ART I DON'T LIKE: AN ANTI-RECOMMENDER SYSTEM FOR VISUAL ART

Sarah Frost
Computational Media
UC Santa Cruz
Santa Cruz, California, U.S.
smfrost@ucsc.edu

Manu Mathew Thomas
Computational Media
UC Santa Cruz
Santa Cruz, California, U.S.
mthomas6@ucsc.edu

Angus G. Forbes
Computational Media
UC Santa Cruz
Santa Cruz, California, U.S.
angus@ucsc.edu

Abstract

Recommender systems show users' products, songs, and political points of view that are similar to content that they have already seen or previously indicated that they like. While these methods are useful for generating sales and clicks, they are not necessarily successful at exposing users to disparate content. *Art I Don't Like* is a Web-based interactive art experience that provides personalized content to users and emphasizes the introduction of disparate content. We suggest a new "anti-recommender" system that provides content that is aesthetically related in terms of low-level features but challenges the implied conceptual frameworks indicated by initial user selections. Furthermore, we demonstrate an application of recommender technologies to visual art in an effort to expose users to a broad range of art genres. We present details of a prototype implementation trained on a subset of the WikiArt dataset, consisting of 52,000 images of art from 14th- to 20th-century European painters, along with feedback from users. *Art I Don't Like* is on the web at http://www.artidontlike.com.

Keywords: recommender systems, visual art, personalization, deep learning

Introduction

This work is motivated by a concern about current recommender system technology and personalization algorithms. As people spend a considerable amount of time on the Internet, their views on politics and social issues are shaped by the information they consume online. Internet users may not realize that algorithms have been developed to give them personalized content. By removing access to opposing viewpoints, personalization can lead to *filter bubbles*—a term coined by Eli Pariser to describe a type of intellectual isolation that occurs as a result of personalization algorithms. These algorithms present information to users based on previously viewed content and content viewed by similar users. Users have little exposure to contradicting viewpoints and have become unknowingly trapped in a digital bubble (Pariser, 2011). The recommender systems community is aware of this tendency, and researchers have explored ways to mitigate content isolation. Some solutions include improving the transparency of these systems (Cosley, 2003; Cacheda, 2011; Forbes, 2013), giving the user control over the settings of the personalization algorithms (Hirschmeier, 2018), and using new technologies to make recommender-system technology more understandable (Forbes, 2012). However, these solutions do not necessarily challenge the core underlying assumption of such methodologies, which is that users want to be presented with content that is as similar as possible to content they have indicated that they like. In addition, these solutions are not always clearly explained to the public, making them less aware of the impact that recommender systems have on their Internet experience.

