and benefit of personalised recommendations of visual arts has been seen in the works of [14], [15] [16]and [18]. Similarities and relationships among paintings can normally be inferred based on common high level features such as colour, material, style, artist...etc. However, such features are not expressive enough as they can not fully capture abstract concepts that are latent/hidden in the paintings. In order to capture non-obvious contexts by a machine, effective data representation is very crucial [3]. In general prominent works

in visual art recommendation rely on ratings and manually curated metadata (i.e. color, style, mood,...etc) in order to drive recommendations. Recently works such as [1] started to use latent visual features extracted using deep neural networks (DNN) and also use pre-trained models for making visual art recommendations. According to results reported by [14] DNN-based visual features perform better than manually curated metadata. However, these latent visual features do not have a direct interpretation and can not be used to explain recommendations [20]. Hence they hinder user acceptance.

In this work we propose a personalised VA recommendation strategy by learning latent semantic representations. In particular we adopt a Latent Dirichlet Allocation (LDA) based representation learning approach [4]. LDA is a topic analysis model that is known to be successful in the domain of Natural Language Processing (NLP) for uncovering hidden semantic relationships among documents. In this approach paintings are represented by documents containing their detailed textual descriptions. Our LDA model is trained on 2368 painting collections from the National Gallery, London and employed for a personalised recommendation task. The objective of this work is two fold. Primarily it introduces a strategy of leveraging textual data for visual art work recommendations. Consequently it offers a method for driving explainable recommendations since the latent topics uncovered though LDA are directly interpretable. The rest of this paper is organised as follows Section II covers a brief literature review. Section III presents the problem formulation. Section IV illustrates the proposed solution introducing the LDA training and recommendation strategy. Section V presents experimental evaluation of the proposed approach and a comparative analysis against visual feature based recommendations extracted using Deep Neural Networks. Finally Section V presents a concluding discussion.

II. RELATED WORK

The earliest works in Information Retrieval (IR) and Natural language processing (NLP) have been using vector space models to represent documents as a vector of key words [3]. However, such representations offered very limited reduction of description length and had a limited ability to capture inter/inter-document structures. To this end further techniques have been developed to tackle the curse of dimensionality by capturing hidden semantic structures in document modeling. In 2003 an unsupervised generative probabilistic model called Latent Dirichlet Allocation (LDA) was proposed [4]. LDA demonstrated superiority over several other proposed models. Following this Latent variable models became widely accepted




Fig. 1. Latent Dirichlet Allocation plate diagram

strategies to make inference about hidden semantic relationship between variables. Particularly in Recommendation systems LDA has been applied on several tasks such as Online Courses Recommendation [20], personalized hashtag recommendation [21], Scientific Paper Recommendation [22], Similar TV User Grouping and TV Program Recommendation [23]. LDA has been proven to be successful over several recommendation tasks. Hence, in this work we propose a strategy to leverage LDA in the domain of Visual Art recommendation.

III. PROBLEM FORMULATION

Let $P = \{p_1, p_2, \dots, p_m\}$ be the set of paintings in a museum, $D = \{d_1, d_2, \dots, d_m\}$ be the set of documents containing textual descriptions of the paintings in P and $P^u = \{p_1^u, p_2^u, \dots, p_n^u\}$ be the set of paintings a user u has rated where $P^u \in P$. $W^u = \{w_1^u, w_2^u, \dots, w_n^u\}$ are the weights representing the preference of u for the paintings in P^u . The task of personalised painting recommendation is to recommend a ranked list of paintings based on what paintings the user have liked or rated. However, considering the realistic scenario in a museum, access to user preferences is usually very limited. Hence, providing a personalised recommendation service in such settings needs to rely on content based approaches that can work well under limited preference information. In this work we propose a Latent Dirichlet Allocation (LDA) model, to build a content based personalised painting recommendation system. Our trained LDA model is capable of uncovering hidden semantic structures in the paintings to derive explainable recommendations. The LDA training and recommendation strategy is discussed in the next section.

