A critical analysis of the negative consequences caused by recommender systems used on social media platforms

Author: Koen Zoetekouw
University of Twente
P.O. Box 217, 7500AE Enschede
The Netherlands

ABSTRACT,
In a world where social media has a very broad reach and impact on society, it is important to stay alert and keep monitoring content that gets spread so easily. Machine learning algorithms used by social media networks do not always function like they should which can lead to undesirable situations. Social media platforms use recommender systems to personalize content according to the user's preference and therefore tailor the enormous amount of content available on the internet. This thesis will shed light on the limitations of recommender systems used on social media platform and the possible negative consequences that these systems will bring with them. Recommendations on how to either avoid or deal with these consequences will also be provided.

Graduation Committee members:
Dr. A.B.J.M. Wijnhoven (1st examiner)
Dr. M. de Visser (2nd examiner)

Keywords
[Machine learning, recommender systems, collaborative filtering, social media, bias]

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

CC-BY-NC

1. INTRODUCTION

1.1 Background

"How YouTube's algorithm facilitates sexual exploitation of children". This is a headliner of a news article which was published on www.androidpit.com. (Dalul, 2019). It discusses a phenomenon which is called the "down the rabbit hole effect" where the algorithm of YouTube keeps recommending the same type of content to keep you watching, even if this content is bad content. But what kind of algorithm are they talking about? The news article discusses a machine learning algorithm, which is a self-learning algorithm and, in this case, used by the social media network YouTube. This specific type of a machine learning algorithm is called a recommender system, which is used to provide suggestions to the users which interact with the platform and therefore gives personalized content to the user. These recommender systems are widely used by many social media networks, like YouTube, Facebook, Twitter, Instagram etc. However, machine learning algorithms have some drawbacks. The problem with machine learning algorithms is that the creators of those algorithms know what they did to build them, but they don't know how they actually do what they do once they are in use. This problem is called the "black box problem" and is written about in many papers (Zednik, 2019). The main concern with this black box problem is that because of the lack of transparency of the algorithm it is extremely difficult to filter out the problems and, if unattended, the algorithm can become very biased or even "evil" (Vincent, 2016)[1].

1.2 Problem definition

The big question now is: how is this a relevant problem? Social media platforms nowadays are getting more and more influence and power than ever before. As of January 2019, according to The Global State of Digital Report 2019, there are 3.484 billion active social media users ("The Global State of Digital in 2019 Report," 2019). Of those social media users 2.320 billion people are Facebook users, and 1.9 billion people are YouTube users ("Most popular social networks worldwide as of April 2019, ranked by numbers of active users (in millions)," 2019). Compared to traditional media, social media is a faster medium to spread news, more easily accessible, more unfiltered, and able to reach a younger audience and therefore has a broader reach. A report from 2018 on media use of children explains that a lot of children (from the age of 3 to 15) already possess either a mobile phone or a table (or both). They also start using social media networks from a young age, as 71% of children between age 12 to 15 have a social media profile. Social media networks like YouTube are used even earlier on, as 45% of children between the age of 3 and 4 are already watching videos on YouTube (Ofcom, 2019). The ease of use and accessibility of social media networks has many advantages when compared to traditional media and it contains some abilities which can be used to achieve great things. However, every advantage has its disadvantage, and that is also the case with social media networks.

In the current era of fake news and influence from the outside, like the infamous Russian interference in the 2016 United States elections, these networks need to be carefully used and monitored to make sure that this big audience that these networks have gathered will not be influenced or informed in a wrong way.

If these algorithms used in social media networks to provide recommendations to their users are spiraling out of control, then there could be a very likely chance that certain problems will slowly emerge. Unfortunately, the article mentioned above is not the only article that talks about these algorithms creating problems. The following articles contain several examples of social media network algorithms which create problems.

"YouTube's algorithm is spreading a series of unfortunate far-right events." (Scott, 2019)

"YouTube's recommended videos algorithm is promoting controversial content." (Mason, 2018)

"Facebook's ad-serving algorithm discriminates by gender and race." (Hao, 2019)

"How Twitter's algorithm is amplifying extreme political rhetoric." (Darcy, 2019)

If we combine these incidents with the knowledge of the broad reach social media has these days, it is the recipe for disaster and if this remains unchecked it could lead to some very negative consequences. Fake news could reach the big audience that social media networks have gathered far more easily, and bad content also can find its way better to the targeted people.

1.3 Research questions

This research is aimed at shedding a new light on the implications of recommender systems used in social media networks. The goal is to make people aware of the implications and, in the end, provide recommendations on how to avoid or deal with these implications. The main research question is:

What are the negative consequences of recommender systems used in social media networks and how can people avoid or deal with them?

To get to the answer of this research questions some sub questions are needed which provide structure. These sub questions are listed below.

a. *What are the challenges that limit the recommender systems used in social media networks from working optimally?*
b. *How do these challenges translate to negative consequences?*
c. *What is the worst-case scenario for every negative consequence?*
d. *How can people avoid or deal with these consequences to limit the damage done?*

1.4 Thesis structure

In this thesis, the recommender systems used in social media networks will be investigated to identify where these systems experience challenges and what the implications would be. The results will be evaluated, and recommendations will be made to avoid or deal with these consequences. The thesis will be structured as follows: firstly, a literature will be conducted to make sure that the reader can get a grasp of the various concepts that are used in the recommender systems of social media networks. This will be done by introducing concepts like machine learning, general recommender systems, collaborative filtering, and giving a real-life example, the YouTube algorithm, to connect the theory in context. After this, a second literature review will be conducted to identify challenges in these recommender systems. To conclude this thesis, a case-based analysis will be conducted to connect both the theory and the challenges with each other to identify the possible negative consequences that emerge from using a recommender system on social media platforms. Recommendations on how to deal with or avoid these consequences will also be given.

[1] Article on "The Verge" about Microsoft's AI chatbot who, on the same day of its launch, began making racist comments.

2. METHODOLOGY

The goal of this thesis is to address the problem stated in the introduction and answer the research question. To achieve this goal, this research will be structured according to the sub questions stated in chapter 1.3. First, a thorough literature review needs to be conducted concerning the mechanisms of machine learning and recommender systems. The literature review will provide a theoretical framework for this thesis, which will be used to assist with the identification of possible challenges that may be present in the recommender systems used in social media platforms. The literature used in this thesis will primarily be sources from databases like Google Scholar and Scopus with keywords like "Machine learning", "Recommender system", and "Collaborative filtering". Secondly, a literature review will be conducted on the challenges that recommender systems offer, which will bring new insights about the possible negative consequences that these algorithms bring with them. The literature used for this part will overlap with the same literature as was used in the first part as most papers with these algorithms as the subject already touch on the challenges that come with them. At last, the theoretical framework and the challenges identified in the algorithms will be used to identify the possible negative consequences that come with recommender systems used in social media. To provide context to these consequences, a case-based analysis will be performed to verify that these consequences exist.

3. MACHINE LEARNING AND RECOMMENDER SYSTEMS

Machine learning and recommender systems consist of a lot of concepts and terms which may not be known to everyone, which might cause confusion later in this thesis. The introductory theory is therefore needed to serve as a theoretical framework for this thesis and to explain the concepts and terms that are used later on.

3.1 introduction to machine learning

Machine learning (ML) is a subfield of artificial intelligence and has evolved out of the need to teach computers how to automatically learn a solution to a problem (Essinger & L. Rosen, 2011). It evolved from studying pattern recognition and computational learning theory (Simon, Singh Deo, Selvam, & Babu, 2016). The goal of machine learning is to learn algorithms to carry out tasks by providing them with a couple of examples (what they need to do, and what they don't need to do) (Richert & Coelho, 2013). The need for Machine learning was created by realizing that there are tasks that are not within human capabilities to accomplish (Shalev-Schwartz & Ben-David, 2014). These tasks namely consist of the analysis of very large and complex data sets, which can be analyzed by machine learning algorithms and are able to detect patterns, take conclusions, and create certain outputs with increasing speed and accuracy. A paper by Grace, Salvatier, Dafoe, Zhang & Evans (2017) even suggests that it will not be very long before AI is able to outperform humans (Grace, Salvatier, Dafoe, Zhang, & Evans, 2017). Machine learning involves two types of tasks, which are supervised machine learning and unsupervised machine learning. Supervised machine learning is "trained" on a pre-defined set of "training examples". Through this "training" the algorithm can learn to discriminate between given concepts and therefore can reach more accurate conclusions when it encounters new data in the form of real life "examples" (Carbonell, 1989). Examples of supervised learning are the recommendation algorithms at Amazon, Facebook, Google, Netflix, YouTube, etc. Unsupervised machine learning is used to decide which examples belong to which classes, and what those classes are (Gennari, Langley, & Fisher, 1989). The most important uses for unsupervised machine learning are clustering and discovery (Shavlik & Dietterich, 1990). Examples of unsupervised learning are spam filters, face recognition, or speech recognition.

Recommender systems, which are studied in this thesis, are part of the supervised learning "track", so the unsupervised learning "track" will remain uncovered further.

3.2 Recommender systems

This thesis focuses on a specific type of machine learning algorithms used in social media networks which are called recommender systems. In this part recommender systems will be briefly introduced and an overview of the different types of recommender systems. will be given.

3.2.1 Definition

Recommender systems are software tools and techniques providing suggestions for items to be of use to a user. These "items" can range from books and clothing, to videos and news articles (Deng, 2019). It uses recommendations from people as inputs, which the system aggregates and then sends through to the appropriate recipients. Some systems put the emphasis more on making good matches between people who put in recommendations and people who seek for recommendations. (Resnick & Varian, 1997). Recommender systems consist of a machine learning algorithm that interacts in a way with the users. The goal of the system is to provide recommendations to the user which are suited to the user's preferences. According to a paper, recommender systems differ from other machine learning algorithms based on two principles. Principle one is that a recommender system is personalized, which means that it is not meant to represent group consensus but is meant to optimize the experience of one user. Principle two is that a recommender system is intended to help the user with decision making, especially in cases where the items are already known (R. Burke, Felfernig, & H. Göker, 2011).

