UNIVERSITY OF CALIFORNIA
SANTA CRUZ

MACHINE LEARNING FOR SOCIAL GOOD

A thesis submitted in partial satisfaction of the
requirements for the degree of

Master of Science

in

COMPUTATIONAL MEDIA

by

Sarah Frost

September 2020

The Thesis of Sarah Frost
is approved:

Angus G. Forbes, Chair

Katherine Isbister

Quentin Williams
Interim Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures **iv**

Abstract **v**

1 Introduction **1**

2 Art I Don't Like **4**

2.1 Introduction 4
2.2 Related Work 6
2.3 Method 8
2.4 Results 11
2.5 Re-Design 12

3 CompostNet **19**

3.1 Introduction 19
3.1.1 Problem Definition 20
3.1.2 Motivation 21
3.1.3 Terminology 22
3.2 Related Work 23
3.2.1 Image Classification 23
3.2.2 Waste Classification 23
3.3 Method 24
3.3.1 Dataset 24
3.3.2 Models 25
3.4 Results and Evaluation 27
3.4.1 Models 27
3.4.2 Deployment 29
3.4.3 Future Work 31

4 Conclusion **32**

Bibliography **34**

List of Figures

2.1 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. 6
2.2 A Bank of Canal by Picasso is an example from our training set. 9
2.3 We leverage the output from both our artist and genre classifiers to select the dissimilar output in the Art I Don't Like application. 10
2.4 The updated interface of our system shows one painting. Users can click on the images to see them in their original resolution. 14
2.5 After the user has selected their preferred number of paintings, they can click "Find Art" to see the dissimilar image returned by the system. . . 14

3.1 Waste receptacles at the National Institutes of Health in Bethesda, MD. 21
3.2 Three images from our compostable class. 24
3.3 Architecture for CompostNet - Version A. 26
3.4 A test image and the confidence scores returned by Version A model. . . 28
3.5 Test and train accuracy for CompostNet - Version A. 28
3.6 Test and train accuracy for CompostNet - Version B. 29
3.7 The iOS app after the system has classified an image. 30

Abstract

Machine Learning for Social Good

by

Sarah Frost

In this thesis, I present a portfolio of projects centered around the theme of Machine Learning for Social Good (ML4SG). I present two fully functional prototypes for machine-learning based systems targeting societal challenges. The first project is Art I Don't Like, an anti-recommender system for visual art that I built in Fall 2018 and Winter 2019. This anti-recommender system encourages users to reflect on their relationship to personalization algorithms on the internet through their art preferences. The second project is CompostNet, an image classification model for personal meal waste. I developed CompostNet in CMPS 290T with Rakshit Agrawal in Spring 2019. CompostNet is a mobile application that includes an image classifier for meal waste, to help individuals divert their personal waste from landfills. In the introduction I discuss some categories ML4SG projects share - method, intended users, and phase of deployment. I reflect on my projects to determine that Art I Don't Like is a Social good-first, individual-oriented, project at small-scale deployment, while CompostNet is a Social Good-First, Individual-oriented, undeployed project.

Chapter 1

Introduction

There has been an increase in conversations and work about ML4SG for the last several years. After Machine Learning became widely available in the early 2000s, it was several years before the technology improved to the point that individuals, small groups, and universities could apply ML to various problems. In recent years, the terms "AI for Social Good" and "Machine Learning for Social Good" have become part of initiatives from large companies like Microsoft, Google, and McKinsey [23], and the subject of reports and workshops at leading AI conferences.

The process of developing a ML4SG project can be broken down into four steps: Identification, in which the societal challenges are identified and analyzed, Definition, in which the problem is defined by a combination of learning tasks, Solution, in which a ML method is determined, and Deployment, in which the system is developed and delivered [13]. ML4SG projects can be classified by method, intended users, and phase of deployment.

Two dichotomous methods for conducting ML4SG work have been identified by Dr. Agrawal [13]. One method involves identifying a social good need, and a stakeholder population, and developing the ML solution necessary to address the problem: "Social Good First". The other method involves using a machine learning method that already exists, and finding a Social Good context in which the method can be applied: "Machine Learning First".

ML4SG projects can also be organized by the intended audience or users of the ML solution. Some projects are designed to be applied by a city or municipality, other solutions are designed to be used by an individual, and other solutions are designed to be deployed by a group of people or a community, such as doctors. I organized the intended users or beneficiaries of each ML4SG project into three groups: Individual, Group, and City or Country.

Finally, ML4SG projects can be categorized by phase of deployment. Some machine learning solutions that can be used to address social good problems, but have not been deployed, or have only been deployed in a small case study [35] [9] [27] [47]. Other projects have been operating for years, and have been deployed in multiple geographic locations [21] [61]. A core component of ML4SG projects is deployment, without that, an argument can be made that there is no social good being done, and the project cannot truly be called an ML4SG project.

This thesis presents a portfolio of projects centered around the theme of Machine Learning for Social Good (ML4SG). I present two fully functional prototypes for machine-learning based systems targeting societal challenges. The first project is Art

I Don't Like, an anti-recommender system for visual art that I built in Fall 2018 and Winter 2019. The second project is CompostNet, an image classification model for personal meal waste. I developed CompostNet in CMPS 290T with Rakshit Agrawal in Spring 2019.

The text of this thesis includes texts of the following two previously published papers, Art I Don't Like: An Anti-Recommender System for Visual Art and CompostNet. Art I Don't Like was written by Sarah Frost, Manu Mathew Thomas, and Angus G. Forbes, and presented at Museums and the Web in Boston, Massachusetts in April 2019. CompostNet was written by Sarah Frost, Bryan Tor, Rakshit Agrawal, and Angus G. Forbes and presented at IEEE Global Humanitarian Conference in Seattle, Washington in October 2019. Dr. Forbes directed and supervised the research which forms the basis for this thesis.

Chapter 2

Art I Don't Like

In this section I discuss Art I Don't Like: An Anti-Recommender System for Visual Art, which I presented at Museums and the Web in April 2019 in Boston. I developed this work with Dr. Angus Forbes and Manu Mathew Thomas.

2.1 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 [50]. 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 [50]. 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 [24], [20], [29], giving the user control over the settings of the personalization algorithms [36], and using new technologies to make recommender-system technology more understandable [30]. 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 users less aware of the impact that recommender systems have on their Internet experience.

Art I Don't Like (http://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, recommender system concept has been expanded to allow for serendipity and exploration. An example of output based on user input is shown in figure 2.1.