IV. PROPOSED APPROACH

A. LDA Model

Latent Dirichlet allocation (LDA) is an unsupervised generative probabilistic model for collections of text corpora proposed by [4]. The intuition behind LDA is that documents are represented as random mixtures over latent topics, where each topic is characterized by a distribution over words. Many works have shown that LDA is capable of capturing latent semantic information from a collection of documents. Particularly it can identify similar documents by uncovering abstract topics that occur in a collection of documents.[5][6][7][8]. Our painting recommender system is built by creating a document description for each painting in the collection and training an LDA model on the corpus that can uncover semantic structures hidden in the paintings. This refers to everything that is not known a priori and hidden in the paintings. Then the trained LDA model can be queried to recommend most similar paintings for paintings a user have liked.

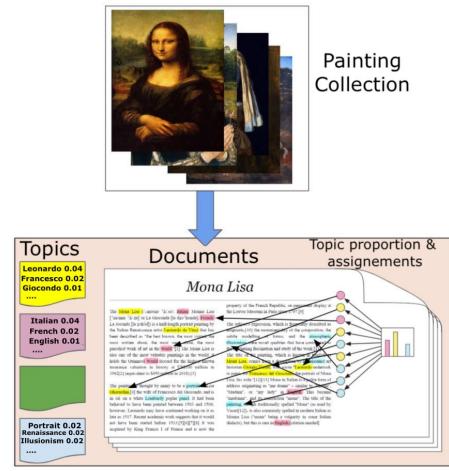


Fig. 2. The intuition behind Painting LDA model

The intuition used in LDA is that each document can be seen as a combination of multiple topics. If we take paintings as an example, paintings can be described as a mixture of several concepts such as religion, nudity, portrait,... etc. In LDA each document is characterized by a predefined set of latent topics. In essence, each document is a distribution of topics and each topic is a distribution of words. This means each word in each document comes from a topic and the topic is selected from a per-document distribution over topics. Prominent words in each latent topic explain the nature of the topic and prominent latent topics related to each document explain the nature of the document (i.e. paintings). For instance, let us assume that latent topics are "religion", "still life", and "landscape". A painting may have the following distribution over the topics : 70% "religion", 10% "still life" and 20% "landscape". Moreover, each topic has a distribution over the words in the vocabulary. For topic "religion", the probability of the word "Saint" would be higher than in the topic "landscape". The LDA model is represented as a probabilistic graphical model in Figure 1. α is the per-document topic distribution, β is the per-topic word distribution, θ is the topic distribution for a document d , φ is the word distribution for the topic K and z the topic for the n^{th} word in the document d and finally, W is a word. In LDA, each topic is a Multinomial distribution over the vocabulary in the collection of documents. To represent a topic, only the top- n words are considered based on their probability. The procedure of building LDA model for a collection of paintings is explained as follows:

- Pre-processing: we construct a collection of documents that contain detailed textual description of the paintings.
- Initialization:
 - 1) We assume there is a defined number of topics k in the collection of documents
 - 2) Attribute a topic to each word W in the collection of documents where $\theta_i \sim Dir(\alpha)$ with $i \in \{1, \dots, M\}$ and $Dir(\alpha)$ is a Dirichlet distribution. This initializes a topic model
- Learning:
 - 1) We assume that the topic assigned to a word

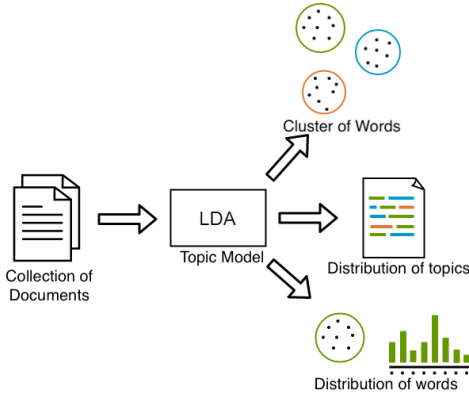


Fig. 3. Topic modeling with LDA

W is wrong but that all the others are correct which consists in computing the conditional probabilities $p(\text{topic } t \mid \text{document } d)$ (probability that the document d is assigned to the topic t) and $p(\text{word } w \mid \text{topic } t)$ (probability that the topic t is assigned to the word w)