3.2.2 Taxonomy

There are several types of recommender systems known at this moment, which all serve a different purpose and are used in different domains. In this part a compact summary of the literature that covers different types of recommender systems will be provided. The difference between the recommender systems is usually made based on the addressed domain, the knowledge that is used. This means that the different types of recommender systems gather their user data from different sources like video metadata, geographic area, demographics, or even relationships between users (Ricci, Rokach, & Shapira, 2011).

The overview is provided by a paper by Burke (2002; 2007) which discusses the different types of recommender systems. The paper distinguishes between six different types of recommender systems. These are content-based, collaborative filtering, demographic, knowledge-based, community-based, and hybrid systems (R. D.Burke, 2007; R. D. Burke, 2002). I will now elaborate on each type and give some practical examples of where these systems are used.

3.2.2.1 Content-based

Content-based recommender systems learn to recommend items that are similar to the items that the user positively rated in the past. Most of the time the user creates a profile on a certain web page, which will save the types of items that are of the user's liking. Then future items will be compared to the user profile to determine which items to recommend. An example of content-based recommendation is the Amazon "favorites" section, where items are recommended by keeping track of the categories of items purchased by users (M.J. Pazzani & Billsus, 2007).

3.2.2.2 Collaborative filtering
Collaborative filtering recommends items to users based on the opinions of other people. This type of recommendation algorithm measures similarity by looking at the rating history (what certain users like or do not like) of users with a similar taste and then recommends items with a high similarity (B. Schafer et al., 2007). Collaborative filtering is the most widely used method for recommendation systems and can be found in the recommendation algorithms of big web platforms like Amazon, YouTube, Facebook, etc. In the next part of the literature review a more detailed explanation of collaborative filtering algorithms will be given, because these algorithms are mostly used in the recommender systems of social media networks.

3.2.2.3 Demographic
As it's in the name already, demographic based recommender systems, recommend items based on the demographic profile of the user. The creator of the algorithm gathers data in the form of attributes of a person and uses this to create demographic classes (Michael J Pazzani, 1999). In this way, different recommendations can be given for different demographic niches. An example of this kind of algorithm is given by a paper from Krulwich about LifestyleFinder (Krulwich, 1997).

3.2.2.4 Knowledge-based
Knowledge-based recommender systems make recommendations based on some kind of inference. These systems have knowledge about how a particular item meets a particular user need. This means that it can take a conclusion about the relationship between the need and a possible recommendation (R. Burke, 2002). A simple example is a search query on Google where the user types in a question and Google delivers the most relevant recommendations with a solution to the question, but it can take on many more forms.

3.2.2.5 Community-based
Community-based recommender systems are using the main concept of relationships between users (Ben-Shimon et al., 2007). These systems exploit relationships and trust between users to make better recommendations, as users tend to trust recommendations from friends or friends from friends more (Sinha & Swearingen, 2001)

3.2.2.6 Hybrid recommender systems
Hybrid recommender systems consist of (a) combination(s) of the above-mentioned types of systems. Burke (2007) gives a good description about hybrid recommender systems but it is not really worthwhile to dig much deeper into this type of recommendation system as the main types of recommender systems are already explained (R. Burke, 2007).

3.3 Collaborative filtering
As stated in the previous part, collaborative filtering uses the opinions of other people to recommend items to users. In this part the collaborative filtering technique for making recommendations will be explained with more detail. Namely the process and methodology of the collaborative filtering method will be explained without elaborating too much on the technical side to make sure that everything will be easily understandable.

There are three different types of collaborative filtering, which are memory based-, model based-, and hybrid collaborative filtering.

3.3.1 Memory based collaborative filtering
Memory based collaborative filtering can be split up into two of the most traditional methods. One is user-user collaborative filtering, which was the first of the (automated) collaborative filtering methods and introduced in a paper about an article recommender (Resnick, Iacovou, Suchak, Bergstrom, & Riedl, 1994).

3.3.1.1 User-user collaborative filtering
User-user collaborative filtering finds other users with a similar past rating behavior (a so called "neighborhood) and then uses these patterns to predict what you (as the current user) will like (Ekstrand, Riedl, & Konstan, 2011). The user-user collaborative filtering method unfortunately has some drawbacks, which are mainly focused on sparsity and scalability (J. Ben Schafer, Konstan, & Riedl, 2001). The sparsity problems are caused by the lack of historic data (because a small portion of the available products have already been bought) and therefore the system is unable to make accurate recommendations. Scalability problems exist because the amount of computations needed for the recommendations grow linearly, which can cause the system to have a hard time to process everything (Sarwar, Karypis, Konstan, & Riedl, 2001).

3.3.1.2 Item-item collaborative filtering
This leads us to the other traditional method which tries to fix one of the drawbacks that user-user collaborative filtering has. This other method is item-item collaborative filtering, which was first talked about in two papers (Karypis, 2001; Sarwar et al., 2001). This method uses the rating patterns of items instead of the rating behavior. It identifies other users who like and dislike the same items. Then the system finds items that already have been rated (positively) by those users but not by the current user and then recommend these items to the current user (Sarwar et al., 2001). With this method, the scalability drawback is addressed because it does not need to identify the "neighborhood" of the active user and therefore makes recommendations at a much faster rate (Sarwar et al., 2001). There are many more "perfections" of these traditional methods, but as this thesis concerns some specific algorithms, more detail on these perfections will not be given.

3.3.2 Model-based collaborative filtering
Model-based collaborative filtering makes recommendations with a model of user ratings which is developed beforehand. These models are built by different machine learning algorithms, which are for example a Bayesian network, clustering, and rule-based approaches (Breese, Heckerman, & Kadie, 1998). These machine learning algorithms are used on training data, which are offline datasets used to "prepare" the model for real life data. Model-based collaborative filtering is mostly used in combination with memory-based collaborative filtering to limit the drawbacks of the system.

3.3.3 Hybrid collaborative filtering
At last, hybrid collaborative filtering systems are comprised of a mix of both memory- and model based collaborative filtering systems or even a mix of a collaborative filtering system and another recommender system type (which are discussed in the previous part) (R. D. Burke, 2002). These systems are used to combine the best of both worlds and eliminate the problems that both different approaches entail (Ghazanfar, Prügel-Bennett, & Szedmak, 2012). The drawback of hybrid systems is that they are very expensive and complex.

3.3.4 The methodology of collaborative filtering
After a short introduction to collaborative filtering methods this part will be concluded by briefly explaining the methodology of a collaborative filtering system to get to understand what the system actually does.

As already mentioned before, the goal of collaborative filtering is to recommend new items to users. It does that by using the opinion of other people with similar taste. A paper by Sarwar, Karypis, Konstan, and Riedl (2001) describes the general process of a collaborative filtering model very well. They describe that a typical collaborative filtering system has a list of users and a list of items where the user has expressed his/her opinions about (they indicate whether they liked those items or disliked those items). These opinions can be formatted as a rating score (numerical scale) or in the form of purchase records, watch time, clicks, etc. One user from the list is distinguished as the "active user", which is the person who the algorithm is finding an item for that can be a prediction, a numerical value which expresses the predicted likeliness of an item for the active user, or a recommendation, which is a list of items that the active user will like the most. This cannot be an item that the user already purchased, watched, or listened to. These two things are called the output of the algorithm (Sarwar et al., 2001). What the output is, differs from algorithm to algorithm and the platform it is used on.

3.4 The YouTube Algorithm

The previous parts of this thesis are all discussing recommendation systems in a very general manner. Because this thesis is focusing on the recommendation systems of social media networks, it is wise to take an example of a recommendation system in a social media network. There is a paper that was published in 2010 and talked about the new and improved recommendation system of YouTube (James et al., 2010). Because the paper explains the algorithm very well without getting too technical, it is the perfect example for this thesis. However, caution is advised as to whether the authors of the paper (all google employees) are presenting biased information or even withholding several key features of the algorithm. Taking these things in consideration, the paper explains the concept of the YouTube recommendation system and discusses the inputs and outputs of the algorithm, which can give the reader an idea of the basic ideas behind the system.

The recommendation system of YouTube is designed to provide personalized recommendations that will help users find videos that are of high quality and relevant to their interests. They do this by striving to recommend videos that are recent and fresh, as well as diverse and relevant to the user's preferences. The list (or set) of recommendations in the form of videos is generated by the user's past activity. This set is then ranked by the system based on relevance and diversity.

User activity is based on two sources of data, which are content data (like video streams and video metadata, like title, description, genre, etc.) and user activity data. User activity data can be split into two categories, which are implicit and explicit activities. Explicit activities are when users actively perform actions like rating a video (liking or disliking) or subscribing to an uploader. Implicit activities are when users are watching or interacting with a video. This type of activities can include for example the watch time of a video (how long the user watches the video).

Parallel to this process there is another process, which compiles a set of related videos based on the video that the active user is watching. In the YouTube algorithm they use a technique which is called association rule mining (Agrawal, Imieliński, & Swami, 1993). This process considers the count of how often a pair of videos are watched within a given period of time, which are called co-visitation counts. Then based on that they make a set of related videos which is used for the next step of the recommendation system. One little side note is that sometimes videos have a too low overall view count and therefore also a low amount of co-visitation counts, which means that a reliable related video set cannot be computed.

In the next step, which is called the candidate generation, the user activity and the related video process are combined. The union of these videos is called the "seed set". From this seed set the candidate recommendations are formed. To make sure that the candidate recommendations are diverse enough there is a system in place which includes videos in the candidate recommendations set that are new to the user.

The last step of the recommendation system is the ranking of the candidate recommendations set and finally making a list of videos which are going to be recommended to the active user. The videos from the candidate recommendations set are ranked with the use of three categories of signals, which are video quality, user specificity, and diversification. Video quality is based on how positive users are on that video. This is computed by taking view count, favorites, likes, sharing activity and so on in account. User specificity is used to give videos which match the user's preferences a higher ranking, which is based on the user's historical activity. Lastly, the subset of videos is optimized by balancing relevancy with diversity. Videos that are too similar are removed from the subset. The paper explains a way where the system puts constraints on videos recommended from the same uploader.