Increasingly, museums are using technology to digitize collections, reach new patrons, and disseminate information about art and cultural objects [33] [55] [51]. This

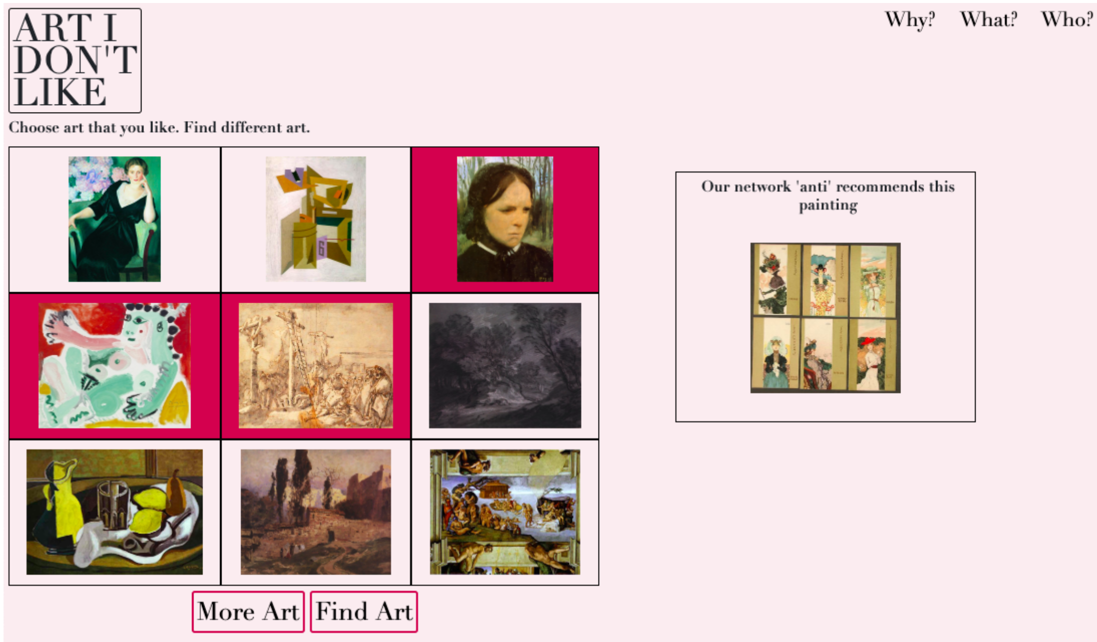

Figure 2.1: Users are asked to select the paintings they like from a group of nine paintings. Our anti-recommender system scans each of them and recommends a painting which is dissimilar to the selections made by the user.

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.

2.2 Related Work

Recommender systems are usually classified into three types: content-based, collaborative, and hybrid [63]. Content-based recommender systems show the user items similar to the ones the user rated highly in the past [12]. 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 [53]. 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 [63]. Finally, hybrid systems utilize a mix of both content-based and collaborative system technology [53]. 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 [53]. 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 [20]. 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 [20]. Conversely, content-based recommender systems require a great deal of feedback from users and a significant level of user involvement [24]. Users can receive targeted content that is not balanced and skews towards content that companies want users to see [36]. 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 [18].

2.3 Method

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. as part of a project to improve conditional image synthesis [66]. 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. 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 [59]. 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 2.2 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 2.2: A Bank of Canal by Picasso is an example from our training set.

Figure 2.3 (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 2.3 (right) shows the results of the genre neural network after being run with the painting.

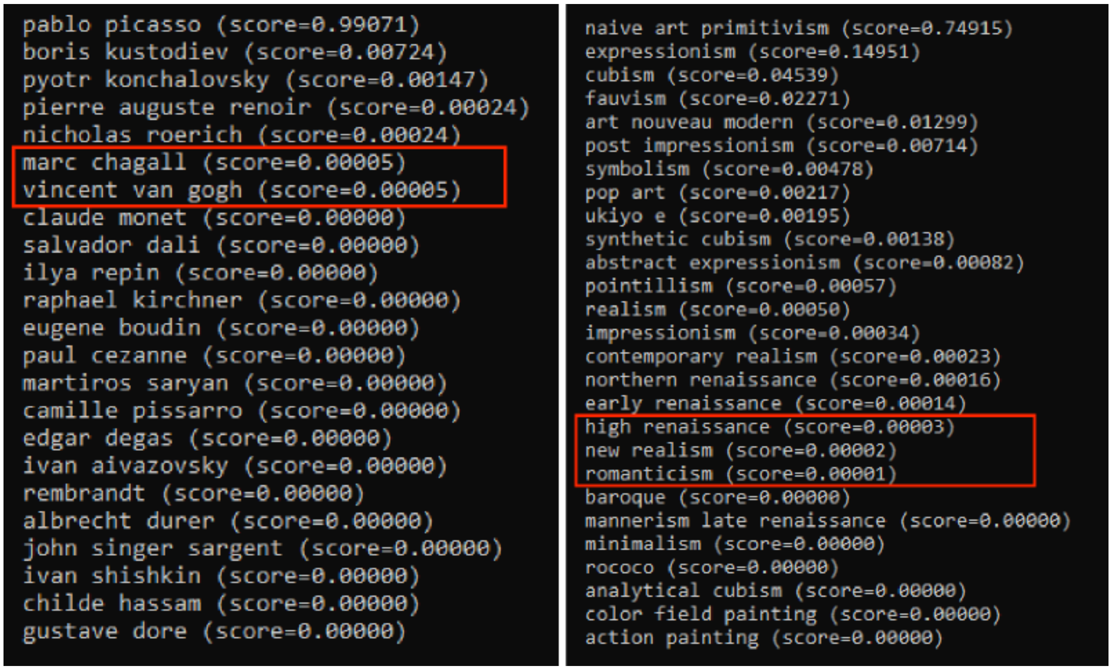

Figure 2.3: We leverage the output from both our artist and genre classifiers to select the dissimilar output in the Art I Don't Like application.

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 3. 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.

2.4 Results

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. 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 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.

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.”

2.5 Re-Design

We received many pieces of positive feedback, but users were also confused about how to select pieces of art, and the purpose of the “Randomize” button. With all this in mind, we decided to create and implement a new website design that showed the user one piece of art at a time. This design would also be streamlined and clean. After implementing this new design, we wanted to get feedback from users and learn about the emotional affect of the website. First we created Adobe XD wireframes with the help of Alejandro Calderon Lopez and Anna Sofia Frattini, two visiting students in the Creative Coding lab. We then modified the code to implement the new design on my local machine. After several iterations of the design, I finished programming and moved to develop the user test with the help of Aurora Alparaz, a Cognitive Science undergraduate working in the Creative Coding lab. We then conducted the user tests and analyzed our findings.

To implement the new design, we moved the code from an AWS EC2 instance to my laptop. The first step was to create a smaller dataset. Our initial dataset consisted of 52,000 paintings from WikiArt. Our mini dataset consisted of 300 images, representing the 23 artists and 27 genres. Our new algorithm for choosing art allows the user to choose the amount of art they like, with a minimum of three paintings. We decided not to cap the amount of pieces of art that a user could look at before clicking the "Find Art" button.

The new design also includes information for the resulting piece of art. The images from the WikiArt dataset were named with the artist, name of the work, and sometimes, date of completion. For example, Edgar Degas's 1877 work, Women on a Cafe Terrace in the Evening, has this image name: edgar-degas_women-on-a-cafe-terrace-in-the-evening-1877.jpg. Using JavaScript and regular expressions, we isolated the pieces of information and displayed them for the user to see. The redesigned site is now live at http://www.artidontlike.com.