- 2) We update the topic of the document which is now the topic that has the highest probability to be assigned to this document ($p(\text{topic } t \mid \text{document } d) \cdot p(\text{word } w \mid \text{topic } t)$)

Assigning the right number of topics as well as the hyperparameters however is not a trivial task. In literature Gibbs sampling algorithm is widely used to estimate parameters of LDA [9][10][11]. In our implementation we used an (optimized version of) collapsed gibbs sampling from MALLET [12]. We refer the reader to the detailed discussion about LDA formulation in [4] and [5]. Figure 2 and 3 illustrate the intuition behind our Painting LDA model construction procedure. As it is illustrated in the figures a collection of documents is used as an input to the LDA algorithm. LDA creates k topics that can be seen as clusters of words. Each document of the collection is represented as a distribution of the topics which are themselves, distributions of words.

B. Personalised Painting Recommendation

Once the LDA model is trained over the entire corpus we get a matrix of documents by topics which expresses latent topic distribution of each painting. From this we can generate an $M \times M$ similarity matrix for all the paintings in the dataset. In our study user preferences towards paintings are modeled by a normalized weight transformed from a 5 star user ratings. (i.e. we assign a weight $w_i^u \in [0, 1]$ for every painting p_i^u a user has rated). The task of personalised Painting recommendation is to recommend most similar paintings to a user based on few of the paintings a user has rated. This is done by expanding user preferences towards unseen paintings and predicting a similarity score for the paintings in the dataset. The predicted score $S(p_i, u)$ for a painting p_i in the dataset, according to the preferences of user u is calculated based on the weighted average similarity score (distance) from all other paintings that

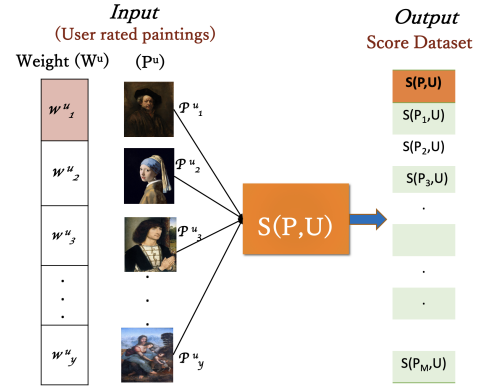


Fig. 4. Painting Scoring Strategy

have been rated by the active user given by:

$$S(p_i, u) = \frac{1}{N} \sum_{j=1}^N w_j * d(p_i, p_j) \quad (1)$$

where $d(p_i, p_j)$ is the similarity between painting p_i and p_j , $p_j \in p^u$ according to the LDA similarity matrix. The summation is taken over all user rated paintings. This scoring strategy is illustrated on Figure 4. Once the scoring procedure is complete the paintings are sorted in a descending order and the first K paintings are selected to generate a recommendation list. Our LDA based personalised painting recommendation procedure is sketched in Algorithm 1.

Algorithm 1 The LDA-BASED PERSONALISED PAINTING RECOMMENDER algorithm; $\mathcal{R} = \{p_1, \dots, p_k\}$ is a set of top k recommendations.

- 1: **procedure** LDA-BASED PERSONALISED PAINTING RECOMMENDER($D, P, P^u, W^u, \mathcal{R}, S(p, u)$)
- 2: ; Initialization
- 3: ; repeat
- 4: **for** d_1, \dots, d_m **do**
- 5: Train LDA model
- 6: **for** p_1, \dots, p_m **do**
- 7: Compute Cosine distance
- 8: **end for**
- 9: **end for**
- 10: $\mathcal{R} \leftarrow \emptyset$
- 11: ; repeat
- 12: **for** p_1, \dots, p_m **do**
- 13: $S(p_i, u) \leftarrow \text{COMPUTE}(\frac{1}{N} \sum_{j=1}^N w_j * d(p_i, p_j))$
- 14: **end for**
- 15: Sort $P \leftarrow \text{DESCENDING}(S(p_i, u))$
- 16: **return** $\{p_1, \dots, p_k\} \leftarrow \text{SELECT TOP } K(\mathcal{R})$
- 17: **end procedure**

V. EXPERIMENT & DISCUSSION

This section presents the experiments and discussions addressing the obtained results. First, we briefly describe the dataset. subsequently, we introduce the metric used to evaluate the proposed topic model. Then, we present an analysis regarding the capability of our approach in terms of explainability. Finally, we describe our method to assess the quality of our approach in the context of a recommendation task.