Of course, there are many complicated techniques in use to make the recommendation system more efficient, but this overview gives a good general idea about how the system works and how YouTube generates the recommended videos for you every time you watch a video.

4. CHALLENGES OF RECOMMENDER SYSTEMS

In this part of the thesis, the challenges concerning recommender systems will be explained. The challenges will be derived from literature and will be explained by using the theory from the first part of this thesis. The explanations in this part will be limited to describing the phenomenon and putting the phenomenon into context by explaining how that would work in the algorithm. Then, for some of the challenges, a few references to research that has been already done will be given to verify that the phenomenon exists within recommender systems and thus poses a problem. The next part will explain why it is a problem and what implications and negative consequences can emerge from them.

4.1 Content diversity

As already stated in this thesis, the goal of recommender systems is to recommend items to the user that will be of their preference. Its aim therefore is to analyze the load of information that certain platforms have to offer and reduce it to a certain set of items which the user will like. A study by Goldman (2008) confirms that these recommender systems do have a positive impact on content diversity, as they promote new or novel items which the user may like (Goldman, 2008). However, according to a paper by Fleder and Hosanager (2008), most recommender systems also have a tendency to keep recommending items from the same category, which in turn decreases the diversity of content over time (Fleder & Hosanagar, 2008). This phenomenon is called concentration bias (Adamopoulos & Tuzhilin, 2014). When a recommender system contains an inclination bias, it has an inclination or prejudice for a certain type of content and therefore is unfair against other categories of content or items. Although the goal of a recommender system is not to be fair and have a perfect balance between different types of content, as this would damage the functionality of these systems, there should be a good and balanced mix between relevancy and diversity of the content recommended (Vargas & Castells, 2011).

Many researchers have tackled the challenge of content diversity concerning the use of recommender systems. Bradley and Smith where the first to mention the balance between item relevance and diversity. Fleder and Hosanagar (2008) experimented with different kinds of recommender system and found out that most of them caused a decrease in diversity because they focus too much on the accuracy of the recommendations (Fleder & Hosanagar, 2008) and Nguyen et al., (2014) conclude that there is a narrowing effect present in collaborative filtering algorithms (Nguyen, Hui, Harper, Terveen, & Konstan, 2014). Jannach, Lerche, Kamehkhosh, and Jugovac experimented with recommender systems by testing them on diversity and found a phenomenon which is called popularity bias (Jannach, Lerche, Kamehkhosh, & Jugovac, 2015). Popularity bias is present in a recommender system when the system keeps recommending content that is already popular and therefore has an inclination or prejudice to popular items, which in turn causes an unfair inclination or prejudice to less popular or long-tail end items (Anderson, 2006). Lastly, Kunaver and Požrl (2017) provide an overview of the evolution of experiments done on recommender systems concerning content diversity which confirms that it is currently still a challenge (Kunaver & Požrl, 2017).

4.2 Personalization

Personalization in recommender systems goes further than "normal" recommending as it uses historic data to provide more personalized recommendations. It creates a user profile, whether the user profile is an actual profile like your own Facebook page or YouTube account, or whether it traces your actions as patterns like web cookies do (Zhou, Xu, Li, Josang, & Cox, 2012). The personalization function uses both explicit and implicit inputs to make sure that the profile is well put together. Explicit inputs are the inputs that users make themselves like rating, subscribing or just filling in their preferences. Implicit inputs provide data that is gathered passively by "observing" the user. This system of user profiling greatly improves the effectiveness of the recommender system, but it comes, most of the time, with some problems.

A paper by Zhou et al., (2012) gives an overview of the issues with the user profiling process, described above, and they even suggest that the user profiling is the weakest link in the system (Zhou et al., 2012). They mostly discuss technical issues, which will be explained in more detail with the next challenge (chapter 4.3). These technical challenges, which are data problems, are called the cold start problem and the problem of data sparseness, which both emerge because of a lack of (accurate) data (Schein, Popescul, Ungar, & Pennock, 2002). Other problems are touched on by Morita and Shinoda, who suggest that while explicit input data are more accurate to create recommendations, it is harder to gather because users are more aware of their actions in contrast to implicit gathered data. A paper by Nichols (1998) also concludes that implicit data is less accurate than explicit data (M. Nichols, 1998). At last, Burke, Mobasher, Williams, and Bhaumik (2006) discuss the problem of malicious rating, which is a technique used by companies or organizations to "attack" the content of their competitors and promote their content instead (R. Burke, Mobasher, Williams, & Bhaumik, 2006). The problems seem to touch on the same problems seen at the previous challenge, which makes it look like it may be a negative moderator to content diversity.

4.3 Data problems

The two problems that are going to be discussed are already mentioned in the previous part. These two problems are called the problem of data sparseness, and the cold start problem, which both concern a lack of data. The problems will be discussed separate of each other.

4.3.1 Data sparseness

Recommender systems rely on historic data from users which are most of the time formatted in ratings. The problem of data sparseness emerges because a user in the recommender system only rates a small portion of the available items on the platform. Data sparsity may cause unreliable recommendations made to users, which in turn can decrease user satisfactions (Grčar, Mladenič, Fortuna, & Grobelnik, 2006).

4.3.2 The cold start problem

The cold start problem faces the same problem as the data sparseness problem, namely a lack of data. The cold start problem can be divided into three categories, which are recommendations for new communities, new users, and new items (Bobadilla, Ortega, Hernando, & Bernal, 2012; Lika, Kolomvatsos, & Hadjiefthymiades, 2014).

The new community problem occurs with starting social media platforms when there is not a sufficient amount of data available to make reliable recommendations as there are not enough users or items on the platform yet (Schein et al., 2002). As this thesis covers already established social media platforms, this problem is less relevant than the other two.

The other two problems are called the new user problem and the new item problem, which both face the same challenge but for a different subject. The new user problem occurs when a new user opens a profile or visits the platform and thus has not yet given any ratings or established any preferences (Rashid et al., 2002). The first set of recommendation therefore are, most of the time, not very accurate which can in turn decrease user satisfaction (Bobadilla et al., 2012). The new item problem emerges in the same fashion. When a new item enters the system, it does not possess any ratings yet, which results in the item not getting recommended easily and thus noticed by users (Park & Tuzhilin, 2008). The item may end up directly in the long-tail (Anderson, 2006)

Both problems can prevent the recommender systems from working optimally as they influence the accuracy of the recommendations which can cause negative consequences. Therefore, these data problems can be seen as a challenge for recommender systems.

4.4 Metrics

Every social media platform has its own different recommender system with their corresponding algorithms. The creators of these systems therefore influence how the algorithm should fundamentally behave to tailor the system to the likings of the social media platform. Putting it into the context of social media platforms, the creators of recommender systems can impose certain rules on the system that determine whether or not content will be recommended. Most of these rules are used to determine which content will be popular, and thus are named popularity metrics. A simple example of a metric is the number of views on a video (Szabo & Huberman, 2010). where a higher number of views corresponds with a higher chance on popularity. Social media platforms, most of the time, do not disclose the exact metrics for determining popularity to avoid possible manipulation of the algorithm (Tatar, de Amorim, Fdida, & Antoniadis, 2014)

However, researchers have discovered metrics on different platforms that correlate with content popularity. Chatzopoulou, Sheng and Faloutsos (2010) found four key popularity metrics that are highly correlated with video popularity on YouTube (Chatzopoulou, Sheng, & Faloutsos, 2010). A paper by Ma, Sun, and Cong (2013) discusses the factors that play a role for a hashtag to become the Trending Topic, which is a much talked about hashtag (ma, Sun, & Cong, 2013). Lastly, a paper by Moro, Rita, and Vala contains research which examined the factors that contribute to the popularity of Facebook posts and how this could predict the popularity of posts (Moro, Rita, & Vala, 2016). Furthermore, certain tools, which are most of the time offered by the platform itself, allow content creators to analyze viewer behavior. This can lead to the content creators identifying some of the popularity metrics by themselves. Three of the biggest social media platforms, Facebook, YouTube, and Twitter, offer this type of tools which can gather data and perform various forms of analysis on your content (Facebook, n.d.; Twitter, n.d.; YouTube, n.d.)

The challenge lies in the fact that discovery of these metrics could lead to manipulation of the algorithm and therefore disturb the natural working of the recommender system. Research done on these metrics and the tools available to essentially discover these metrics confirm that they can be identified and therefore can become a weak link in the system.

4.5 External influences

The algorithms in recommender systems are not the only things that are faced with challenges. External factors may as well exercise influence on the recommender system and therefore pose a challenge against the recommender system. In this part four different "groups" of people will be discussed which all pose a challenge against the well-functioning of recommender system. These groups are: human operators, the government, and third parties like advertisers and companies.

4.5.1 Human operators

Firstly, human operators, which are the people who monitor content on a social media platform, can influence both recommender systems and content on the respective social media platform. This phenomenon is called gatekeeping, which found its origin in reporting and journalism (Groshek & Tandoc, 2017). The human operator is empowered with the ability to remove or temporarily block content if the content is not conformed to the guidelines that the social media platform has set. This results in the content on the platform being subject to human bias and judgement.

A recent example of human operators removing content on social media, or being caught in the act of "censoring" the platform is an article from Wilson on Christian Today (2019), where Facebook is acting as an "enforcer" and flagging (marking the content as not within the guidelines of Facebook) or removing content from the platform (Wilson, 2019). Especially when topics regarding religion or politics, which are very sensitive, come into discussion, content cannot be censored or removed that easily without causing an outrage.