User Testing

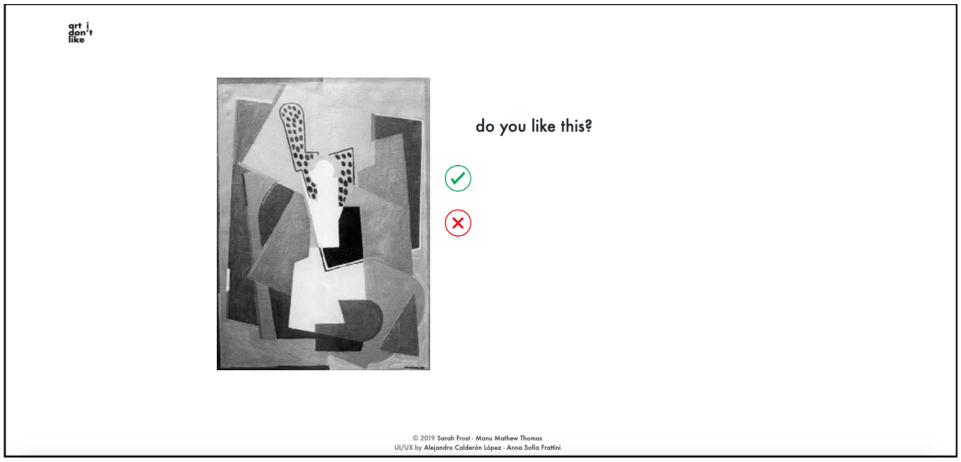

Figure 2.4: The updated interface of our system shows one painting. Users can click on the images to see them in their original resolution.

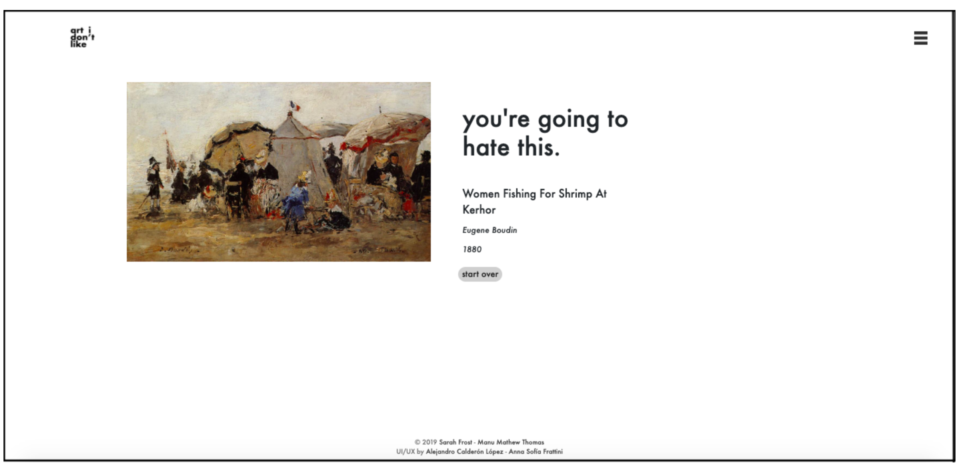

Figure 2.5: After the user has selected their preferred number of paintings, they can click "Find Art" to see the dissimilar image returned by the system.

We developed a study to gauge user emotion. We wanted to know if this updated design would increase feelings of curiosity in users. We decided to use the PANAS evaluation instrument, a questionnaire developed to allow users to self-report feelings through Likert scales. [74]. We made two adjustments: we included a "0" rating of a feeling or emotion to denote "Not at all", and changed "1" to denote "Very slightly" instead of "Very slightly or not at all". We added curious as a feeling, but chose to report changes in feelings of curiosity separately, and not include the amount in the general positive and negative affect.

We conducted the study with Aurora Alparaz, a Cognitive Science undergraduate student working in the Creative Coding lab. We ran the study with six users who were graduate students or Post-Doctoral researchers in the Computational Media Department, who had heard little or nothing about this project. Because we had collected the names of the students who completed the questionnaire for the original design, we made sure that those users were not included in this study.

We began each user test with a brief overview of the project. To control for users' state of minds before looking at the system, we conducted a pre-use PANAS test. After approximately five minutes of use, we stopped each user and conducted a post-use PANAS test. Finally, we asked the user to fill out a survey with several open ended questions. The time instruction with both the pre-use and post-use PANAS tests was "Today: You have felt this way today" [74]. Some general observations we had across the users were that they liked the simple, clean design, they did not look at the information section of the website, and finally, that they chose to like or dislike many pieces of art before clicking “Find Art” and seeing the resulting image. Although we did not track specifically how many pieces of art each user looked at before requesting a result, it was many more than the nine images presented to the user in the original design. In the survey we asked the users, “How many pieces of art should a user like before being presented with a disliked piece of art?” The answers we received were:

“Probably dozens, maybe hundreds”

“All of them”

“Depending on the system...the number should increase as more factors are considered”

“50”

“At least 15”

“5 minimum”

This was surprising to us, because we assumed that users would not want to spend a significant amount of time evaluating pieces of art. In a follow up study, we would track how long users spend looking at each piece of art, and use that to make decisions about the optimal length of the experience. We also asked users to tell us how likely they were to recommend Art I Don’t Like to a friend. On a scale of 1 to 7, with 1 denoting “not at all”, and 7 denoting “very much so” the mean was 2.5. This is lower than we were hoping, and has made me think about redesigning the project again - and considering alternative modes of experience. In terms of the positive and negative affect change between the pre- and post-use tests, there was an average decrease of 2.5 points over all 10 positive feelings/emotions, and an average decrease of 4.83 points over all

10 negative feelings/emotions. There was an average increase of .5 points for "curious". Although we would have liked to conduct the study with more participants, given the events of the Winter quarter, we did not have enough time to conduct more.

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 to 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.)

Chapter 3

CompostNet

This chapter describes a project I conducted in the class CMPS 290T - Machine Learning For Social Good in Spring 2019. I completed this project with the help of Bryan Tor, Dr. Angus Forbes, and Dr. Rakshit Agrawal. I then presented this work in Fall 2019 at the IEEE Global Humanitarian Technology Conference in Seattle, Washington.

3.1 Introduction

Many businesses, cafes, and outdoor spaces provide trash, recycling, and composting bins, requiring consumers to decipher instructional text, icons, or images in order to sort their waste accurately. It can be confusing to know what pieces of waste go in which bin. Moreover, different areas may have different rules for how to separate waste, and people often inadvertently throw their trash in the wrong bin. Machine learning solutions can help us more quickly and accurately choose the proper receptacle for our waste by classifying a photograph of the waste. This paper presents a novel image classification model that categorizes the types of waste produced after eating a meal, which can be used in mobile applications to encourage users to correctly sort waste. Building on recent work in deep learning and waste classification, we introduce CompostNet, a convolutional neural network that classifies images according to how they should be appropriately discarded. We provide details about the design and development of CompostNet, along with an evaluation of its effectiveness in classifying images of waste. Further, we discuss two different approaches to the design of our system, one using a custom model and the other augmenting a pre-trained image classification model (MobileNet) through transfer learning, and and how we achieved greater success with the transfer learning approach. To the best of our knowledge, CompostNet is the first waste classification system that uses a deep learning network to identify compostable, recyclable, and landfill materials. CompostNet is an application of machine learning for social good, and supports United Nations Sustainable Development Goal 12: Responsible Consumption and Production.