A. Dataset

In this study we used a dataset containing 2368 paintings from the National Galary of London. Each painting is represented by a set of attributes which are summarized in Table I. For our approach we specifically focus on painting description attribute. These descriptions are provided by museum experts and carry complementary information about the paintings such as stories and narratives that can be exploited to capture the semantic of a painting. These descriptions provide concise information about each painting. For the task of topic modeling using LDA, we decided to enrich the descriptions dataset D by concatenating the paintings descriptions with keyword attributes such as the artist name, the painting title, the technique used, the publication date, the format (landscape, portrait) and the size (small, medium, ...) and also additional information from other sources to better train the model. Hence we generated a second dataset DE enriched with additional textual descriptions and stories related to the paintings, collected from various sources such as Wikipedia and books. To avoid "Garbage in, garbage out" we performed textual pre-processing on both datasets (i.e. removal of punctuation, stop words, bi-grams, Lemmatization) as they do not add any value to the topic models.

Attributes	Description
artist_name	Artist name
painting_title	Painting title
painting_id	Painting identifier
painting_description	Description
publishing_date	Publication date
format	(Landscape, Portrait)
size	(Small, Medium, Large)
technique	(Oil, tempera, ...)

TABLE I. ATTRIBUTES & DESCRIPTIONS OF THE PAINTINGS DATASET

B. Model Evaluation

In topic modeling, *Topic Coherence* is a commonly used technique to evaluate topic models. It is defined as the sum of pairwise similarity scores on the words w_1, \dots, w_n used to describe the topic, usually the top n words by frequency $p(w|t)$ [5], [13].

$$Coherence = \sum_{i < j} Score(w_i, w_j) \quad (2)$$

Ideally, a good model should generate coherent topics. (i.e the higher the coherence score the better the topic model is [13]). In order to identify the optimal number of topics as well as which data set D or DE generates the best topic model, our implementation resorted the topic coherence pipeline from gensim library *CoherenceModel*. Figure 5 shows the evolution of the topic coherence as a function of the number of topics for each of the two dataset. From the analysis presented in figure 5 we can make two observations. Firstly, the data set DE (orange) generally gives slightly better topic coherence score and thus, a better topic model. Secondly, we can observe that with 10 topics we obtained a topic coherence of approximately 0.46 for the data set DE which is the best score for this window. Having too many topics requires more resources as well as time for a result that is not significantly better. Thus, we decided to limit the number of topics to 10 and choose to work with the dataset DE since DE provides slightly better results than D . We believe that it is due to the fact that the

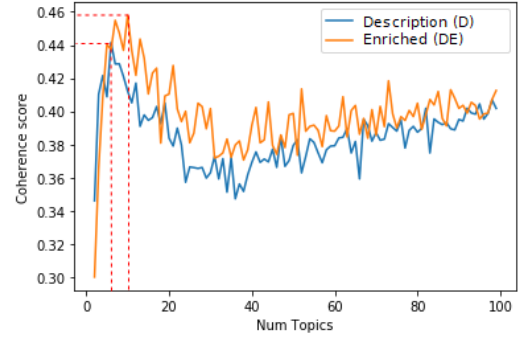


Fig. 5. Comparative topic coherence analysis of the two data sets




Fig. 6. Inter-topic Distance Map in a 2-dimensional space of 10 topics

dataset DE contains more information than D . Moreover, what is it worth noticing that a mixture of keywords and paintings descriptions seems to generate a better topic model than paintings descriptions used alone.