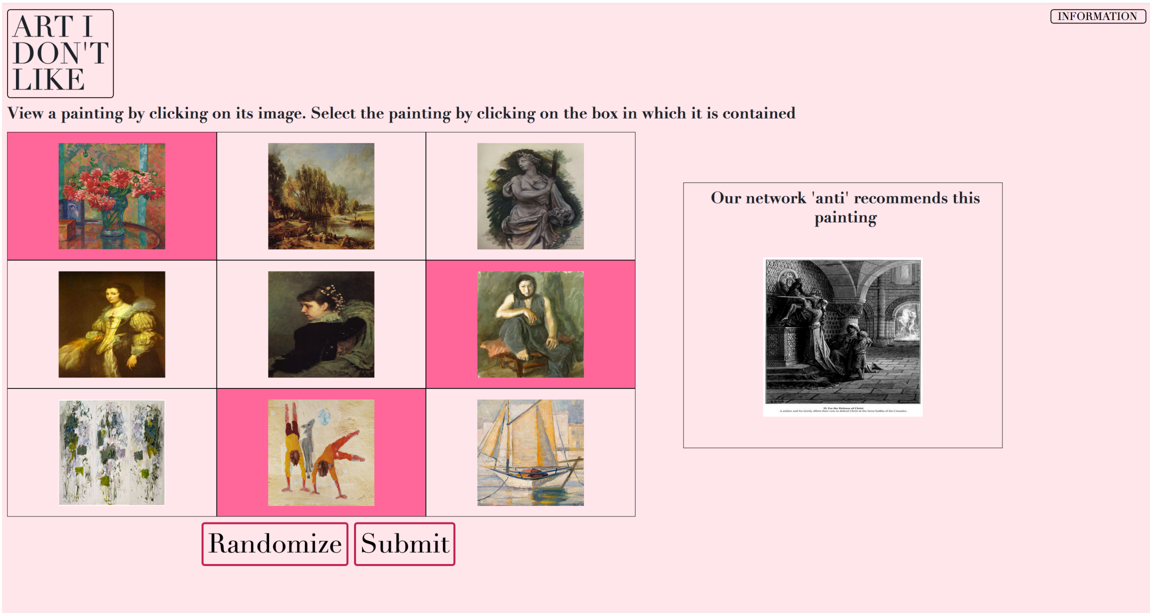

Figure 1: Users are asked to select the paintings they like from a group of nine paintings. Our anti-recommender system scans each of them and recommends a painting which is dissimilar to the selections made by the user.

Art I Don't Like (artidontlike.com) is a Web-based interactive art experience that provides personalized content to users. *Art I Don't Like* suggests content by prompting users to pick artworks that they find visually appealing. This anti-recommender system then returns an artwork that is *dissimilar* to the selected content. This system will expose users to a broader range of art, and shed light on how recommendations are made. This project gives users a digital space to view and interact with art that they have not specifically searched for and probably would not search for. In this way, the recommender system concept has been expanded to allow for serendipity and exploration. An example of output based on user input is shown in Figure 1.

Increasingly, museums are using technology to digitize collections, reach new patrons, and disseminate information about art and cultural objects (Gentile, 2011; Risseeuw, 2016; Petrelli, 2017). This project offers another way that museums can utilize emerging technologies for visual art dissemination. This project also responds to people's desire for personalized art experiences. Many museums want to provide users with a digital, online space to view and interact with art. Users can access the site from their homes on a personal computer or via a digital kiosk or installation in a museum or gallery space.

Related Work

Recommender systems are usually classified into three types: content-based, collaborative, and hybrid (Su, 2009). Content-based recommender systems show the user items similar to the ones the user rated highly in the past (Adomavicius, 2005). For example, a movie recommender system classifies movies based on genre, actors, general rating, and other characteristics. It then looks for other movies with classifications similar to movies the user has viewed and rated highly in the past. Over time, content-based systems learn about user taste and preferences either implicitly or explicitly (Ricci, 2015). On the other hand, collaborative recommender systems show the user items that people with similar tastes and preferences have liked in the past. Amazon's system for recommending books is an example of a collaborative recommender system (Su, 2009). Finally, hybrid systems utilize a mix of both content-based and collaborative system technology (Ricci, 2015). Hybrid systems can combine content-based and collaborative recommender systems in several ways. They can apply both methods separately and merge the predictions from both systems or incorporate some characteristics of one approach into the other (Ricci, 2015). This combination is used to avoid some of the limitations of content-based and collaborative recommender systems.

The three categories of recommender systems have several known drawbacks and limitations. For collaborative recommender systems, mapping similar users can be ineffective when providing recommendations. Users with similar opinions about many items and, therefore, considered similar, can totally disagree about others (Cacheda, 2011). Similarly, capturing users' opinion as a number on a scale can be ineffective. Many recommenders represent a user's opinion about an item as a single number on a rating scale. These scales vary widely in their granularity (Cacheda, 2011). Conversely, content-based recommender systems require a great deal of feedback from users and a significant level of user involvement (Cosley, 2003). Users can receive targeted content that is not balanced and skews towards content that companies want users to see (Hirschmeier, 2018). An issue with all recommender systems is the potential lack of transparency for users. The way in which recommender systems connect similar items or similar users can be unclear to users. Users might not be aware that Facebook is presenting different information to them based on personalization algorithms (Zuiderveen, 2016).

