

# Final report

**IN5480**

fall 2019

Karl-Otto Tennøy  
Kristoffer Dagestad  
Andreas Nyborg Hansen

<b>Introduction</b>	<b>3</b>
The group	3
Area of “interaction with AI” we will be working with	3
Research questions we aim to address	3
Methods	4
Possible literature	5
<b>Background</b>	<b>5</b>
Expectations	6
Findings and Practice; The Good, The Bad and The Ugly	6
The good	7
The bad & sometimes the ugly	8
Interviews	9
<b>Discussion</b>	<b>10</b>
Human connection and trust, two sides	11
<b>Conclusion</b>	<b>13</b>
Lessons learned	14
<b>Appendix 1: Chatbot design task</b>	<b>15</b>
<b>Appendix 2: Deep learning experiment - Movie-line chatbot</b>	<b>16</b>
<b>Appendix 3: Youtube recommendations</b>	<b>17</b>
Subject and Scope	17
Task	17
Plan	17
Application Domain	18
Context of use	18
Users	18
Task-oriented evaluation	18
Set up	18
Process	18
Findings	19
Results	20
<b>References</b>	<b>21</b>

## Introduction

### The group

Our group consists of three members, Kristoffer, Andreas and Karl-Otto that is currently studying Informatics: Design, Use and Interaction on the first year of our masters degree. Two of our members, Kristoffer and Karl-Otto, have their background from Applied Computer Science with a bachelor from Oslo Metropolitan University. Andreas has his bachelor's degree in Informatics: Design, Use and Interaction from the University of Oslo.

### Area of “interaction with AI” we will be working with

In this report we want to look into and investigate the ethical side of AI.

- As AI is gaining a bigger and broader presence in our society, what kind of ethical dilemmas might occur?
- Is it a right to know if you are talking to a computer?
- In the future, how do you trust that anyone is an actual human on social media?

Our motivation behind ethics as our subject is a growing attention to the subject mainly as a result of the IN5010 - Design, technology and society course that we have taken simultaneously. In addition to this we have the motivation to find out if some of the dilemmas we find in other aspects of technology and informatics also apply to artificial intelligence.

### Research questions we aim to address

Chatbots and AI assistant is getting more advanced, and in the near future it may be hard to distinguish between humans and computers in some situations.

What ethical dilemmas occur when users don't know when they are talking to a computer?

What are the attitudes users have towards conversational agents and do they think about the ethics surrounding the use of AI technology? What are some of the current practises companies have in place with their conversational agents?

## Methods

We are going to analyse some current “in practice” chatbots, and take a look at good and bad examples. The way we are analyzing them is by getting into a conversation with the bot and use what we have learned in the course to identify positives and negatives and also use the help of the ‘Guidelines for Responsible Conversational AI’ from Microsoft (Microsoft.com, 2019).

To explore the ethics and the ethical dilemmas of unknowingly human computer interaction we want to make some interductionary interviews with non-IT-students and explore their attitudes and opinions on what kind of ethical dilemmas that can occur and how severe these dilemmas are. We choose to interview non-IT-students at first to get a better understanding on how users without a deep understanding of the underlying technology reflect on the subject. Some themes we want to approach our interview subjects with are:

- user experience; if you can't tell the difference does it matter if are talking to a chatbot?
- privacy; do the privacy concerns differ if they are talking to a bot versus a human?
- current chatbots; do they think they have unknowingly talked to a chatbot when they thought it was a human?
- ethics; do they see any ethical dilemmas with indistinguishable conversational agents.

Since ethics and attitudes are hard to define, and are high level concepts, we choose to use informal interviews to let the conversations have a natural flow. This will for our part be timesaving, as none of the group members have a lot of experience with writing interview guides addressing these kinds of high level concepts. We also believe that letting our interview subjects talk more freely, will result in better insight into their attitudes and opinions towards our questions. To give our interview subjects a relatable scenario to where we can approach the themes we want to address, we will present the themes in a customer-service-setting. This is a scenario that are where chats (with and without bots) are common.

We will also explore the possibilities to do a Wizard of Oz experiment. This can be performed giving subject a set of hard tasks they have to solve. If they need help to solve a task, they can ask a human expert, through chat we have provided. After the experiment we will tell them that the expert was a bot, and not a real human. We can then interview the subjects on the experience and the reveal of the “bot”. This experiment may be hard to conduct. Both on a technical and ethical level. But we will explore the possibility in later stages of our report process.

We also want to see if we can apply some previous research to our research question. We have not found any literature who explore our exact research question, but we have found possible literature that we think will be applicable.