3.1.1 Problem Definition

Junk, waste, rubbish, garbage, trash. Whatever its called, humans produce a lot of it. The United States generates 624,700 metric tons of solid waste each day [2]. Recycling and composting programs exist (see figure 3.1), but are vulnerable to recycling contamination. One in four items placed in a recycling container are not actually recyclable [2] which contaminates the surrounding materials, making it unable to be processed. When individuals sort their waste after a meal, they may not know

Figure 3.1: Waste receptacles at the National Institutes of Health in Bethesda, MD.

what is recyclable, what is trash, and what is compostable. When people are unsure about how to sort their waste, more of it will be misplaced, resulting in recycling contamination and otherwise recyclable material being sent to landfills.

3.1.2 Motivation

This work is motivated by a concern about trash. Waste management and organization is a growing concern for many groups. Goal 12 of the United Nations Sustainable Development Goals is "Responsible Consumption and Production" [1], with target 12.5 aiming to "Substantially reduce waste generation through prevention, reduction, recycling and reuse" [1]. Similarly, the European Commission has an environmental policy that sets several priority objectives for waste policy [3]. Extensive research has been done to study waste management across the globe [4]–[7]. In this project, we use machine learning, as it can be used to classify images with a high rate of accuracy. Although efforts to improve waste sorting accuracy must be multifaceted, this system can be used at the first point of differentiating types of waste, and will help people learn how to correctly sort their waste.

3.1.3 Terminology

In this paper, we use the term ‘waste’ to refer to all material that is discarded. This term encompasses all of the materials that we are classifying. We use the labels ‘landfill’, ‘recyclable’, and ‘compostable’ to refer to the locations or processes that these materials will go to or undergo after they are correctly disposed. We also use the term ‘trash’ as a synonym to ‘landfill’. There are different guidelines for recyclable and compostable materials, based on the recycling and composting fa- cilities for a municipality. We use the San Francisco Recology guidelines [8]. Recology is a San Francisco-based integrated resource recovery company that processes compostable waste in an industrial composting facility that can break down fish, meat, dairy, and bio-plastics [9]. In addition, some recycling facilities will not recycle plastic utensils, while Recology will accept plastic utensils [9].

3.2 Related Work

3.2.1 Image Classification

Image classification is one of the major applications of artificial intelligence. Recent image classification models often rely on deep neural networks, specifically Convolutional Neural Networks (CNNs). CNNs are neural network variants that learn by performing convolutions and have shown stellar performance for image classification tasks [11]. Image classification relies on supervised learning, which requires labeled data to train networks. After the network is trained, it can classify images into discrete classes [11]. Using image classification, these networks can answer questions related to the visual properties of an image [12].

3.2.2 Waste Classification

Waste classification can be addressed in many ways - from educating individuals about sorting household trash [5] to using a hyperspectral imaging system to analyze attributes of waste products at compost or recycling facilities [7]. Few researchers have studied the use of CNNs to develop image recognition models for classifying waste. However, there are two examples of waste classifying CNNs. Gary Thung and Mindy Yang built the CNN "TrashNet" to classify waste into five classes of recyclable content and trash [10]. Spot- Garbage [13] is a mobile application designed by researchers at the Indian Institute of Technology. This app allows users to identify garbage in the street around Indias urban centers. SpotGarbage uses a CNN called GarbNet, which has been trained on an annotated dataset called Garbage In Images. Although both are examples of CNNs used for waste classification, they do not classify food waste, or compostable material, separately.

Figure 3.2: Three images from our compostable class.

3.3 Method

3.3.1 Dataset

We began with the data collected by Thung and Yang for "TrashNet". Their dataset consists of 2527 images in six classes: glass, plastic, cardboard, metal, paper, and trash [10]. We wanted to train our models on images of compostable content, in addition to recyclable and landfill content. We kept the subcategories of recyclables, but were more concerned with the three categories of 'trash', 'recyclable', and 'compostable'. We augmented their dataset by adding 175 photos of food waste and 49 photos of landfill waste, for a total of 2751 images. Three example images are seen in figure 3.2. We also moved images previously classified as "trash" to the "compostable" class, bringing the total image count to 177. We followed the methods for data collection that Thung and Yang outline in their project. We took photos of the waste against a white poster board, used natural or overhead lights, and focused on one piece of waste in each photo. Although Thung and Yang resized their images to 512 by 384 pixels, we resized our images to 400 by 300 pixels for our version B model. Our version A model required images to be resized to 224 by 224 pixels.

3.3.2 Models

1. CompostNet - Version A

Our first version of CompostNet utilized a pretrained MobileNet model that was partially retrained on our dataset. MobileNet is a lightweight mobile-first convolutional neural network trained on the ImageNet database [14]. We chose MobileNet because it balances efficiency and accuracy. We trained CompostNet - Version A on Tensorflow v. 2. This pre- trained model serves as the first layer in our model. After the MobileNet model layer we have 4 layers which have been trained on our dataset. The last dense layer has a softmax function to present us with the seven outputs corresponding to the waste classes. The system architecture, including the MobileNet layer and our retrained layers, is shown in Figure 3.3.

2. CompostNet - Version B

This version of CompostNet was built around three convolutional layers. Our first layer breaks down the input image into 32 output matrices. It then goes through a max pooling layer with a 2x2 filter, halving the dimensionality of the output matrices, which reduces computation costs and helped us avoid overfitting. We apply a dropout rate