In addition to the evaluation of the topic coherence, we visualize the topics by using a visualization tool called pyLDavis. In figure 6, we see each topic represented by a circle annotated with a digit from 0 to 9. The size of the circle represents the prevalence of a topic, i.e, the popularity of a topic among the paintings. The distance represent the similarity between topics. In fact it is an approximation to the original topic similarity because we are using a 2-dimensional scatter plot to best represent the spatial distribution of the topics. The objective here is to have topics that are overlapping as little as possible. With 10 topics, we can observe that the topics are evenly popular while being sufficiently distinct from each other. This means that the topics are sufficiently different which is what it is intended.

C. Explainability

In figure 7.a and 7.b we present two paintings¹. The painting in 7.b) is the most similar painting to the painting 7.a) based on our LDA topic model. We observe in figure 7.c and 7.d that their topic distribution is very similar and one particular topic stands out: topic 8. The fact that this topic stands out from the others for these paintings implies that the words found in this topic are more likely to be found in the paintings descriptions than the words from the other topics.

When we have a closer look at the topics descriptions in Table II, we can see that the topic 8 is very well defined as it is obvious that there is a coherence between the words that

¹These paintings are available under Creative Commons License

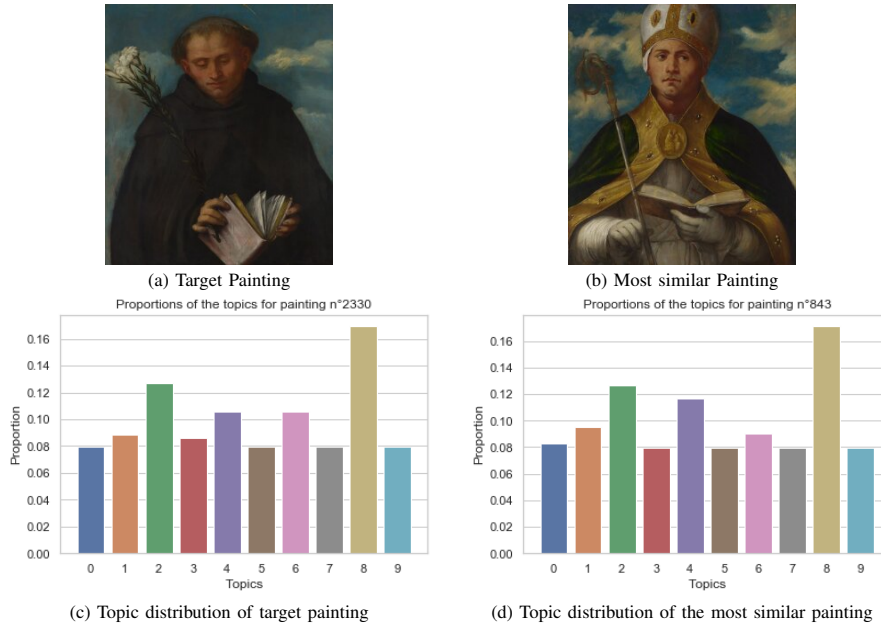


Fig.7. Target painting (Top left), its most similar painting (Top right), Topic distribution of the target (bottom left) and its most similar painting (bottom right).



Fig. 8. Example of prominent concepts shared among recommendations (i.e Explanation)

are used. In fact, topic 8 can be described as the “christian” topic since many of the words in this topic are usually found in christian corpora. When we look at the paintings, it is obvious that there are many references to Christianity and therefore, we can assume that their descriptions contain vocabulary that refers to a religious context. On the other hand, when we look at the descriptions of topics 0 and 5, what is interesting is that the words found in these topics do not seem to share the same semantic at first sight. However, the underlying intuition behind LDA is that words are in the same topics because they are *frequently* found together and therefore, they are used in a same context. Thus, words “Paris” and “flower” found in the topic 5 that do not seem to have any obvious relation are nonetheless in the same topic because many paintings descriptions in our data set describe Paris and flowers together. LDA is able to find relations between words and at a broader extent, paintings that are not always obvious to us but explainable. Therefore for the task of recommendation using LDA we can easily identify prominent topics and reason out shared semantics among recommendations in order to help users better understand recommendations. In our experiment we offered explanations to users as a word cloud of prominent concepts shared among the recommendations matching their personal preferences. Fig. 8 an example of interest word cloud.