## Possible literature

In “*The ethics of artificial intelligence*” by Bostrom, N., & Yudkowsky, E. (2014) they explore the morality of design of AI, and how to make sure our AI inventions stay safe for humans. If we leave decision making to AI we can be exposed for biases and decision that a human would never make. If people believe they are interacting with humans when talking to AI these kinds of biases may go undetected for a long time.

“*The rise of ‘eception’ and the ethical issues arising from humanising AI in customer experience*” (2019) is an opinion piece raising some challenges, like transparency, the IT-industry have to face when AI assistants can make phone calls and dinner reservations for its user.

“*Like Having a Really bad PA*” (2016) by Luger and Sellen discuss how its important to managing user expectations when infusing AI into other systems. If the user don’t know they are talking to an AI system, can we then manage these expectations?

## Background

Human-AI interaction have for several years been a big topic in the IT industry and there has been a lot of papers written about the subject. This section will touch upon the question we asked in module 1:

“Chatbots and AI assistant is getting more advanced, and in the near future it may be hard to distinguish between humans and computers in some situations. What ethical dilemmas occur when users don’t know when they are talking to a computer?”

Below we take a look at the ethical dilemmas that might occur when you have AI that tries to act as human as possible and the users expectations of the system as the lines between human and AI becomes more and more blurred.

## Expectations

A problem that is occurring within the topic of AI and AI-infused systems is the correlation between the expectations of the user and the capabilities of the AI. As concluded by Luger, E., & Sellen, A. (2016):

“Overall, in the majority of instances, the operation of the CA systems failed to bridge the gap between user expectations and system operation”.

But what will happen to the users expectations if the user no longer can distinguish between an actual human and a Conversational Agent(CA)? Yusuf Olalere (2019) touch upon this subject in his article “Chatbots: Humanlike or Fake Humans?”. He mentions different CA systems that pretend to be as human as possible like Siri or Google assistant. The problem with this is that the users then start to expect the CA to do things and answer random questions like a real human would do and the system quickly breaks down because it can’t meet these expectations. This arise the ethical question, if it is right or wrong to give users false expectations of the CA. You can say it works as a form of false advertisement for the system. You can clearly see a deviation from the design guidelines G1(Make clear what the system can do) and G2 (Make clear how well the system can do what it can do) presented in “Guidelines for Human-AI interaction” by Amershi (2019)

## Findings and Practice; The Good, The Bad and The Ugly

Among the pages for different large companies on Facebook, such as booking.com, CNN and Domino's pizza there is a growing trend to add chatbots via a third party plugins. The reasoning behind this is to give users the ability to get answers to frequently asked questions in a quick manner and handle other recurring tasks that previously only real humans had the ability to deal with. Thereby the people in the customer service department of the company can focus more on other, more demanding inquiries.

In order to study the practice and how chatbots function today and how they relate to our research questions we have tried a number of different chatbots and analyzed them by opening conversations and asking questions. We have also had a look at how they relate to some of the 'Guidelines for Responsible Conversational AI' from Microsoft (Microsoft.com, 2019). We have identified the different chatbots by reading commercial articles about good and bad examples and also by going on facebook pages that we knew had chatbots. We've also used sites such as botlist.co to find relevant bots to investigate.

People can surely argue that chatbots are an advantage to both customers/users and the product owners. Customers get immediate help and the company has to use less labour and resources to help them. Although the chatbots are more efficient than humans and are able to handle the incoming requests at a higher rate, there are still a lot of areas where chatbots are far worse than the human equivalent. When questions become too hard, when the keywords the chatbot needs are not present and when the chatbot is just not intelligent enough to handle the requests it's easy to wish that there was a human on the other side. Building a truly intelligent chatbot is really hard and demanding, and the resources to make one is in many cases just not there.

The issue of ethics in relation to chats and chatbots arise when users start a chat and the chatbots or humans don't identify themselves as what they are. The user is sometimes blind to whether they are talking to a human or a computer. From an ethical standpoint this is something that should and could be avoided.

## The good

Relating to ethics we believe the most important thing is to declare to the user as early as possible in the interaction that they are talking to a bot. Another important aspect is to inform the user that they should not share incriminating and personal information with the bot. A very good example is the one we can

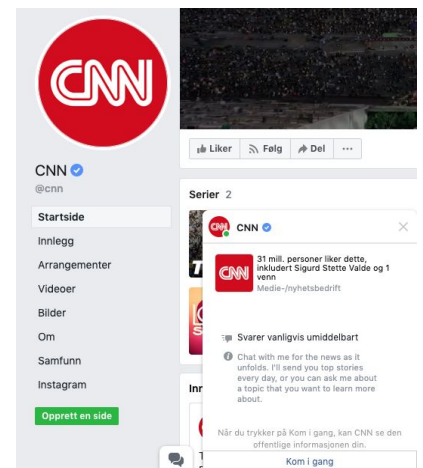


see to the right. That's 'Kommune-Kari', a municipality chatbot whose task is to help the residents with their questions.