Methodology

Data Set and Network Architecture

Although this system can be trained with paintings from any time period or region, our initial dataset consists of 52,000 paintings from WikiArt. This dataset was extracted by Tan, et al. (2018) as part of a project to improve conditional image synthesis. They categorized the art by artist and genre—23 artists and 27 art genres are represented. We transformed the WikiArt dataset to suit our classifiers. We used two helper scripts to group the paintings based on artist and genre, then we created one-hot encodings for artists and genres. We instantiated two neural networks using MobileNet, one for classifying the artist and one for classifying the genre. MobileNet is built on TensorFlow, an open source library for numerical computation developed by Google (Abadi, 2016). MobileNet is a convolutional neural network used for classification trained on a huge image database called ImageNet (Sandler, 2018). This allowed us to apply transfer learning. We began with a model that had been trained on another problem and retrained the last few layers for our specific application. MobileNet has eleven layers, each with a 3 x 3 filter. The first layer of each sequence has a stride 5, and all others use stride 1. The classifiers return a confidence score for each artist and genre. Figure 2 is a digital image of the input painting, “A Bank of Canal” by Pablo Picasso, which we resized to a 128 x 128 pixel .jpeg file before classifying it using our artist and genre networks. Figure 3 (left) shows the results of the artist neural network after being run with this painting. The classifier is 99.071% confident that this piece was painted by Picasso, .724% confident that this piece was painted by Kustodiev, .147% confident that this piece was painted by Konchalovsky, .024% confident that this piece was painted by Renoir, .024% confident that this piece was painted by Roerich, and .005% confident that this piece was painted by Chagall or van Gogh. Figure 3 (right) shows the results of the genre neural network after being run with the painting.

Figure 2: *A Bank of Canal* by Picasso is an example from our training set.

```
pablo picasso (score=0.99071)
boris kustodiev (score=0.00724)
pyotr konchalovsky (score=0.00147)
pierre auguste renoir (score=0.00024)
nicholas roerich (score=0.00024)
marc chagall (score=0.00005)
vincent van gogh (score=0.00005)
claude monet (score=0.00000)
salvador dali (score=0.00000)
ilya repin (score=0.00000)
raphael kirchner (score=0.00000)
eugene boudin (score=0.00000)
paul cezanne (score=0.00000)
martiros saryan (score=0.00000)
camille pissarro (score=0.00000)
edgar degas (score=0.00000)
ivan aivazovsky (score=0.00000)
rembrandt (score=0.00000)
albrecht durer (score=0.00000)
john singer sargent (score=0.00000)
ivan shishkin (score=0.00000)
childe hassam (score=0.00000)
gustave dore (score=0.00000)
```

```
naive art primitivism (score=0.74915)
expressionism (score=0.14951)
cubism (score=0.04539)
fauvism (score=0.02271)
art nouveau modern (score=0.01299)
post impressionism (score=0.00714)
symbolism (score=0.00478)
pop art (score=0.00217)
ukiyo e (score=0.00195)
synthetic cubism (score=0.00138)
abstract expressionism (score=0.00082)
pointillism (score=0.00057)
realism (score=0.00050)
impressionism (score=0.00034)
contemporary realism (score=0.00023)
northern renaissance (score=0.00016)
early renaissance (score=0.00014)
high renaissance (score=0.00003)
new realism (score=0.00002)
romanticism (score=0.00001)
baroque (score=0.00000)
mannerism late renaissance (score=0.00000)
minimalism (score=0.00000)
rococo (score=0.00000)
analytical cubism (score=0.00000)
color field painting (score=0.00000)
action painting (score=0.00000)
```

Figure 3: Our artist classifier gives a confidence score for the input painting with regard to the artists on which we trained (left). We also trained a genre classifier which gives confidence scores for an input painting with regard to the genres on which we trained (right). We leverage the output from both of these classifiers to select the dissimilar output in the *Art I Don't Like* application.

Anti-Recommender System

The networks are used to provide confidence scores for the artists and genres for which they were trained. We are interested in the lowest non-zero scores, which we use to determine dissimilarity. Our system selects artists and genres that have received less than a .01% confidence score. For example, for Picasso's *A Bank of Canal*, our artist network was <.01% confident it was painted by Chagall or van Gogh, and similarly, our genre network was <.01% confident that it was a work belonging to the genres of Romanticism, New Realism, or High Renaissance, as shown in Figure 4. Our anti-recommender system will return a piece of art that is categorized as van Gogh or Chagall, and Romanticism, New Realism, or High Renaissance. Because van Gogh and Chagall do not have paintings in the genres of Romanticism, New Realism, or High Renaissance, an artwork from any artist from the genres Romanticism, New Realism, or High Renaissance will be returned to the user.