The other side of the problem, namely the power of human operators to block content, is discussed in an article by Thielman (2016) published on the website of The Guardian (Thielman, 2016). It involves a document created by Facebook where they wrote down the guidelines for how a topic is going to be a trending topic. Furthermore, it states that there is a team which select or blocks certain trending topics, which confirms that the trending topic page of Facebook is subject to human bias.

The two articles display the influence the human operator has on the content on social media platforms and thus also the recommender system. This confirms that there is a challenge present where a balance needs to be struck between passively monitoring and active censoring.

4.5.2 Governmental interference

The government can also interfere with the natural working of the recommender system. The challenge of governmental interference is twofold. The first problem emerges when the government takes over the role of the human operator and engages in influencing the content available on social media platforms. In some countries, like China, the government has very close control on what content will be displayed on the platform and what content will be instantly removes (MacKinnon, 2011). A paper by Faris and Villeneuve (2008) gives an overview of which countries are using some form of internet filtering , which gives an idea of how many countries the government has influence on the content published on, for example, social media (Faris & Villeneuve, 2008). Furthermore, Deibert and Rohozinski (2010) explain that the tools that are used to control and filter content in the cyberspace are becoming increasingly stricter, which possibly will result in more content removed and filtered (Deibert & Rohozinski, 2010).

The other side of the problem is concerned with e-government. Governmental instances are using social media as a tool in an attempt to reach out better to citizens. They might use social media to campaign, acquire feedback, convey messages, or even warn users in, for example events of extreme weather (Kavanaugh et al., 2012). Bertot, Jaeger, and Grimes (2010) also discuss social media as a tool to promote openness and transparency between the government and citizens (Bertot, Jaeger, & Grimes, 2010). However, the adoption of social media by the government can also have adverse effects, as the government can use social media tools to propagate bad content. King, Pan, and Roberts (2017) show that China fabricates social media posts for strategic distraction (King, Pan, & Roberts, 2017) and an article by El-Khalili (2013) discusses the ways in which the government of Egypt abused social media platforms to propagate their views during the revolution that started on Janurary the 25th of 2011 (EL-Khalili, 2013).

Both problems confirm that governmental interference can be a challenge for the recommender system as the government can influence content on social media platforms and therefore influence the recommender system.

4.5.3 Advertisers

As most social media platforms have a business model, they are generating revenue in some way. Most of the time this revenue is generated from advertisers, who can display their ads on these platforms in exchange for monetary compensation. Furthermore, content creators can also promote their content through ads.

On the website of Feedough, the business models of YouTube, Facebook, and Instagram are explained, which gives an insight in the main sources of income for the respective platforms (Das, 2019; Dutta, 2019; Pahwa, 2019). The website mentions that the three platforms have advertisements as their biggest contributor to their revenue. Furthermore, the platforms offer options for advertising for both advertisers and content creators. Advertisements can be seen before or while watching content, but they can also show up in the part where the user can usually find his or her recommendations, which means that they can indirectly influence the recommender system. Content creators on Facebook and Twitter have the option to temporarily "boost" their content which promotes the content and therefore has more chance to show up in the recommendations section of the user. All the options for advertisements can be found on the websites of the respective platforms ("Ads, how it works," n.d.; "Advertising on Twitter," n.d.; "Facebook-ads," n.d.).

The combination of the options for advertisers to display their advertisements in the recommendations sections of social media platforms and the ability of content creators to promote their content to get featured more proves that advertisement on social media platform forms a challenge for recommender systems. As with the other challenges in this part, it disturbs the natural working of the recommender systems by being influenced by an external factor.

4.5.4 Companies

To conclude this part about external influences on recommender systems some explanation will be given on other third parties as gatekeepers for social media content. With third parties, big companies, especially media companies are meant. A paper by Bauer (2015) examines the issues that come forward when talking about copyright on social media platforms. The writer talks about how social media provides a platform for creativity and inspiration, but this comes at the cost of people who are trying to abuse the openness and velocity of these platforms. The writer describes several cases where creators of content (photographers, music artists, etc.) are suing uploaders of content over huge sums of money because of copyright infringement (Bauer, 2015).

Cases like these cause the fact that owners of the original content are aggressively making sure that people are not using their content without permission which can lead to excessive or false claims of copyright, which is described very well in an article on Forbes by Sands (2019) which studies copyright claims on YouTube (Sands, 2018). This in turn causes certain videos to be taken down due to possible copyright infringement and therefore influences the recommender system as the content on the platform is "attacked". The challenge here lies in the fact that social media platforms need to react adequately on these possibly false claims to make sure that people with bad intentions do not get content removed when it should not be removed.

Table 1: Overview of the challenges

Challenge	Description
Content diversity	Users of social media platforms might get exposed to a narrower range of content over time.
Personalization	User accounts, which use historic data and patterns, may allow the system to over-personalize the recommendations
Data Problems	Data problems present in recommender systems may cause a decrease in the accuracy of the recommendations presented to the user
Metrics	Metrics, used to determine which content will be recommended, can be identified and exploited
External Influences	Human operators, the government, advertisers, and (big) companies, may be able to influence the recommender system

5. IMPLICATIONS OF RECOMMENDER SYSTEMS

In the last chapter of this thesis, the challenges mentioned in the previous chapter will be translated to possible, negative consequences. The consequences will be explained by executing a case-based analysis. This analysis will partially verify that the challenges mentioned in the previous chapter exist in recommender systems used on social media platforms. Furthermore, the analysis will provide context to the consequences which makes it more understandable. To conclude the part recommendations will be given on how to avoid or deal with the negative consequences emerging from the challenges.

5.1 Case#1: The 2016 United States presidential election

5.1.1 Case description

On Tuesday, November 8, of 2016 the 58th United States presidential elections took place. The candidates from the republican and democratic party respectively were Donald Trump and Hillary Clinton. Eventually, Donald Trump won the elections and took office as the 45th president of the United States on the 20nd of January 2017.

During these elections, social media platforms were heavily used as a tool to campaign for the candidates. Platforms such as YouTube, Facebook, Twitter, Instagram and Snapchat were all used by the campaign team of the candidates to gather votes or connect with existing voters. However, the heavy use of social media during the elections means that the candidates will be subject to the bias present in social media platforms, and thus also the recommender systems incorporated in these platforms. During and after the elections the problem of content diversity came to light, which might have influenced voters during the elections.

Firstly, an article from Lang, published on Government Technology (2016) shows that Trump did get more attention by uploading appealing content during the elections which got liked and shared and therefore became trending on social media platforms (Lang, 2016). Another paper from Alcott and Gentzkow (2017) discusses the influence of fake news on the popularity of the presidential candidates and concludes that it did have some influence on the popularity or presence of the presidential candidates (Allcott & Gentzkow, 2017). Lastly, an article by El-Bermawy published on Wired (2016) gives an empirical example on how a user of a social media platform can get stuck in a so-called filter bubble (Pariser, 2011), which means that everything that is recommended for you agrees with your opinion and there is no or little content that disagrees with your opinion. He explains that, while he is a Clinton supporter, his YouTube recommendations are swamped with pro Clinton videos and liberal content. He also explains that there is no or little content that covers Trump or right content, which means that he is essentially stuck in a filter bubble (El-Bermawy, 2016).

5.1.2 Consequences of a lack of content diversity

The challenge of content diversity is clearly visible in this case. The articles from Lang (2016), Alcott, and Gentzkow (2017) prove that popularity bias is present in this case. The presence of popularity bias in this case and in general has two big implications.

The first problem concerns the vicious circle that emerges as popular content will get recommended, which in turn leads to the content becoming more popular, which in turn leads to the already popular content being recommended to even more users. This may lead to less popular, but relevant, content being left out of the recommender system which decreases the content diversity and may possibly influence the opinion of the user. The second problem builds on the first problem as users may risk being trapped in a filter bubble after they watched some popular content, promoted through the recommender system, on a social media platform. The filter bubble is a figurative "place" where users will mostly see content that agrees with one opinion and does not show content that displays opinion from the other "side". These two problems together may cause the user to be influenced by the recommender system which can play a part in shaping the opinion of the user.

The worst-case scenario would be that recommender systems used on social media platforms do play a part in influencing and shaping the opinion of the users. This would mean that the US elections of 2016 could have been influenced through popularity bias present in the system. Furthermore, more events could have been influenced in the same way.

5.1.3 Recommendations

Prevention is always better than letting it happen and fixing it afterwards. Therefore, the responsibility for prevention lies with the creators of the recommender systems and the people who monitor the content on social media platforms. The creators can look at implementing functions which make sure that content diversity is maintained, and long-tail items are promoted to the user to avoid popularity bias or the filter bubble effect. They can also promote algorithm transparency to make the user more well-informed. The human operators can avoid these consequences by carefully monitoring the platform, especially during important events, but without interfering too much, as human bias can be present too.

If the bias is already present in the system, the responsibility lies with the user to limit the damage the bias can do. The user needs to be well informed of the possible lack of content diversity that may exist in the system and needs to act proactively. They can do this by actively searching for the other “side” of the story and therefore making sure that the algorithm also promotes other opinions than only the popular ones. Wijnhoven and Brinkhuis (2015) discuss internet information triangulation tools which may assists users in achieving the goal of finding contradictory statements (Wijnhoven & Brinkhuis, 2015).

5.2 Case #2: The Rabbit Hole effect

5.2.1 Case description

The Rabbit Hole effect is a phenomenon where users are exposed to incrementally more extreme content, as the recommender system in the social media platform adjusts itself to the preferences of the user and wants to recommend new content from the same category. Because this phenomenon is not present in one big event, several articles discussing the subject will be presented. An article from Albright (2018) published on Medium displays a “walk down the rabbit hole” as the writer starts at a video on YouTube about the high school shooting in Florida (Laughland, Luscombe, & Yuhas, 2018) and ends up at many controversial and violent videos (Albright, 2018). Another article discusses how watching videos on Donald Trump and Hillary Clinton causes increasingly either far right or far-left videos to be recommended (Tufekci, 2018b).