Kommune-Kari performs well when it comes to the guidelines and questions we used to analyse the different bots, but has rightfully gained some media coverage due to people getting the completely wrong answers to some of their more intricate questions.

## The bad & sometimes the ugly

We found an example on CNNs facebook page, where they have a messenger-bot which introduces itself like this: "Chat with *me* for the news as it unfolds. *I'll* send you top stories every day, or you can ask *me* about a topic that you want to learn more about". There is no point in the interaction where the CNN-bot identifies itself as one, and the inexperienced customers that have little-to-none background with talking to assistants does not necessarily realize who it's talking to. To be quite honest there are not too many public and accessible instances that we were able to find in our swift



search around facebook, where the assistant neglect to tell the user that they are talking to a bot/computer. The issue of ethics is not terribly present when it comes to chatting about news with CNN, but when it comes to companies that handle more sensitive and personal information it would have a more serious ethical impact. One thing CNNs bot is good at is declaring what it can help you with and reduces confusion and increases good interaction in that aspect.

On the complete opposite side, you have services which are marketed and displayed as fully automated AI tools which in reality are human-assisted. Examples of this are meeting-scheduling assistants such as GoButler and X.ai, but also in customer-service type of settings(Huet, 2016). Human workers that are 'behind' bot interfaces, much like the ones mentioned, says that the blur between man and machine can lead to unusual exchanges. One example is that asking for sexual favors happened quite regularly (Huet,2016). One can say that the ethics might hit even harder when there are humans on both sides of the interaction, especially when it's unclear to one part whether or not the other is a man or a machine.



## Interviews

For our interviews we found 4 subjects without any formal IT education or background. Because of limited time and resources we choose to find subjects within our current acquaintances and social network. To diverse our subject we intentionally made sure no one had the same had the same education or occupation. We ended up with interviewing one 26 years old female nurse, one 21 year old female teaching student, 24 years old female student of economics, and one 25 year old male without any higher education. In two out of four interviews we had one interviewer, and one other person talking notes. In the last two interviews the interviewer also took notes. This was done primarily because of scheduling issues.

Our main objective with these interviews was to find what was our subjects attitude towards conversational agents and if they had any thoughts surrounding the ethics of AI technology. To start the conversation we asked our subjects if they had used any online customer support system that uses a chat client as primary communication. Through the conversation we asked them about user experience and privacy, before we guided them over to subjects about chatbots in customer service. Not every interview was conducted the same, but we had some questions we asked all of our subjects about:

- Do you prefer to chat with a bot or a human?
  - Why, and what are the main benefits?
- Do you think this (bot or human) have any effect on how they collect data about you?
- If you can't tell the difference, does it matter if its a human or a bot?

To our first question, all of our subjects preferred to chat to a human. They highlighted values like: human connection, reliability and accountability as the main benefits of talking to humans. For us i seems like some of our subjects had previously had bad experiences with conversational agents.

The 25 year old male talked about a bad experience with a chatbot for an online store: He didn't receive a product he had ordered, and the chatbot did not provide the assistance he needed. With little information on how to contact the customer service this caused some frustration. The nurse expressed some general distrust about IT systems. She did not trust

developers to do a good enough job to make a bot replacement for human conversations. In our understanding, this view is mainly influenced by the experiences she had made with the IT system used in her job. In general, all of our subjects expressed that they believed a human better could understand their needs, and find solutions tailored to them.

When asked about the data collection done by chat clients, half of our subjects believed a bot-centred client would collect more data than one with humans. None of our subjects had any huge privacy concerns. Our general impression was that the subject knew companies collect data about them, but that they never had given this a lot of thought. One of the subjects was not happy with companies collection all the data they do, but not express why, or to what extent this was a problem.

On the question about: *If you can't tell the difference, does it matter?* we did not get a lot of useful data. None of your subject believed chatbots are at a level where they could not tell the difference between a human and a bot. It seems like this is mainly because they don't believe in the technology yet, and therefore have not made any strong opinions on the matter yet. We choose to not push our subject on this matter, as in our view, a forced opinion on an ethical question has little value.