Topic 0	Topic 5	Topic 8
LANDSCAPE	PARIS	CHRIST
SCENE	LIGHT	SAINT
VIEW	FORM	ALTARPIECE
17th C LANDSCAPE	COLOUR	PANEL
PEASANT	SKETCH	JESUS
DUTCH	STUDIO	NEW TESTAMENT
LANDSCAPE OIL	FLOWER	EVANGELIST
TREE	COMPOSITION	CROSS
TOWN	STUDY	CHURCH
RIVER	19th CENTURY LANDSCAPE	CRUCIFIXION

TABLE II. DESCRIPTION OF THREE TOPICS

D. Recommendation Quality

In order to evaluate the performance of our algorithm on a recommendation task we performed a controlled experiment with 15 real users. Each user provided a rating for 80 paintings from the dataset to be used for profiling. As a baseline recommender we used visual features extracted using Residual Networks (ResNet-50) [24]. ResNet-50 is a 50-layer deep convolutional neural network trained on more than a million images from the ImageNet database [25]. Thus, it has learned rich feature representations for a wide range of images. We used the pre-trained ResNet-50 to extract latent visual features from our painting dataset. The extracted features are then used to identify similar paintings to user preferred ones through similar scoring mechanism used for LDA. The *User-Centric* evaluation was done through a questionnaire where participants had to express their opinion in a Five point 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 known otherwise.” (Serendipity); “The recommended paintings are diverse” (Diversity). Additionally users were asked if explanations offered by LDA helped them to better understand recommendations. Interestingly all




Fig. 9. Comparison of LDA Vs ResNet-50 in terms of Accuracy, Novelty, Serendipity and Diversity

participants responded “Yes”, this shows that explainable recommendations have a positive impact on user experience. Figure 9 summarizes the average values of user ratings for the two recommendation pipelines. As reported on the figure, LDA achieved significantly higher diversity values (4/5) compared to ResNet-50 (2.3/5). This is due to the fact that the notion of similarity in LDA is directly related to semantically dominant concepts shared among paintings. i.e. LDA also uncovers relationship between paintings that do not necessarily have a resembling visual features. Hence, LDA captures semantic similarities that can not be discovered through latent visual features. This can also justify the higher values of serendipity and Novelty since LDA based recommendation can contain visually diverse but semantically related paintings. In terms of matching user preferences, LDA also performs significantly better (4.1/5) than ResNet-50 (3.1/5).

VI. CONCLUSION

In this work we proposed a personalised Visual Art recommendation strategy based on learning latent semantic representations. In particular we trained an LDA based algorithm which leverages textual descriptions of paintings to uncover latent semantic relationships among paintings. Our Algorithm compares favourably with state of the art VA recommendation method based on visual features extracted using Deep Neural Networks. The latent features extracted with LDA can directly be interpreted by humans. Thus, it enables a personaliser to offer explainable recommendation which is proved to be helpful for users to better understand recommendations. Globally it is evident from the controlled experiment that, LDA outperforms the baseline in all the evaluation matrices. As a future work, further validation with larger set of user groups and also a hybrid approach combining visual features could be a relevant extension.