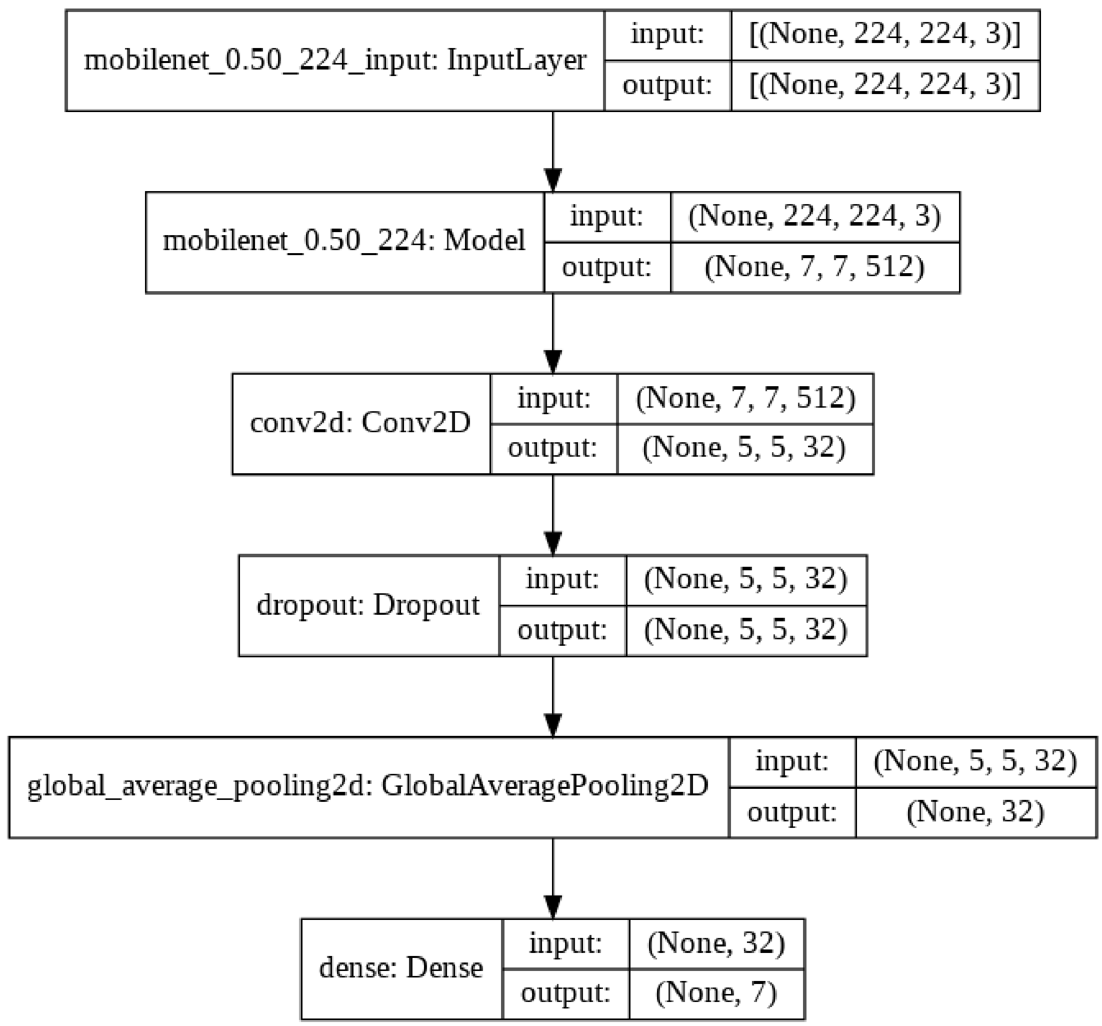

Figure 3.3: Architecture for CompostNet - Version A.

of 30% after this step to further avoid overfitting. In our testing, we found that this dropout rate was most effective in improving the accuracy of this model. These layers are replicated twice, but with different output matrices of 64 and 128 respectively. Finally, we flattened the output and passed it through a densely- connected neural network layer. The network architecture can be viewed on our Github.

3.4 Results and Evaluation

Thung and Yang achieved an accuracy of 75% using their CNN 'TrashNet'. We took that as our baseline and aimed to exceed that level of accuracy.

3.4.1 Models

A. CompostNet - Version A

We split the dataset into two groups, 80% of images were used to train the network, and 20% of the images were excluded from training to test the network after it had finished training. We trained our model for 20 epochs, with a batch set of 64. We then tested the accuracy of the model on test image classification. Our network returned confidence scores for each of our seven classes. These values add up to one, and demonstrate how the model has classified the object. Once the network has been trained, it can classify images in approximately two seconds. An example test photo and the network output scores is shown in Figure 3.4. It is clear that the dataset must be expanded because of the amount of items that will need to be classified if a hardware system is deployed with this model. After achieving a high training and test accuracy, the final test accuracy was 77.3%, as shown in Figure 3.5.

B. CompostNet - Version B

For this model, our train-test split was also 80/20. We ran the model for 20 epochs, with a batch size of 8. Training the model any further significantly decreased the accuracy, likely because it was overfitting to the data. The model's final test accuracy

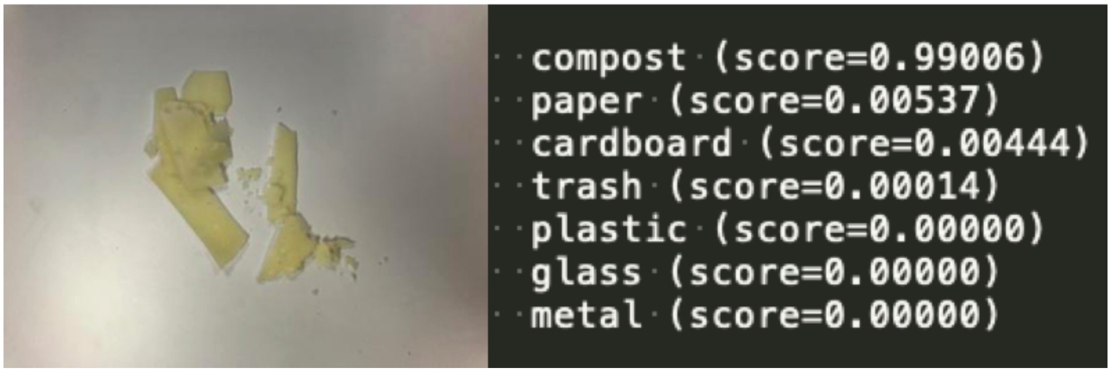

Figure 3.4: A test image and the confidence scores returned by Version A model.

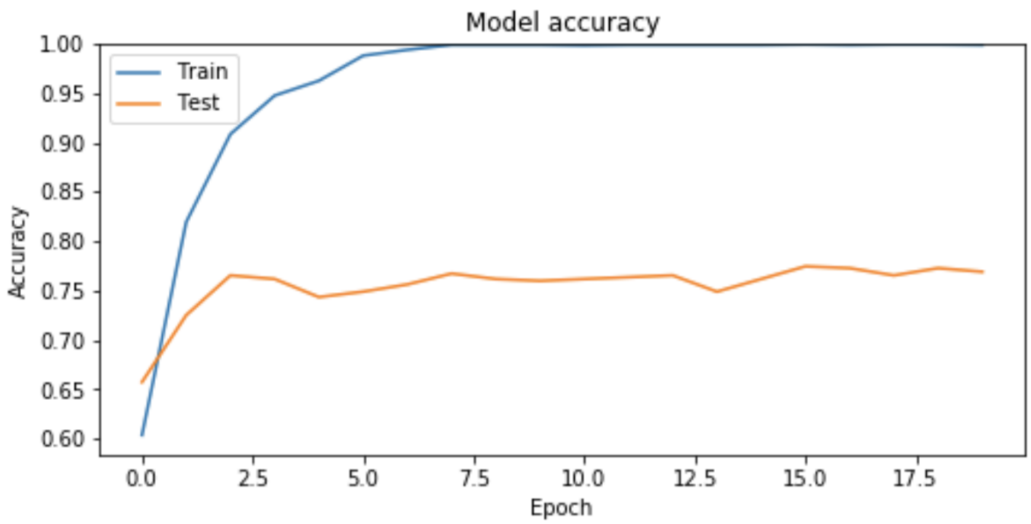

Figure 3.5: Test and train accuracy for CompostNet - Version A.

was 22.695%, see Figure 3.6. Version A had an accuracy of 77.3% while Version B had an accuracy of 22.695%. This difference is not surprising. Version A has 159 layers and has been trained on the ImageNet dataset, which has over 1.3 million images [15]. Version B is significantly less robust. Expanding our dataset and adjusting the hyperparameters of either model may improve the accuracy.