User Interaction

The landing page of the website shows nine pieces of art from our dataset. In this grid, the pieces have been resized for continuity, but the user can click on each piece to see the original size, as shown in Figure 4. The user is prompted to select the pieces of art that he/she likes, and submit their choices. Our network classifies each piece of art and determines the most dissimilar piece in the dataset. After the user is presented with the first piece of dissimilar art, he/she has the option of choosing more pieces of art to seed the process again. Our users can visit the "information" section to understand how our system made those recommendations are made. They can learn about our recommender system, and how we've used neural networks to classify art. We hope that the process of showing users new pieces of art, and exposing them to other users' opinions, will challenge their opinions of the type of art that they like.

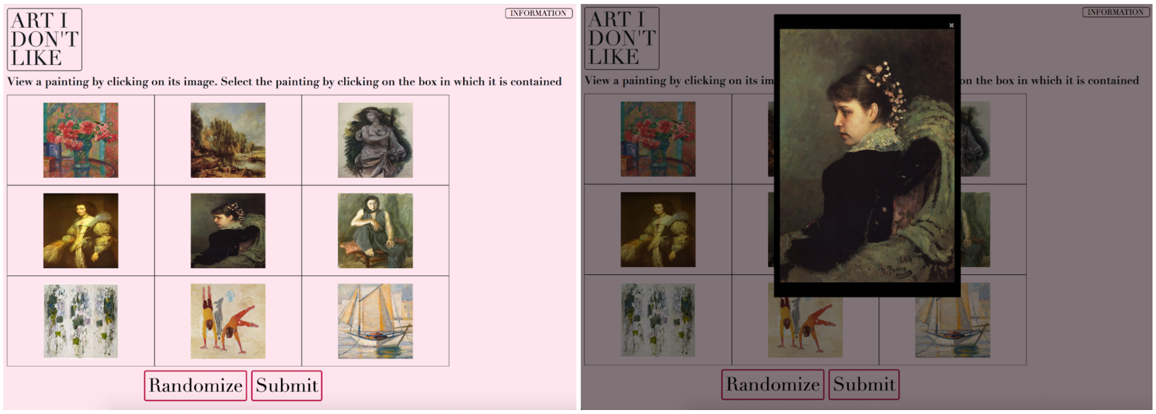

Figure 4: The interface of our system shows nine paintings in a grid format. Users can click on the images to see them in their original resolution.

Discussion and Further Work

We conducted a series of interviews with three arts experts to obtain initial feedback about *Art I Don't Like* and to generate ideas for future iterations of the project. Our experts were Stacy Kamehiro and Kyle Parry, faculty of the History of Visual Arts and Culture department at University of California, Santa Cruz, and Nina Simon, head of the Santa Cruz Museum of Art and History. Each of the experts appreciated our efforts to create an engaging interactive application to increase audience awareness of art history. Kamehiro found *Art I Don't Like* to be a promising tool for fostering learning about lesser known artists and periods of art history. She also suggested that we present users with paintings that range from most to least similar based on the paintings that the user chose that they liked. Parry gave helpful feedback on the some of the design choices of our current implementation. He also acknowledged that while our initial dataset was adequate for the developing the neural networks used in the project, we should consider including a wider range of artworks. He further recommended that we retrain our neural networks when introducing new data to ensure that the classifiers are not biased. Nina Simon was very interested in the underlying idea of using visual art as a way to encourage users to broaden their opinions, and encouraged us to study the use of *Art I Don't Like*in real world scenarios in order to refine the project.