5.2.2 Consequences of over-personalization

The articles discussed in the case description all cover the challenge of personalization in recommender systems. Personalization, or over-personalization in this case, comes with some negative by products which can lead to negative consequences. The personalization features in recommender systems can lead to the content narrowing down too much and therefore decreases the content diversity of the recommendations offered to the user. Furthermore, the “down the rabbit hole effect”, may cause the recommendations of certain user to become increasingly more extreme as the system needs to make new recommendations to the user while the user profile is becoming more personalized.

Worst-case scenarios would be that the recommender systems used on social media platforms expose the user to extreme, disturbing, or in general bad content and may cause them to become trapped in a filter bubble through over-personalized user profiles. This could eventually influence the opinions of users and may cause them to form a false image of reality, as extremist or conspiracy content will get recommended to them. As already stated in the introduction, users of social media platforms start their profiles at a very young age, at which they are very easily influenced, which can cause a big problem in this case.

5.2.3 Possible recommendation

The consequences of over-personalization in recommender systems can possibly be prevented by the creators of these systems and the people who monitor them. As with the previous case, content diversity as well as algorithm transparency can be promoted through the system to prevent users from being trapped in a filter bubble and bad content can be removed from the platform to prevent users from being able to be exposed to this type of content.

However, if these people are not able to prevent the problem from happening, users should be informed about the problem and stay critical of the content they are being recommended. Especially in this era where fake news is more common than you think, users should fact check content to make sure they do not just believe some opinion offered to them (Tufekci, 2018a). Lastly, parents should monitor young kids on what they are watching to make sure they do not accidentally stumble on bad content. YouTube already acted on the behalf of kids to solve that part of the problem with a platform called YouTube kids (Alba, 2015).

5.3 Case #3: The consequences of data problems

5.3.1 Description

Data problems in recommender systems cause a decrease in accuracy of recommendations to new users and new items. This decrease in accuracy can open the user up to, almost, all the consequences mentioned in this chapter, therefore no specific case can be given for data problems, but the magnitude of the consequences are worth mentioning.

5.3.2 Consequences of a lack of accuracy

The consequences of these data problems are that users will sometimes, especially in the beginning, get inaccurate recommendations. This can be innocent but can also prove dangerous as, if a user does not qualify yet for personalized recommendations, they will probably get the most popular videos recommended, which can be influenced by factors mentioned in the previous cases like popularity bias, external influences, algorithm exploiters, etc. Besides this, items that do not have enough ratings to be recommended may experience the problem of ending in the long-tail end of content, which means that it will be hard for those items to be recommended to the user and therefore have little chance to become popular. This can mean that valuable contradictory content can get “skipped” by the recommender system whereas the other side of the content keeps getting recommended.

5.3.3 Recommendations

Several researchers have already tried to tackle the cold start problem and the problem of data sparsity. However, there is no evidence of recommender systems used on social media platforms have found a permanent solution to this problem. This means that the creators of these systems need to keep developing their systems to make sure to minimize the technical difficulties that can be present. Emphasis should also be on incorporating a mechanism that also gives long-tail content a chance in the recommender system to make sure that every bit of content gets included in the process. Beside this, especially new users need to make sure that they remain critical when watching recommended items and not just watch content because it is popular and recommended to them, especially when they create a profile or visit the platform for the first time.

5.4 Case #4: Exploiting the recommender system

5.4.1 Case description

On the internet, several guides can be found which explain step by step how you can understand how the algorithm, or the recommendation system works to use this knowledge to make your content appear in the recommendations or on the trending page. An article by Gielen, published on Tubular Insights (2016) explains several different factors which increase the chance of your video being recommended. Factors include video length, title choice, description choice, keywords, tags, etc. (Gielen, 2016). The same type of article is published on BuzzSumo by Moeller (2019) and explains how you can get more people engaged with your Facebook post. She made this guide by analyzing a huge amount of Facebook posts and therefore analyzing the algorithm to come up with the factors that determine how a post gets more likes and thus becomes more popular (Moeller, 2019). Lastly, an article by Narang on Socialalert Blog (2017) discusses how to get your hashtag to become a trending topic in less than 48 hours. The writer states that you need to understand the algorithm to take advantage of it and get your hashtag to the trending topic page (Narang, 2017).

5.4.2 Consequences of algorithm exploiting

As can be seen from the case, there are "guides" which give a step by step plan on how to make your content more popular on social media platforms by effectively exploiting the algorithm. Large studies, as can be seen from the article about Facebook from the case, can identify patterns in content popularity. Together with the performance metrics tools, like discussed in chapter 4.4, provided by social media platforms it opens opportunities for content creators to identify more popularity metrics.

When these metrics are discovered, content creators can exploit this information by tailoring their content exactly to match these metrics and therefore gain popularity in an unfair manner. Content creators can therefore partially influence the natural working of the recommender system. Negative consequences will then emerge as content creators, who want to push bad content, get a hold of these metrics and then are able to achieve this goal. This would especially be terrible in the case of, for example, the event of elections, terrorist organizations trying to radicalize people, hooligans trying to gather a mob, etc.

5.4.3 Possible recommendation

In contrast to the previous two cases, the algorithm is not to blame here, but the creator. The creator develops the rules incorporated in the recommender system, and therefore needs to make sure that these rules are not easily identified by content creators. They also need to make sure that these rules or metrics are, kind of, randomized to prevent big studies from analyzing patterns in video popularity, at least concerning the algorithm.

The social media user can also be critical of what gets recommended to them, especially when something gets recommended that does not really fit with their viewing history.

5.5 Case #5: The aftermath of the Egyptian revolution of 2011

5.5.1 Case description

The Egyptian revolution of 2011, starting at January the 25th of 2011 and ending on the 11th of February 2011, was an event where millions of protestors demanded the overthrow of Egyptian President Hosni Mubarak. The revolution was sparked by a vlog from a woman called Asmaa Mahfouz, where she urged protesters to protest on Tahrir Square. This vlog was shared everywhere on the internet, especially via social media and therefore social media became the tool for starting a revolution. The revolution ended with the resigning of President Mubarak and a horrific number of casualties.

However, this was not the end. A paper by El-Khalili (2013) discusses how the government used social media in post-revolution Egypt as a tool for propaganda (EL-Khalili, 2013). They created official accounts as communication tools between the government and the citizens which were promoted by anonymous people, employed by the government, to like these posts and increase their popularity. Furthermore, a few articles state that the Egyptian government began censoring the internet, by removing posts, penalizing uploaders and even jailing some people (Cappon, 2019; "Egypt to regulate popular social media users," 2018; Toulas, 2019). The events from the Egyptian revolution sparked the importance of social media as was seen in the increase of Egyptian Facebook users after the revolution , but the government apparently thought of this too and reacted harshly ("Egypt's Facebook users double: Ministerial report," 2012).

5.5.2 Consequences of external influences

The case shows clearly the challenge of external influences on the content offered on social media platforms and thus also on the recommender system. They were able to both censor content on these platforms and propagate through the same channels. As discussed in chapter 4.5, human operators, advertisers, and companies also share these characteristics of having the power of either censoring the platform or propagating through it or even both. The goal of social media platforms is to remain open and transparent, which is violated in this case and in every case where external influences disturb the natural working of the recommender system. The negative consequences will be that users expect this openness and transparency to be maintained whereas in the meantime they are being "fed" content that is heavily influenced by outside factors and therefore these users are also being influenced in their opinions and views without them even knowing.

5.5.3 Possible recommendation

The external influences described in chapter 4.5 can partially be avoided. The people behind social media platforms can take a critical look how much options for advertisements they make available to ensure openness and transparency. They also need to assess copyright claims more critically to avoid malicious copyright striking (Wodinsky, 2019). The government influence, especially in the case of Egypt is harder to avoid, as social media companies do not have the theoretical firepower to go against the government when they issue orders.

When the system is influenced by external parties, users need to be informed and stay critical to spot these influences. Users should actively pursue openness and transparency on the platforms and need not to be afraid to let them hear their voice when they think that these values are compromised. Especially in the case of government influence, users need to empower themselves and others and find other ways to express their opinions.

Table 2: Overview of the consequences and recommendations per challenge

Challenge	Consequences	Recommendations
Content diversity	- Most recommender systems cause a decrease of content diversity over time - Popular content keeps getting more popular - Users might become trapped in a filter bubble which filters out contradictory content - *WCS: recommender system might influence the opinion of users through a lack of content diversity*	- **The creators of the recommender system** need to make sure that the algorithm promotes content diversity and takes care of the long-tail problem. Algorithm transparency should also be promoted. - **Operators** need to monitor content diversity during big events and interfere if content diversity is not maintained. - **Users** need to be proactive and stay critical of the items they get recommended especially in times of big events. They need to actively search for contradictory content if diversity is absent. Possibly with assistance of tools (Wijnhoven & Brinkhuis, 2015).
Personalization	- Over-personalization in recommender systems - (Young) users might be exposed to harmful content - Over-personalized content might lead to a decrease in content diversity which can cause user to become trapped in a filter bubble - *WCS: recommender system might influence the opinion and thoughts of users through over-personalized recommendations.*	- **The creators of the recommender system** need to control how much bad content gets promoted to users and how narrow the recommendations provided to the users will be over time. Algorithm transparency should also be promoted. - **Operators** need to monitor the platform for bad content and remove it if possible. - **Users** need to be more proactive and actively search for contradictory statements when diversity is absent. Users also need to fact check content if the subject seems doubtful. Possibly with assistance of tools (Wijnhoven & Brinkhuis, 2015). - **Social media platforms** need to follow the example of YouTube concerning their kid-friendly function
Data problems	- Lack of accuracy through data problems within the algorithm - Less popular content might not be considered as a recommendation by the system. - *WCS: non-personalized recommendations might lead to the user being exposed to the most popular content at the moment which in turn may be subject to the other consequences mentioned in this table*	- **The creators of the RS** need to keep developing methods to compensate for the lack of data and decrease in accuracy. - **Users** need to remain critical of the content that gets recommended to them, especially when creating a profile or visiting the social media platform for the first time.
Metrics	- Upon discovery of these metrics, content creators can exploit the recommender system. - People who possess this kind of information are able to influence the popularity of their content and thus also the recommendations. - *WCS: content creators might use these metrics to push harmful content to the users. People can also use these metrics to influence users during big events like elections.*	- **The creators of the recommender system** need to make sure that the “rules” incorporated in the RS are not easily identifiable. - **Users** need to be critical of the content they are getting recommended, especially during big events. They also need to be critical when content is being recommended that does not fit with their viewing history.
External influences	- Several parties (human operators, government, advertisers, and companies) can exercise influence on the content provided by the social media platform. - These influences can disturb the natural working of the recommender system. - *WCS: the goals of openness and transparency might be replaced with corporate, promoted, or ordered content which in turn might influence the opinion and thoughts of users Furthermore, user opinions could be suppressed and censored.*	- **Social media platforms** need to have a critical look at how much space there is available for people to promote content on their platform and if it strikes an even balance with “normal” user content. - **Human operators** need to assess copyright claims more critically to avoid malicious copyright striking. - **Users** need to stay critical of content that gets recommended and does not fit with their viewing history. When faced with governmental influence, users can move to another platform to find other ways to express their feelings.