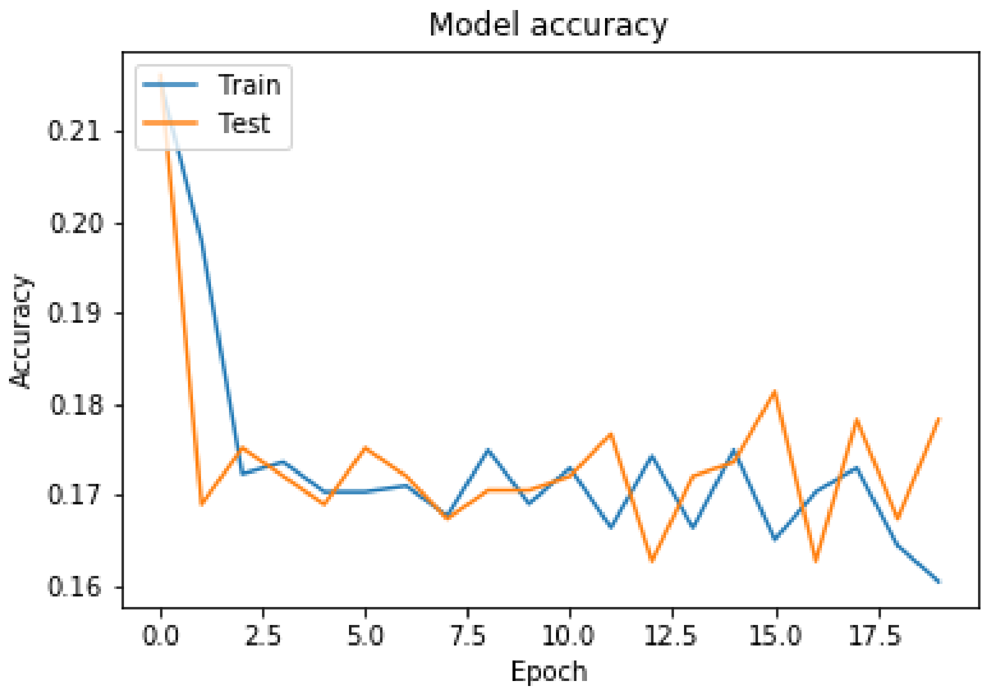

Figure 3.6: Test and train accuracy for CompostNet - Version B.

3.4.2 Deployment

We have validated that our system classifies images with an accuracy above 75%, which was our baseline. To deploy this model, we have built a prototype of an iPhone app that users can use to identify if their waste can be recycled.

We saved our Version A model and converted the saved model to a TensorFlow lite compatible format. We used our model in the TensorFlow Lite image classification iOS project developed by Google [16]. This code project provides the user interface for the ML model. The user can open the app and hover the camera over a object. Our model is then used to classify the image, and the top three classes are returned for the user to see, as shown in Figure 3.7. This app has limited functionality, for example, it does not allow the user to choose a photo from his or her photo library. Although this iOS application is fairly simple, this deployment is exciting and represents a proof of concept.

Figure 3.7: The iOS app after the system has classified an image.

3.4.3 Future Work

We would like to expand the types of compostable materials (bio-plastic, paper plates, bamboo utensils, etc.) in our dataset to improve the accuracy of the CompostNet model. The Office of Sustainability at the University of California, Santa Cruz has expressed interest in this project as it aligns with the University of California sustainability goal of 90% of waste diverted from landfill by 2020 [17]. In an email sent to the UCSC community, Associate Chancellor Ashish Sahni wrote to students: "Campus recycling is temporarily being landfilled due to high rates of contamination, with the highest rates in our residential and dining halls...If you are unsure if something is recyclable, it is better to throw it into the landfill bin!" [18]. Our CompostNet system can help educate students about what can be recycled and composted at the end of a meal, to reduce the amount of material that is incorrectly sorted, preventing further cases of recycling contamination.

We have analyzed two models for image classification of three categories of waste - landfill, recyclable, and compostable. We demonstrate that with a small dataset, a model based on transfer learning returns good results for recognizing images of waste. The growing amount of waste sent to landfills is a rising concern addressed by the United Nations Sustainable Development Goal 12.5, as well as other countries and organizations. It is incumbent upon all of us to divert waste from landfills, and reduce the amount of waste we create. Our dataset and code are available for download at https: //github.com/sarahmfrost/compostnet.

Chapter 4

Conclusion

During my time as a Master's student in the Computational Media department, I developed an interest in using programming and technology to tackle social and environmental issues. In Spring 2019, I took CMPS 290T - Machine Learning for Social Good with Rakshit Agrawal. In this class I increased my foundation of knowledge in Machine Learning. It also became clear to me that the vocabulary and epistemology of the ML4SG field is still in flux. "Machine Learning for Social Good", "Tech For Good", and "Artificial Intelligence for Social Good" were used somewhat interchangeably. It was unclear what features these projects share and don't share, and as more and more projects are developed, it is important to have a clear understanding of the features shared by these projects. I present an initial characterization of attributes of ML4SG projects.

As a representation of my time in the Computational Media Department, this thesis is the realization of a computational media project. The strength of Computational Media derives from its synthesis of four different types of work - technical, creative, interpretive, and collaborative [73]. The two sections of this thesis combine various amounts of three of these. Art I Don't Like and CompostNet are examples of creative, technical, and collaborative work. Additionally, the mission statement of Computational Media includes two core principles that resonated with me as I researched and wrote my thesis: the idea that computational media practitioners "actively perform research and education that benefit society at large and advance social goals...We are unafraid of handling real-world problems that are naturally messy" [2]. It is these values that inspired me to develop technical projects that address social and environmental issues.

I have presented two fully functional prototypes for machine-learning based systems targeting societal challenges. The first project is Art I Don't Like, an anti-recommender system for visual art. Using the categories developed with Dr. Agrawal, I determine that Art I Don't Like is a Social good-first, individual-oriented, project at small-scale deployment. The second project is CompostNet, an image classification model for personal meal waste. CompostNet is a Social Good-First, Individual-oriented, undeployed project. This categorization has been integral in helping me to determine what I value as a researcher, and what I will focus on in the future.

Bibliography

[1] Ai commons.

[2] Computational Media: Mission Statement.

[3] *Goal 12: Responsible Consumption and Production.*

[4] *How a Japanese cucumber farmer is using deep learning and TensorFlow | Google Cloud Blog.* Google.

[5] *A Simple List of What Can and Cannot Be Recycled.*

[6] Sustainability. Technical report, University of California.

[7] Waste. Technical report, European Commission.

[8] *The world's trash crisis, and why many Americans are oblivious.* Los Angeles Times, April 2016.

[9] *INTERNATIONAL CONFERENCE ON INNOVATIVE COMPUTING AND COMMUNICATIONS.* SPRINGER Verlag, SINGAPOR, Place of publication not identified, 2019. OCLC: 1128102295.

[10] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: A System for Large-Scale Machine Learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283, Savannah, GA, 2016. USENIX Association.

[11] Flavia Sofia Acerbo and Claudio Rossi. Filtering informative tweets during emergencies: a machine learning approach. In *Proceedings of the First CoNEXT Workshop on ICT Tools for Emergency Networks and DisastEr Relief - I-TENDER '17*, pages 1–6, Incheon, Republic of Korea, 2017. ACM Press.

[12] G. Adomavicius and A. Tuzhilin. 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*, 17(6):734–749, June 2005.