Additionally, we presented a post-use survey to graduate students in the Computational Media department at University of California, Santa Cruz. We received eight responses to the open-ended questions which included "How would you use this website?," and "What are your initial thoughts about the website?" Representative feedback included the following comments:

- *I like the idea of a art website that exposes me to art that I might not normally encounter, because I wouldn't seek it out/identify with similar things.*

- *The idea itself is cool, but it's also valuable to notice that the system itself forces the user to become aware of "echo chambers" in their own interactions. The information page is very useful in that sense. I would hope that general users would take the time to read through it and not just stay on the anti-recommender page.*
- *I like it! I would have my friends do it too and see what they get—it could generate discussions between us.*
- *I think having more options to choose from would be good—I'm not sure if the current sample size is enough to determine my taste in art.*

The current prototype of *Art I Don't Like* is designed to be used on a personal computer via a web browser. However, it can also be used in museums and gallery spaces. For example, it can be integrated into a welcome kiosk in a museum atrium. To do this, we can upload the collections on display at the museum. The user can choose pieces that they have seen and like during past visits, and the system will return pieces of art that the user might not be familiar with. Along with this, we will develop the functionality to map out the locations of these pieces of art in the museum and send the map to the user. This will create a personalized museum experience for the users who want an opportunity to view new art and learn about their own art preferences.

For future iterations of *Art I Don't Like*, we plan to continue exploring approaches to introducing lesser-known artists and genres to users. For instance, we will investigate alternative ways define our dissimilarity metric to retrieve additional types of artworks. Art from dissimilar genres can at the same time be similar in other ways, such as sharing subject matter or compositional arrangement. Currently, we make the assumption that users will have less exposure to art that is of a different genre than art that they profess to like, and we will explore additional ways to generate profiles of user interests. As noted by our experts, a limitation of our initial implementation of *Art I Don't Like* concerns our default dataset of artworks. The dataset consists entirely of paintings from 14th-to-20th century European painters. Of course, this is not a comprehensive database of artworks, and we will extend our database to include artworks from many cultures and to include contemporary artworks. Additionally, user interaction does not currently have any impact on the results of the recommender system for later users. We plan to incorporate user feedback in training the neural network to understand similar and dissimilar artworks.

Personalization algorithms and recommender systems connect users and the information, products, or experiences they seek. We present *Art I Don't Like* as an example of a recommender system that maximizes the distances between objects and pushes toward the boundaries of similarity, which emphasizes the need for serendipity and diversity. The system can be accessed online (http://www.artidontlike.com), and source code and data is also available. (https://github.com/CreativeCodingLab/ArtIDontLike.)

Acknowledgements

We would like to thank Daniel Shapiro of the Computational Media Department at UC Santa Cruz for his support and feedback. We also thank our panel of experts—Stacy Kamehiro and Kyle Parry of the History of Art and Visual Culture Department at UC Santa Cruz, and Nina Simon of the Santa Cruz Museum of Art and History—for providing feedback on the project. Finally, we are grateful to Dan Frost, Professor Emeritus of Teaching at UC Irvine's Department of Informatics, for his constructive criticism throughout the development of this project.

References

Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D. G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., & Zheng, X. (2016). "TensorFlow: A system for large-scale machine learning." 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16). Consulted December 4, 2018. Available at: https://www.usenix.org/system/files/conference/osdi16/osdi16-abadi.pdf

Adomavicius, G. & Tuzhilin, A. (2005). "Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions." *IEEE Transactions on Knowledge and Data Engineering*, Vol. 17, No. 6.

Cacheda, F., Carneiro, V., Fernandez, D., & Formoso, (2011). "Comparison of collaborative filtering algorithms: Limitations of current techniques and proposals for scalable, high-performance recommender systems." *ACM Trans*. Web 5, 1, Article 2. Consulted January 2, 2019. Available at: http://doi.acm.org/10.1145/1921591.1921593

Cosley, D., Lam, S. K., Albert, I., Konstan, J., & Riedl, J. (2003). "Is Seeing Believing? Recommender Interfaces Affect Users' Opinions." *CHI* 2003, Ft. Lauderdale, Florida, USA. ACM 1-58113-630-7/03/0004.