6. CONCLUSION

To conclude this thesis, the research question stated in the beginning needs to be answered. The research question is: "What are the negative consequences of recommender systems used in social media networks and how can people avoid or deal with them?". To answer this research question, the subquestions stated under the research question will be answered one by one. The first subquestion is concerned with the challenges that come with recommender systems used in social media networks. Chapter 4 gives an overview of the challenges that limit the recommender system from working optimally, which are content diversity, personalization, data problems, metrics, and external influences (also summarized in table 1)

Negative consequences emerge from these challenges, which are analyzed by a case-based analysis in chapter 5. The negative consequences and worst-case scenarios can be grouped into three categories, which are the users of social media platforms and the creators of the system. The main consequence, for all the challenges, for social media users is that they can be influenced by the recommender system. Through a lack of content diversity, over-personalization, algorithm exploiting, and external influences, they can be exposed to content that may create an opinion for them, influence their opinion, or make them doubt their own opinion. The main theme that comes back is the theme of a lack of content diversity, where certain content is pushed to the user and there is no contradictory content recommended which can influence the opinion of the user. Content creators can also push content to the user by exploiting the algorithm behind recommender systems, promoting content through advertisements, or by order from the government. Besides this, several groups of people also have the power to censor content, like the government and human operators behind the platform. Through these methods, worst-case scenarios like influence in the US elections, massive social media censoring, and radicalizing through going down the rabbit hole, can become a reality. Lastly non-accurate recommendations because of data problems present in the algorithm can also moderate the effect of the user being recommended content that is being pushed by external parties. The creators of these systems also face negative consequences, as external parties can try to either exploit the algorithm, or exercise force to make sure that the algorithm is exploited. They can be vulnerable to "attacks" on their systems which can in turn cause negative consequences for users.

Recommendations on how to avoid or deal with the negative consequences can be grouped under the factors that are to blame for causing the negative consequences. These factors are the algorithm or recommender system itself, the creators of the systems, the people monitoring the content on social media platforms, external forces that influence the content on social media platforms, and the users of the platforms. Starting with the recommender systems themselves, several challenges like a lack of diversity, over-personalization, and the data problems, cause negative consequences because the algorithm is not "smart" enough to balance everything out and make accurate recommendations while also keeping in mind to keep the content diverse and novel. The system and the algorithms behind it are not capable of thinking that way (yet!), therefore it is unfair to put the blame on the system. Because of this reason, the factors that can make a difference in avoiding or limiting the negative consequences will be discussed.

Firstly, the creators of recommender systems and the people who monitor the working of the systems and content need to be aware of what the recommender system can do. As they deploy the recommender system, they are becoming responsible for the "actions" of the algorithm and thus need to be careful with how it works. Consequences like a lack of diversity, over-personalization, exploiting of the system and too much options on promotion can be contained if the platform constantly gets monitored and tested on these exact factors. Of course, social media platforms cannot be perfectly balanced all the time, but the cases mentioned in this thesis shed light on when it does go wrong and there should have been interference. They need to actively inform users when there are problems concerning their system and promote transparency of their algorithms if negative consequences do show up.

Secondly, the users of social media platforms also need to be aware of what the recommender system can do. The recommender system provides a useful tool for giving you content that you might like so you do not have to make your way through the enormous information stream of content that the internet has to offer. This also means that you will get exposed to the possible negative consequences that come with a recommender system used on social media platforms. The main recommendation for users is that they remain critical of the content they encounter on social media platforms and put it into question if they suspect that the content, they are being recommended, is very one-sided. They should always try to check out the other side of the story too, perhaps with tools available on the internet, to make sure that they do not get trapped in filter bubbles or going down the rabbit hole.

Lastly, the creators of the recommender system, the people who monitor the system, and the users of the system, need to be aware of possible external influences that may have an impact on the natural working of the recommender system. Governmental interference, people exploiting the recommender system, human operators behind the system, companies and advertisers are all able to influence the content. Influencing content can be removing content, uploading propaganda, or paying large sums of money to promote content. This can change the composition of the content on the social media platform to such an extent that it can disturb the natural working of the recommender system and therefore needs to be considered by the parties named. For the creators of the system, they need to be aware of possible orders from the government, and companies trying to disturb their platform and interfere with it. The users need to be aware of how these external influences trying to interfere with the platform can disturb the recommender system and therefore, as with the other challenges, need to stay critical of what content they get recommended by the system.

The main conclusion to this thesis is that every party involved or influenced by a recommender system used on social media platforms needs to be aware of the possible negative consequences, which are described in this thesis, that can be present in these systems. They need to be aware of the fact that there are many factors that can play a role in disturbing the natural working of the recommender system, which can in turn can cause disastrous events in the worst-case scenarios. Recommendations from social media platforms are almost never questioned and just taken for granted when they are seen by the user. However, this thesis shows that there are challenges present in these systems and the cases that are analyzed show that consequences can emerge that are not that innocent. As mentioned in the introduction, social media is a huge part of the lives of people and kids of a very young age are already using social media, which possibly exposes a massive amount of people to these negative consequences. To conclude this, awareness and perhaps lessons in algorithm safety may be the first step into taking a more critical look at recommender systems and the possible negative consequences that they bring with them and use this to be more careful and thoughtful when using these systems on social media platforms.

7. LIMITATIONS AND RECOMMENDATIONS FOR FURTHER RESEARCH

To end this thesis, the limitations of the research and recommendations for further research will be given.

This research is subject to some limitations. The introduction of this thesis already touched on some of the limitations that this research possesses. The technical magnitude of this subject makes it hard to explain some of the concepts concerning machine learning and recommender systems without making it too complex to understand for the reader. This in turn, causes another limitation where the explanation of these technical concepts is too abstract to understand for readers who possess the technical expertise. This research tried to strike a balance between not making the concepts to complex or abstract, but this might cause some issues in readability.

The last limitation of this study is that there was a lack of academic research on the possible negative consequences of recommender systems used on social media platforms. Of course, the goal of this research was to create new knowledge about that subject, but it could have provided this thesis with more solid cases.

This research provided a basis for the limitations of recommender systems used on social media platforms. To confirm that this is also the case in reality these challenges could be tested against current recommender systems which would validate the challenges being present in these systems.

More research could also identify more challenges of recommender systems used on social media platforms, as this research could bring researcher to new insights. Furthermore, more research on the role of recommender systems used on social media platforms can be conducted, as social media remains a big player in the media world.

Lastly, further research could be done concerning the recommendations on how to avoid or deal with the negative consequences of recommender systems used on social media platforms. More concise recommendations or even guides on “algorithm safety” could be produced to better inform all the parties who are influenced by these algorithms to make sure that they can respond well to possible negative consequences.

8. ACKNOWLEDGEMENTS

I would like to thank my first supervisor, Dr. Fons Wijnhoven for providing feedback and guidance during both the bachelor thesis proposal and the bachelor thesis program itself. I would also like to thank Dr. Matthias de Visser for being my second supervisor. Lastly, I would like to thank my thesis circle for providing valuable feedback and discussions.

9. REFERENCES

Adamopoulos, P., & Tuzhilin, A. (2014). On Over-specialization and Concentration Bias of Recommendations: Probabilistic Neighborhood Selection in Collaborative Filtering Systems. *Proceedings of the 8th ACM Conference on Recommender Systems*, 153–160. https://doi.org/10.1145/2645710.2645752

Ads, how it works. (n.d.). Retrieved from YouTube.com website: https://www.youtube.com/intl/nl/ads/how-it-works/

Advertising on Twitter. (n.d.). Retrieved from Twitter Ads website: https://ads.twitter.com

Agrawal, R., Imieliński, T., & Swami, A. (1993). Mining Association Rules Between Sets of Items in Large Databases. *SIGMOD Rec.*, *22*(2), 207–216. https://doi.org/10.1145/170036.170072

Alba, D. (2015). Google Launches “YouTube Kids”, A New Family-Friendly App. Retrieved from Wired website: https://www.wired.com/2015/02/youtube-kids/

Albright, J. (2018). Untrue-Tube: Monetizing Misery and Disinformation. Retrieved from Medium website: https://medium.com/@d1gi/untrue-tube-monetizing-misery-and-disinformation-388c4786cc3d

Allcott, H., & Gentzkow, M. (2017). Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives*, *31*, 211–236. https://doi.org/10.1257/jep.31.2.211

Anderson, C. (2006). *The Long Tail: Why the Future of Business Is Selling Less of More*. Hyperion.

Bauer, I. (2015). When Copyright and Social Media Meet: Zooming in on Current Issues and Cases. *Undergraduate Law Journal*, *2*(12).