[13] Rakshit Agrawal. Applied ML for Social Good - Lecture 1, April 2019.

[14] Everaldo Aguiar, Himabindu Lakkaraju, Nasir Bhanpuri, David Miller, Ben Yuhas, and Kecia L. Addison. Who, when, and why: a machine learning approach to prioritizing students at risk of not graduating high school on time. In *Proceedings of the Fifth International Conference on Learning Analytics And Knowledge - LAK '15*, pages 93–102, Poughkeepsie, New York, 2015. ACM Press.

[15] Sultan A. Alharthi, Olaa Alsaedi, Zachary O. Toups, Joshua Tanenbaum, and Jessica Hammer. Playing to Wait: A Taxonomy of Idle Games. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*, pages 1–15, Montreal QC, Canada, 2018. ACM Press.

[16] Ferran Altarriba Bertran, Samvid Jhaveri, Rosa Lutz, Katherine Isbister, and Danielle Wilde. Making Sense of Human-Food Interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, pages 1–13, Glasgow, Scotland Uk, 2019. ACM Press.

[17] G Bonifazi, S Serranti, A Bonoli, and A Dall'Ara. Innovative recognition-sorting procedures applied to solid waste: the hyperspectral approach. *WIT Transactions on Ecology and the Environment*, 120:885–894, 2009.

[18] Frederik J Zuiderveen Borgesius, Damian Trilling, Judith Möller, Balázs Bodó, Claes H de Vreese, and Natali Helberger. Should we worry about filter bubbles? *Internet Policy Review*, 5(1):16, 2016.

[19] Scott R Braithwaite, Christophe Giraud-Carrier, Josh West, Michael D Barnes, and Carl Lee Hanson. Validating Machine Learning Algorithms for Twitter Data Against Established Measures of Suicidality. *JMIR Mental Health*, 3(2):e21, May 2016.

[20] Fidel Cacheda, Víctor Carneiro, Diego Fernández, and Vreixo Formoso. Comparison of Collaborative Filtering Algorithms: Limitations of Current Techniques and

Proposals for Scalable, High-performance Recommender Systems. *ACM Trans. Web*, 5(1):2:1–2:33, February 2011.

[21] João Caldeira, Alex Fout, Aniket Kesari, Raesetje Sefala, Joseph Walsh, Katy Dupre, Muhammad Rizal Khaefi, Setiaji, George Hodge, Zakiya Aryana Pramestri, and Muhammad Adib Imtiyazi. Improving Traffic Safety Through Video Analysis in Jakarta, Indonesia. *arXiv:1812.01106 [cs, stat]*, 1038:642–649, 2020. arXiv: 1812.01106.

[22] Francois Chollet. *Deep Learning with Python.* Manning Publications Co., Greenwich, CT, USA, 1st edition, 2017.

[23] Michael Chui, Martin Harrysson, James Manyika, Roger Roberts, Rita Chung, Pieter Nel, and Ashley van Heteren. Applying artificial intelligence for social good, Nov 2018.

[24] Dan Cosley, Shyong K. Lam, Istvan Albert, Joseph A. Konstan, and John Riedl. Is Seeing Believing?: How Recommender System Interfaces Affect Users' Opinions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '03, pages 585–592, New York, NY, USA, 2003. ACM. event-place: Ft. Lauderdale, Florida, USA.

[25] Ayhan Demirbas. Waste management, waste resource facilities and waste conversion processes. *Energy Conversion and Management*, 52(2):1280–1287, 2011.

[26] Javier Di Deco, Ana M. Gonzalez, Julia Diaz, Virginia Mato, Daniel Gar-

cia–Frank, Juan Alvarez–Linera, Ana Frank, and Juan A. Hernandez–Tamames. Machine Learning and Social Network Analysis Applied to Alzheimer's Disease Biomarkers. *Current Topics in Medicinal Chemistry*, 13(5):652–662, April 2013.

[27] Mahika Dubey, Alan Peral Ortiz, Rakshit Agrawal, and Angus G. Forbes. Predicting Biker Density at Bikeshare Station Intersections in San Francisco. In *2019 IEEE Global Humanitarian Technology Conference (GHTC)*, pages 1–7, Seattle, WA, USA, October 2019. IEEE.

[28] Natalia Efremova, Dennis West, and Dmitry Zausaev. AI-based evaluation of the SDGs: The case of crop detection with earth observation data. *arXiv:1907.02813 [cs, eess]*, July 2019. arXiv: 1907.02813.

[29] Angus Forbes, Mihai Surdeanu, Peter Jansen, and Jane Carrington. Transmitting narrative: An interactive shift-summarization tool for improving nurse communication. In *Proceedings of the IEEE Workshop on Interactive Visual Text Analytics (TextVis)*, volume 2, 2013.

[30] Angus Graeme Forbes, Saiph Savage, and Tobias Höllerer. Visualizing and verifying directed social queries. In *IEEE Workshop on Interactive Visual Text Analytics. Seattle, WA, Citeseer*, 2012.

[31] Sarah Frost, Manu Mathew Thomas, and Angus G. Forbes. Art I Don't Like: An Anti-Recommender System For Visual Art. 2019.

[32] Huiji Gao, Geoffrey Barbier, and Rebecca Goolsby. Harnessing the Crowdsourcing Power of Social Media for Disaster Relief. *IEEE Intelligent Systems*, 26(3):10–14, May 2011.

[33] Antonio Gentile, Salvatore Andolina, Antonio Massara, Dario Pirrone, Giuseppe Russo, Antonella Santangelo, Eleonora Trumello, and Salvatore Sorce. A Multi-channel Information System to Build and Deliver Rich User-Experiences in Exhibits and Museums. In *2011 International Conference on Broadband and Wireless Computing, Communication and Applications*, pages 57–64, Barcelona, Spain, October 2011. IEEE.

[34] Heloïse Greeff, Achut Manandhar, Patrick Thomson, Rob Hope, and David Clifton. Rural infrastructure health monitoring system: Using ai to increase rural water supply reliability.

[35] Anees Ul Hassan, Jamil Hussain, Musarrat Hussain, Muhammad Sadiq, and Sungyoung Lee. Sentiment analysis of social networking sites (SNS) data using machine learning approach for the measurement of depression. In *2017 International Conference on Information and Communication Technology Convergence (ICTC)*, pages 138–140, Jeju, October 2017. IEEE.

[36] Stefan Hirschmeier and Vanessa Beule. Compliance of Personalized Radio with Public-Service Remits. In *Adjunct Publication of the 26th Conference on User Modeling, Adaptation and Personalization*, UMAP '18, pages 15–21, New York, NY, USA, 2018. ACM. event-place: Singapore, Singapore.

[37] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.

[38] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.

[39] Avishek Kumar, Syed Ali Asad Rizvi, Benjamin Brooks, R. Ali Vanderveld, Kevin H. Wilson, Chad Kenney, Sam Edelstein, Adria Finch, Andrew Maxwell, Joe Zuckerbraun, and Rayid Ghani. Using Machine Learning to Assess the Risk of and Prevent Water Main Breaks. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 472–480, London United Kingdom, July 2018. ACM.