Forbes, A.G., Savage, S., & Höllerer, T. (2012). "Visualizing and verifying directed social queries." *Proceedings of the IEEE Interactive Visual Text Analytics (TextVis)*. Seattle, Washington, USA. Consulted January 2, 2019. at: https://creativecoding.soe.ucsc.edu/pdfs/Forbes_DSQ_VisWeek2012.pdf

Forbes, A.G., Surdeanu, M., Jansen, P., & Carrington, J.M. (2013). "Transmitting narrative: An interactive shift-summarization tool for improving nurse communication." *Proceedings of the IEEE Interactive Visual Text Analytics (TextVis)*. Atlanta, Georgia, USA. Consulted January 2, 2019. at: https://creativecoding.soe.ucsc.edu/pdfs/Forbes_TransmittingNarrative_VIS_INTERACTIVE_TEXT_2013.pdf

Gentile, A., Andolina, S., Massara, A., Pirrone, D., Russo, G., Santangelo, A., Trumello, E., & Sorce, S. (2011). "A Multichannel Information System to Build and Deliver Rich User-Experiences in Exhibits and Museums." *International Conference on Broadband, Wireless Computing, Communication and Applications*, BWCCA 2011, Barcelona, Spain, October 26-28, 2011. 57-64. 10.1109/BWCCA.2011.14.

Hirschmeier, S. & Beule, V. (2018). "Compliance of Personalized Radio with Public-Service Remits." In *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization* (UMAP '18). ACM, New York, NY, USA, 15-21. Consulted December 20, 2018. DOI: Available at: https://doi.org/10.1145/3213586.3226204

Metz, R. (2018). “An ex-Google engineer is scraping YouTube to pop our filter bubbles.” Last updated April 12, 2018. Consulted January 20, 2019. Available at: https://www.technologyreview.com/s/610760/an-ex-google-engineer-is-scraping-youtube-to-pop-our-filter-bubbles/

Pariser, E. (2011). “The Filter Bubble: What The Internet Is Hiding From You.” London: Penguin UK.

Petrelli, D., Marshall, M. T., O’brien, S., Mcentaggart, P., & Gwilt, I. (2017). “Tangible data souvenirs as a bridge between a physical museum visit and online digital experience.” *Personal Ubiquitous Comput*. 21, 2 (April 2017), 281-295. Consulted February 1, 2019. DOI: Available at: https://doi.org/10.1007/s00779-016-0993-x

Ricci, F., Rokach, L., & Shapira, B., eds. (2015). *Recommender Systems Handbook*. New York: Springer Science+Business Media.

Rinehart, W. (2017). “The Election of 2016 and the Filter Bubble Thesis in 2017.” Consulted January 20, 2019. Available at: https://medium.com/@willrinehart/the-election-of-2016-and-the-filter-bubble-thesis-in-2017-51cd7520ed7

Risseeuw, M., Cavada, D., Not, E., Zancanaro, M., Marshall, M. T., Petrelli, D., & Kubitza, T. (2016). “Authoring Augmented Digital Experiences in Museums.” In P. Buono, R. Lanzilotti, & M. Matera (Eds.) *Proceedings of the International Working Conference on Advanced Visual Interfaces* (AVI ’16), ACM, New York, NY, USA, 340-341. DOI: Available at: https://doi.org/10.1145/2909132.2926064

Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L. (2018). “MobileNetV2: Inverted Residuals and Linear Bottlenecks.” Consulted January 15, 2019. Available at: https://arxiv.org/pdf/1801.04381.pdf

Su, X. & Khoshgoftaar, T. M. (2009). “A Survey of Collaborative Filtering Techniques. Advances in Artificial Intelligence,” vol. 2009, Article ID 421425. Available at: https://doi.org/10.1155/2009/421425

Tam, S. (2018). “Pop goes the filter bubble?” Last updated April 4, 2018. Consulted January 20, 2019. Available at: https://slate.com/technology/2018/04/personalization-algorithms-could-actually-help-pop-filter-bubbles.html

Tan, W. R., Chan, C. S., Aguirre, H. E., & Tanaka, K. (2018). “Improved ArtGAN for Conditional Synthesis of Natural Image and Artwork.” arXiv: 1708.09533.

Zuiderveen Borgesius, F. J., Trilling, D., Möller, J., Bodó, B., de Vreese, C. H., & Helberger, N. (2016). “Should we worry about filter bubbles?.” *Internet Policy Review*, 5(1). DOI: Available at: https://doi.org/10.14763/2016.1.401