REFERENCES

- Ruining He, Chen Fang, Zhaowen Wang, and Julian McAuley. 2016. Vista: A Visually, Socially, and Temporally-aware Model for Artistic Recommendation. In Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16). ACM, New York, NY, USA, 309–316
- Mayer, Ralph, The Artist’s Handbook of Materials and Techniques, Third Edition, New York: Viking, 1970, p. 375
- I. Bengio, Y. Courville, A. Vincent, P. Representation learning: A review and new perspectives. IEEE Trans. Pattern Anal. Mach. Intell. 35, 1798–1828 (2013).
- Blei, D. M., Ng, A. Y. Edu, J. B. Latent Dirichlet Allocation Michael I. Jordan. Journal of Machine Learning Research 3, (2003).
- Jelodar, H. et al. Latent Dirichlet allocation (LDA) and topic modeling: models, applications, a survey. Multimedia Tools and Applications 78, (2019).
- Krestel, R., Fankhauser, P., and Nejdl, W. (2009). Latent dirichlet allocation for tag recommendation. Paper presented at the Proceedings of the third ACM conference on Recommender systems.
- Kakkonen, T., Myller, N., Sutinen, E., and Timonen, J. (2005). Comparison of dimension reduction methods for automated essay grading. Natural Language Engineering, 1, 1-16.
- Misra, H., Yvon, F., Cappé, O., and Jose, J. (2011). Text segmentation: A topic modeling perspective. Information Processing Management, 47(4), 528-544.
- S.Ganguly and V.Pudi, ‘Paper2vec: Combining Graph and Text Information for Scientific Paper Representation’, Advances in Information Retrieval 2017, pp. 383-395.
- A.Grover and J.Leskovec, ‘node2vec: Scalable Feature Learning for Networks’, ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 855-864.
- P.Ng. “dna2vec: Consistent vector representations of variable-length k-mers”. In: arXiv preprint arXiv: 1701.06279, 2017.
- Graham, S., Weingart, S., Milligan, I., 2012. Getting Started with Topic Modeling and MALLET. <http://programminghistorian.org/lessons/topic-modeling-and-mallet>.
- David Newman, Jey Han Lau, Karl Grieser, and Timothy Baldwin, Automatic Evaluation of Topic Coherence. NAACL 2010. <https://www.aclweb.org/anthology/N10-1012/>
- Pablo Messina, Vicente Dominguez, Denis Parra, Christoph Trainer, and Alvaro Soto. 2017. Exploring Content-based Artwork Recommendation with Metadata and Visual Features. (2017)
- Abigail R. Esman. 2012. The World’s Strongest Economy? The Global Art Market. <https://www.forbes.com/sites/abigailsman/2012/02/29/the-worlds-strongest-economy-the-global-art-market/>. (2012). [Online; accessed 21-March-2017].
- Ruining He, Chen Fang, Zhaowen Wang, and Julian McAuley. 2016. Vista: A Visually, Socially, and Temporally-aware Model for Artistic Recommendation. In Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16). ACM, New York, NY, USA, 309–316.
- Frost, S., Media, C., Cruz, S., Thomas, M. M., Media, C., Cruz, S., ... Cruz, S. (2012). ART I DON’T LIKE: AN ANTI- RECOMMENDER.
- Naudet, Y., Yilma, B.A., Panetto, H.: Personalisation in Cyber Physical and Social Systems: the Case of Recommendations in Cultural Heritage Spaces. 2018 13th International Workshop on Semantic and Social Media Adaptation and Personalization (SMAP). (2018), doi: 10.1109/smap.2018.8501890.
- Katrien Verbert, Denis Parra, Peter Brusilovsky, and Erik Duval. 2013. Visualizing recommendations to support exploration, transparency and controllability. In Proceedings of the 2013 international conference on Intelligent user interfaces. ACM, 351–362.
- Apaza, R. G., Cervantes, E. V., Quispe, L. C. (2014). Online Courses Recommendation based on LDA. 42–48.
- Zhao, F., Zhu, Y., Jin, H., Yang, L. T. (2016). A personalized hashtag recommendation approach using LDA-based topic model in microblog environment. Future Generation Computer Systems, 65, 196–206. <https://doi.org/10.1016/j.future.2015.10.012>
- Amami, M., Pasi, G., Stella, F., Faiz, R. (2016). An LDA-Based Approach to Scientific Paper Recommendation. 2, 200–210. <https://doi.org/10.1007/978-3-319-41754-7>
- Pyo, S., Kim, E., Kim, M., Member, S. (2015). LDA-Based Unified Topic Modeling for Similar TV User Grouping and TV Program Recommendation. 45(8), 1476–1490.
- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In Proceedings of CVPR, pages 770–778, 2016. arxiv.org/abs/1512.03385.
- ImageNet. <http://www.image-net.org>