[40] Himabindu Lakkaraju, Everaldo Aguiar, Carl Shan, David Miller, Nasir Bhanpuri, Rayid Ghani, and Kecia L. Addison. A Machine Learning Framework to Identify Students at Risk of Adverse Academic Outcomes. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '15*, pages 1909–1918, Sydney, NSW, Australia, 2015. ACM Press.

[41] Sandra Lebersorger and Felicitas Schneider. Discussion on the methodology for de-

termining food waste in household waste composition studies. *Waste management*, 31(9-10):1924–1933, 2011.

[42] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.

[43] Kristen Lee and Sarah Frost. *Phone conversation with Kristen Lee - Sustainability Programs Manager at University of California, Santa Cruz*. May 2019.

[44] Kyumin Lee, James Caverlee, and Steve Webb. Uncovering social spammers: social honeypots + machine learning. In *Proceeding of the 33rd international ACM SIGIR conference on Research and development in information retrieval - SIGIR '10*, page 435, Geneva, Switzerland, 2010. ACM Press.

[45] Yun Liu, Timo Kohlberger, Mohammad Norouzi, George E. Dahl, Jenny L. Smith, Arash Mohtashamian, Niels Olson, Lily H. Peng, Jason D. Hipp, and Martin C. Stumpe. Artificial Intelligence–Based Breast Cancer Nodal Metastasis Detection: Insights Into the Black Box for Pathologists. *Archives of Pathology & Laboratory Medicine*, 143(7):859–868, July 2019.

[46] Mohammad Saeid Mahdavinejad, Mohammadreza Rezvan, Mohammadamin Barekatain, Peyman Adibi, Payam Barnaghi, and Amit P. Sheth. Machine learning for internet of things data analysis: a survey. *Digital Communications and Networks*, 4(3):161–175, August 2018.

[47] Claudia Meda, Federica Bisio, Paolo Gastaldo, and Rodolfo Zunino. A machine

learning approach for Twitter spammers detection. In *2014 International Carnahan Conference on Security Technology (ICCST)*, pages 1–6, Rome, Italy, October 2014. IEEE.

[48] Rachel Metz. An ex-Google engineer is scraping YouTube to pop our filter bubbles. *MIT Technology Review*, April 2018.

[49] Gaurav Mittal, Kaushal B Yagnik, Mohit Garg, and Narayanan C Krishnan. SpotGarbage: smartphone app to detect garbage using deep learning. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 940–945. ACM, 2016.

[50] E. Pariser. *The Filter Bubble: What The Internet Is Hiding From You.* Read more. Penguin Books Limited, 2011.

[51] Daniela Petrelli, Mark T. Marshall, Sinéad O'Brien, Patrick McEntaggart, and Ian Gwilt. Tangible data souvenirs as a bridge between a physical museum visit and online digital experience. *Personal and Ubiquitous Computing*, 21(2):281–295, April 2017.

[52] Recology. *What Goes Where.*

[53] Francesco Ricci, editor. *Recommender systems handbook.* Springer, New York, 2011. OCLC: ocn373479846.

[54] Will Rinehart. The Election of 2016 and the Filter Bubble Thesis in 2017. *Medium*, September 2017.

[55] Martin Risseeuw, Dario Cavada, Elena Not, Massimo Zancanaro, Mark T. Marshall, Daniela Petrelli, and Thomas Kubitza. Authoring Augmented Digital Experiences in Museums. In *Proceedings of the International Working Conference on Advanced Visual Interfaces - AVI '16*, pages 340–341, Bari, Italy, 2016. ACM Press.

[56] Caleb Robinson and Bistra Dilkina. A Machine Learning Approach to Modeling Human Migration. *arXiv:1711.05462 [physics]*, November 2017. arXiv: 1711.05462.

[57] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, and others. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.

[58] Associate Chancellor Ashish Sahni. *We need your help with campus recycling.* May 2019.

[59] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobileNetV2: Inverted Residuals and Linear Bottlenecks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2018.

[60] John W Santerre, James J Davis, Fangfang Xia, and Rick Stevens. Machine learning for antimicrobial resistance. *arXiv preprint arXiv:1607.01224*, 2016.

[61] Kaz Sato. How a Japanese cucumber farmer is using deep learning and TensorFlow, August 2016.

[62] Juliana Shihadeh, Anaam Ansari, and Tokunbo Ozunfunmi. Deep Learning Based Image Classification for Remote Medical Diagnosis. In *2018 IEEE Global Humanitarian Technology Conference*, pages 1–8. IEEE, 2018.

[63] Xiaoyuan Su and Taghi M. Khoshgoftaar. A Survey of Collaborative Filtering Techniques. *Advances in Artificial Intelligence*, 2009:1–19, 2009.

[64] Scott Sweeney. *Industrial Composting: What It Is and How It Works*. May 2018.

[65] Stephanie Tam. Pop Goes the Filter Bubble? *Slate*, April 2018.

[66] Wei Ren Tan, Chee Seng Chan, Hernan Aguirre, and Kiyoshi Tanaka. Improved ArtGAN for Conditional Synthesis of Natural Image and Artwork. *IEEE Transactions on Image Processing*, PP:1–1, 2018.

[67] Tensorflow. *tensorflow/examples*. July 2019.

[68] Gary Thung and Mindy Yang. *Trashnet*.

[69] University of Deusto, Begoña Garcia-Zapirain, Yolanda Garcia-Chimeno, and Heather Rogers. Machine Learning Techniques for Automatic Classification of Patients with Fibromyalgia and Arthritis. *International Journal of Computer Trends and Technology*, 25(3):149–152, July 2015.

[70] Michael Veale and Reuben Binns. Fairer machine learning in the real world: Mitigating discrimination without collecting sensitive data. *Big Data & Society*, 4(2):205395171774353, December 2017.

[71] Christian Voigt, Calkin Suero Montero, and Massimo Menichinelli. An Empirically Informed Taxonomy for the Maker Movement. In Franco Bagnoli, Anna Satsiou, Ioannis Stavrakakis, Paolo Nesi, Giovanna Pacini, Yanina Welp, Thanassis Tiropanis, and Dominic DiFranzo, editors, *Internet Science*, volume 9934, pages 189–204. Springer International Publishing, Cham, 2016.

[72] Alex Hai Wang. Detecting Spam Bots in Online Social Networking Sites: A Machine Learning Approach. In David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, Sara Foresti, and Sushil Jajodia, editors, *Data and Applications Security and Privacy XXIV*, volume 6166, pages 335–342. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. Series Title: Lecture Notes in Computer Science.

[73] Noah Wardrip-Fruin and Michael Mateas. Envisioning the Future of Computational Media. Technical report.

[74] David Watson, Lee Anna Clark, and Auke Tellegen. Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology*, 54(6):1063–1070, 1988.

[75] Jiafan Yu, Zhecheng Wang, Arun Majumdar, and Ram Rajagopal. DeepSolar: A Machine Learning Framework to Efficiently Construct a Solar Deployment Database in the United States. *Joule*, 2(12):2605–2617, December 2018.

[76] Ying Zhuang, Song-Wei Wu, Yun-Long Wang, Wei-Xiang Wu, and Ying-Xu Chen. Source separation of household waste: a case study in China. *Waste management*, 28(10):2022–2030, 2008.