Ben-Shimon, D., Tsikinovsky, A., Rokach, L., Meisels, A., Shani, G., & Naamani, L. (2007). Recommender System from Personal Social Networks. *Advances in Soft Computing*, *43*, 47–55. https://doi.org/10.1007/978-3-540-72575-6_8

Bertot, J. C., Jaeger, P. T., & Grimes, J. M. (2010). Using ICTs to create a culture of transparency: E-government and social media as openness and anti-corruption tools for societies. *Government Information Quarterly*, *27*(3), 264–271. https://doi.org/https://doi.org/10.1016/j.giq.2010.03.001

Bobadilla, J., Ortega, F., Hernando, A., & Bernal, J. (2012). A collaborative filtering approach to mitigate the new user cold start problem. *Knowledge-Based Systems*, *26*, 225–238. https://doi.org/https://doi.org/10.1016/j.knosys.2011.07.021

Breese, J. S., Heckerman, D., & Kadie, C. (1998). Empirical Analysis of Predictive Algorithms for Collaborative Filtering. *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*, 43–52. Retrieved from http://dl.acm.org/citation.cfm?id=2074094.2074100

Burke, R. (2002). Hybrid Recommender Systems: Survey and Experiments. *User Modeling and User-Adapted Interaction*, *12*(4), 331–370. https://doi.org/10.1023/A:1021240730564

Burke, R. (2007). *The Adaptive Web* (Peter Brusilovsky, A. Kobsa, & W. Nejdl, Eds.). Retrieved from http://dl.acm.org/citation.cfm?id=1768197.1768211

Burke, R. D. (2002). Hybrid Recommender Systems: Survey and Experiments. *User Modeling and User-Adapted Interaction*, *12*, 331–370.

Burke, R., Felfernig, A., & H. Göker, M. (2011). Recommender Systems: An Overview. *Ai Magazine*, *32*, 13–18. https://doi.org/10.1609/aimag.v32i3.2361

Burke, R., Mobasher, B., Williams, C., & Bhaumik, R. (2006). Classification Features for Attack Detection in Collaborative Recommender Systems. *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 542–547. https://doi.org/10.1145/1150402.1150465

Cappon, L. (2019). Egypt’s new “Facebook law” targets freedom of speech. Retrieved from Coda Story website: https://codastory.com/authoritarian-tech/egypts-new-facebook-law-targets-freedom-of-speech/

Carbonell, J. G. (1989). Introduction: Paradigms for Machine Learning. *Artif. Intell.*, *40*(1–3), 1–9. https://doi.org/10.1016/0004-3702(89)90045-3

Chatzopoulou, G., Sheng, C., & Faloutsos, M. (2010). A First Step Towards Understanding Popularity in YouTube. In *2010 INFOCOM IEEE Conference on Computer Communications Workshops* (pp. 1–6). https://doi.org/10.1109/INFCOMW.2010.5466701

Dalul, S. (2019). How YouTube's algorithm facilitates child exploitation. Retrieved from Androidpit website: https://www.androidpit.com/youtube-algorithm-facilitates-sexual-exploitation-of-children

Darcy, O. (2019). How Twitter's algorithm is amplifying extreme political rhetoric. Retrieved from CNN Business website: https://edition.cnn.com/2019/03/22/tech/twitter-algorithm-political-rhetoric/index.html

Das, S. (2019). How does Twitter Make Money? | Twitter Business Model. Retrieved from Feedough website: https://www.feedough.com/how-does-twitter-make-money/

Deibert, R., & Rohozinski, R. (2010). Beyond Denial: Introducing Next-Generation Information Acess Controls. In *Access Controlled: The Shaping of Power, Rights and Rule in Cyberspace* (p. 6). MIT Press.

Deng, H. (2019). Recommender Systems in Practice. Retrieved May 13, 2019, from February 13 website: https://towardsdatascience.com/recommender-systems-in-practice-cef9033bb23a

Dutta, A. (2019). YouTube Business Model | How Does YouTube Make Money? Retrieved from Feedough website: https://www.feedough.com/youtube-business-model-how-does-youtube-make-money/

Egypt's Facebook users double: Ministerial report. (2012). Retrieved from Ahram Online website: http://english.ahram.org.eg/News/44676.aspx

Egypt to regulate popular social media users. (2018). Retrieved from BBC News website: https://www.bbc.com/news/world-middle-east-44858547

Ekstrand, M. D., Riedl, J. T., & Konstan, J. A. (2011). Collaborative Filtering Recommender Systems. *Found. Trends Hum.-Comput. Interact.*, *4*(2), 81–173. https://doi.org/10.1561/1100000009

El-Bermawy, M. M. (2016). Your Filter Bubble Is Destroying Democracy. Retrieved from Wired website: https://www.wired.com/2016/11/filter-bubble-destroying-democracy/

EL-Khalili, S. (2013). Social media as a government propaganda tool in post-revolutionary Egypt. *First Monday*, *18*(3). https://doi.org/10.5210/fm.v18i3.4620

Essinger, S., & L. Rosen, G. (2011). *An introduction to machine learning for students in secondary education*. 243–248. https://doi.org/10.1109/DSP-SPE.2011.5739219

Facebook-ads. (n.d.). Retrieved from Facebook Business website: https://www.facebook.com/business/ads

Facebook. (n.d.). *Facbook Analytics*. Retrieved from https://analytics.facebook.com/

Faris, R., & Villeneuve, N. (2008). Measuring global Internet filtering. *Access Denied: The Practice and Policy of Global Internet Filtering*, 5–28.

Fleder, D., & Hosanagar, K. (2008). *Blockbuster culture's next rise or fall: The effect of recommender systems on sales diversity*.

Gennari, J. H., Langley, P., & Fisher, D. (1989). Models of incremental concept formation. *Artificial Intelligence*, *40*(1), 11–61. https://doi.org/https://doi.org/10.1016/0004-3702(89)90046-5

Ghazanfar, M. A., Prügel-Bennett, A., & Szedmak, S. (2012). Kernel-Mapping Recommender system algorithms. *Information Sciences*, *208*, 81–104. https://doi.org/https://doi.org/10.1016/j.ins.2012.04.012

Gielen, M. (2016). 7 Expert Tips: How to Get YouTube to Recommend Your Video. Retrieved from Tubular Insights website: https://tubularinsights.com/7-expert-tips-youtube-suggested-videos/

Goldman, E. (2008). Search Engine Bias and the Demise of Search Engine Utopianism. In A. Spink & M. Zimmer (Eds.), *Web Search: Multidisciplinary Perspectives* (pp. 121–133). https://doi.org/10.1007/978-3-540-75829-7_8

Grace, K., Salvatier, J., Dafoe, A., Zhang, B., & Evans, O. (2017). When Will {AI} Exceed Human Performance? Evidence from {AI} Experts. *CoRR*, *abs/1705.0*. Retrieved from http://arxiv.org/abs/1705.08807

Grčar, M., Mladenič, D., Fortuna, B., & Grobelnik, M. (2006). Data Sparsity Issues in the Collaborative Filtering Framework. In O. Nasraoui, O. Zaïane, M. Spiliopoulou, B. Mobasher, B. Masand, & P. S. Yu (Eds.), *Advances in Web Mining and Web Usage Analysis* (pp. 58–76). Berlin, Heidelberg: Springer Berlin Heidelberg.

Groshek, J., & Tandoc, E. (2017). The affordance effect: Gatekeeping and (non)reciprocal journalism on Twitter. *Computers in Human Behavior*, *66*, 201–210. https://doi.org/10.1016/j.chb.2016.09.020

Hao, K. (2019). Facebook's ad-serving algorithm discriminates by gender and race. Retrieved from MIT Technology Review website: https://www.technologyreview.com/s/613274/facebook-algorithm-discriminates-ai-bias/

James, D., Benjamin, L., Junning, L., Palash, N., Taylor Van, V., Ullas, G., … Dasarathi, S. (2010). The YouTube video recommendation system. *Proceedings of the Fourth ACM Conference on Recommender Systems %@ 978-1-60558-906-0*, pp. 293–296. https://doi.org/10.1145/1864708.1864770

Jannach, D., Lerche, L., Kamehkhosh, I., & Jugovac, M. (2015). What Recommenders Recommend: An Analysis of Recommendation Biases and Possible Countermeasures. *User Modeling and User-Adapted Interaction*, *25*(5), 427–491. https://doi.org/10.1007/s11257-015-9165-3

Karypis, G. (2001). Evaluation of Item-Based Top-N Recommendation Algorithms. *Proceedings of the Tenth International Conference on Information and Knowledge Management*, 247–254. https://doi.org/10.1145/502585.502627

Kavanaugh, A. L., Fox, E. A., Sheetz, S. D., Yang, S., Li, L. T., Shoemaker, D. J., … Xie, L. (2012). Social media use by government: From the routine to the critical. *Government Information Quarterly*, *29*(4), 480–491. https://doi.org/https://doi.org/10.1016/j.giq.2012.06.002

King, G., Pan, J., & Roberts, M. E. (2017). How the Chinese Government Fabricates Social Media Posts for Strategic

Distraction, Not Engaged Argument. *American Political Science Review*, *111*(3), 484–501. https://doi.org/10.1017/S0003055417000144

Krulwich, B. (1997). Lifestyle finder: Intelligent user profiling using large-scale demographic data. *AI Magazine*, *18*, 37–45.

Kunaver, M., & Požrl, T. (2017). Diversity in recommender systems – A survey. *Knowledge-Based Systems*, *123*, 154–162. https://doi.org/https://doi.org/10.1016/j.knosys.2017.02.009

Lang, M. (2016). 2016 Presidential Election Circus: Is Social Media the Cause? Retrieved from Government Technology website: https://www.govtech.com/social/2016-Presidential-Election-Circus-Is-Social-Media-the-Cause.html

Laughland, O., Luscombe, R., & Yuhas, A. (2018). Florida school shooting: at least 17 people dead on "Horrific, horrific day." Retrieved from The Guardian website: https://www.theguardian.com/us-news/2018/feb/14/florida-shooting-school-latest-news-stoneman-douglas

Lika, B., Kolomvatsos, K., & Hadjiefthymiades, S. (2014). Facing the cold start problem in recommender systems. *Expert Systems with Applications*, *41*(4, Part 2), 2065–2073. https://doi.org/https://doi.org/10.1016/j.eswa.2013.09.005

M. Nichols, D. (1998). *Implicit Rating and Filtering*.

ma, Z., Sun, A., & Cong, G. (2013). On Predicting the Popularity of Newly Emerging Hashtags in Twitter. *Journal of the American Society for Information Science and Technology*, *64*. https://doi.org/10.1002/asi.22844

MacKinnon, R. (2011). China's "Networked Authoritarianism." *Journal of Democracy*, *22*, 32–46. https://doi.org/10.1353/jod.2011.0033

Mason, D. (2018). YouTube's recommended videos algorithm is promoting controversial content. Retrieved from KitGuru website: https://www.kitguru.net/gaming/damien-cox/youtubes-recommended-videos-algorithm-is-promoting-controversial-content/

Moeller, S. (2019). The 2019 Ultimate Guide to Facebook Engagement. Retrieved from BuzzSumo website: https://buzzsumo.com/blog/facebook-engagement-guide/

Moro, S., Rita, P., & Vala, B. (2016). Predicting social media performance metrics and evaluation of the impact on brand building: A data mining approach. *Journal of Business Research*, *69*(9), 3341–3351. https://doi.org/https://doi.org/10.1016/j.jbusres.2016.02.010

Most popular social networks worldwide as of April 2019, ranked by numbers of active users (in millions). (2019). Retrieved from Statista website: https://www.statista.com/statistics/272014/global-social-networks-ranked-by-number-of-users/

Narang, P. (2017). How To Get Your Hashtag Trending on Twitter In Less Than 48 Hours. Retrieved from Socialalert Blog website: https://blog.socialert.net/how-to-get-a-hashtag-trending/

Nguyen, T. T., Hui, P.-M., Harper, F. M., Terveen, L., & Konstan, J. A. (2014). Exploring the Filter Bubble: The Effect of Using Recommender Systems on Content Diversity. *Proceedings of the 23rd International Conference on World Wide Web*, 677–686. https://doi.org/10.1145/2566486.2568012

Ofcom. (2019). Children and parents: media use and attitudes report 2018. Retrieved from Ofcom website: https://www.ofcom.org.uk/research-and-data/media-literacy-research/childrens/children-and-parents-media-use-and-attitudes-report-2018

Pahwa, A. (2019). Faebook Business Model | How Does Facebook Make Money? Retrieved from Feedough website: https://www.feedough.com/facebook-business-model-makes-money/

Pariser, E. (2011). *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Group , The.

Park, Y.-J., & Tuzhilin, A. (2008). The Long Tail of Recommender Systems and How to Leverage It. *Proceedings of the 2008 ACM Conference on Recommender Systems*, 11–18. https://doi.org/10.1145/1454008.1454012

Pazzani, M.J., & Billsus, D. (2007). Content-Based Recommendation Systems. In P. Brusilovsky, A. Kobsa, & W. Nejdl (Eds.), *The Adaptive Web* (pp. 291–324).

Pazzani, Michael J. (1999). A Framework for Collaborative, Content-Based and Demographic Filtering. *Artificial Intelligence Review*, *13*(5), 393–408. https://doi.org/10.1023/A:1006544522159

Rashid, A., Albert, I., Cosley, D., Lam, S., McNee, S., Konstan, J., & Riedl, J. (2002). Getting to Know You: Learning New User Preferences in Recommender Systems. *International Conference on Intelligent User Interfaces, Proceedings IUI*. https://doi.org/10.1145/502716.502737

Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., & Riedl, J. (1994). GroupLens: An Open Architecture for Collaborative Filtering of Netnews. *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work*, 175–186. https://doi.org/10.1145/192844.192905

Resnick, P., & Varian, H. R. (1997). Recommender systems. *Commun. ACM %@ 0001-0782*, *40*(3), 56–58. https://doi.org/10.1145/245108.245121

Ricci, F., Rokach, L., & Shapira, B. (2011). Introduction to Recommender Systems Handbook. In F. Ricci, L. Rokach, B. Shapira, & P. B. Kantor (Eds.), *Recommender Systems Handbook* (pp. 1–35). https://doi.org/10.1007/978-0-387-85820-3_1 %U https://doi.org/10.1007/978-0-387-85820-3_1

Richert, W., & Coelho, L. P. (2013). Getting Started with Python Machine Learning. In *Building Machine Learning Systems with Python* (pp. 7–8).

Sands, M. (2018). Why Copyright Will Be The Biggest Issue For YouTube In 2019 (Updated). Retrieved from Forbes website: https://www.forbes.com/sites/masonsands/2018/12/30/why-copyright-will-be-the-biggest-issue-for-youtube-in-2019/#30d60eb81c12

Sarwar, B., Karypis, G., Konstan, J., & Riedl, J. (2001). Item-based Collaborative Filtering Recommendation Algorithms. *Proceedings of the 10th International Conference on World Wide Web*, 285–295. https://doi.org/10.1145/371920.372071

Schafer, B., J, B., Frankowski, D., Dan, , Herlocker, , Jon, , … Sen, S. (2007). *Collaborative Filtering Recommender*

Systems.

Schafer, J. Ben, Konstan, J. A., & Riedl, J. (2001). E-Commerce Recommendation Applications. *Data Mining and Knowledge Discovery*, *5*(1 %@ 1573-756X), 115–153. https://doi.org/10.1023/A:1009804230409 %U https://doi.org/10.1023/A:1009804230409

Schein, A. I., Popescul, A., Ungar, L. H., & Pennock, D. M. (2002). Methods and Metrics for Cold-start Recommendations. *Proceedings of the 25th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 253–260. https://doi.org/10.1145/564376.564421

Scott, S. (2019). YouTube's algorithm is spreading a series of unfortunate far-right events. Retrieved from The Drum website: https://www.thedrum.com/opinion/2019/05/14/youtube-s-algorithm-spreading-series-unfortunate-far-right-events

Shalev-Schwartz, S., & Ben-David, S. (2014). When Do We Need Machine Learning? In *Understanding Machine Learning: From Theory to Algorithms* (pp. 21–22). Cambridge University Press.

Shavlik, J., & Dietterich, T. (1990). General aspects of machine learning. In *Readings in Machine Learning* (pp. 1–10). Morgan-Kaufmann.

Simon, A., Singh Deo, M., Selvam, V., & Babu, R. (2016). An Overview of Machine Learning and its Applications. *International Journal of Electrical Sciences & Engineering*, *Volume*, 22–24.

Sinha, R., & Swearingen, K. (2001). Comparing Recommendations Made by Online Systems and Friends. In *In Proceedings of the DELOS-NSF Workshop on Personalization and Recommender Systems in Digital Libraries*.

Szabo, G., & Huberman, B. A. (2010). Predicting the Popularity of Online Content. *Commun. ACM*, *53*(8), 80–88. https://doi.org/10.1145/1787234.1787254

Tatar, A., de Amorim, M. D., Fdida, S., & Antoniadis, P. (2014). A survey on predicting the popularity of web content. *Journal of Internet Services and Applications*, *5*(1), 8. https://doi.org/10.1186/s13174-014-0008-y

The Global State of Digital in 2019 Report. (2019). Retrieved from Hootsuite website: https://hootsuite.com/pages/digital-in-2019

Thielman, S. (2016). Facebook news selection is in hands of editors not algorithms, documents show. Retrieved from The Guardian website: https://www.theguardian.com/technology/2016/may/12/facebook-trending-news-leaked-documents-editor-guidelines

Toulas, B. (2019). Egypt Intensifies Internet Censorship by Blocking Social Media Accounts. Retrieved from TechNadu website: https://www.technadu.com/egypt-internet-censorship-blocking-social-media-accounts/62242/

Tufekci, Z. (2018a). It's the (Democracy-poisoning) Golden Age of Free Speech. Retrieved from Wired website: https://www.wired.com/story/free-speech-issue-tech-turmoil-new-censorship/?CNDID=50121752

Tufekci, Z. (2018b). YouTube, the Great Radicalizer. Retrieved from The New York Times website: https://www.nytimes.com/2018/03/10/opinion/sunday/youtube-politics-radical.html

Twitter. (n.d.). *Analytics*. Retrieved from https://analytics.twitter.com/about

Vargas, S., & Castells, P. (2011). Rank and Relevance in Novelty and Diversity Metrics for Recommender Systems. *Proceedings of the Fifth ACM Conference on Recommender Systems*, 109–116. https://doi.org/10.1145/2043932.2043955

Vincent, J. (2016). Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day. Retrieved from March 24 website: https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist

Wijnhoven, F., & Brinkhuis, M. (2015). Internet information triangulation: Design theory and prototype evaluation. *Journal of the Association for Information Science and Technology*, *66*(4), 684–701. https://doi.org/10.1002/asi.23203

Wilson, G. (2019). If Facebook keeps censoring Christians, perhaps it's time to say goodbye. Retrieved from Christian Today website: https://www.christiantoday.com/article/if-facebook-keeps-censoring-christians-perhaps-time-to-say-goodbye/132592.htm

Wodinsky, S. (2019). YouTube's copyright strikes have become a tool for extortion. Retrieved from The Verge website: https://www.theverge.com/2019/2/11/18220032/youtube-copystrike-blackmail-three-strikes-copyright-violation

YouTube. (n.d.). Analytics and Reporting APIs. Retrieved from https://developers.google.com/youtube/analytics/

Zednik, C. (2019). *Solving the Black Box Problem: A General-Purpose Recipe for Explainable Artificial Intelligence*.

Zhou, X., Xu, Y., Li, Y., Josang, A., & Cox, C. (2012). The state-of-the-art in personalized recommender systems for social networking. *Artificial Intelligence Review*, *37*(2), 119–132. https://doi.org/10.1007/s10462-011-9222-1