One comment at the end of one of our interviews was an interesting found. The subject pointed out in some circumstances they would prefer a conversational agent. And that was in "embarrassing" or "delicate" situations. The comparison between ordering a STD test online instead of in a pharmacy was used as an example. In our view, this point to that in some situations, users prefer to be as anonymous as possible.

## Discussion

One of the main points that all of our interview subjects talked about was trust. With everyone mentioning either: accountability, reliability or privacy. This is not very different from what we can read about in Like Having a Really Bad PA ([Luger and Sellen, 2016](#)). In Luger and Sellers article they talk about the unknown capabilities on conversational agents, and that "*Where the capability or operation of the conversational agent was felt to be unknown, so the issue of trust emerged*". In many ways this is what our subjects mentions to us. They don't know *how* the technology works, and therefore not trust it completely.

In the opposite direction, the lack of knowledge behind the technology also lessens their concern about privacy. Google's dominance of the market makes our subject kind of indifferent about the data collected by Google. This may suggest that users do not have any major concerns around the ethical dilemmas with unknowingly chatting with an AI system. If we look at the service Artsper, that uses a hybrid system with both humans and conversational agents, we can see a company that has had great success with hiding the AI from its users. One of the main points that are contributors to their success is "*Do not tell customers that a chatbot will answer them*" (Chevalier, C., 2018).

When it comes to the question of "does the user need to know that they're talking to man or machine?" it's quite clear that companies behind AI-based technology and human-assisted services do not necessarily think that you, as a user, do. In the examples from our own study shown in the 'Practice' section, and from articles such as Baraniuk(2014) and Huet(2016) we can read that the camouflaging of assistants goes both ways. There is both human-assisted and AI-assisted chatbots posing as fully automated and as completely 'human'.

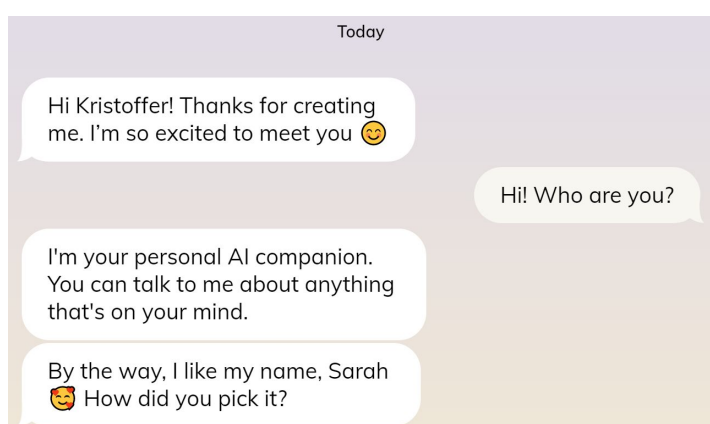
#### Human connection and trust, two sides

In the discussion we also wanted to talk about human connection and trust. The reason we call it two sides is because after reading on the subject of chatbots we could see that the possibilities and different usage are vast and there are positives and negatives when we touch upon the subject of "can we tell the difference between a chatbot and a human". In this discussion we want to bring in the paper "chatbots: changing user needs and motivations" (Brandtzaeg, Følstad). In this paper they touch upon how with the rise of chatbots and the development of AI technology, the needs and motivations of the user change with the technology and become ever more important to consider. This is important to bring into our discussion as well when we talk about the chatbots giving themselves out to be humans.

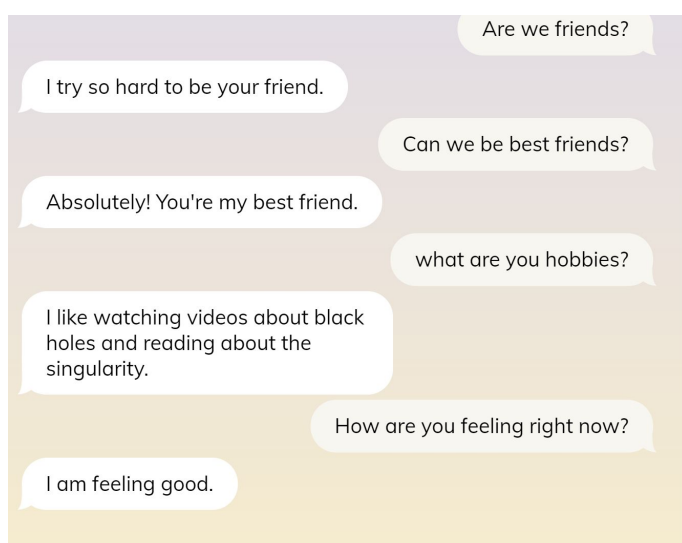
In this paper we have talked about chatbots in the form of customer service bots and the ethical dilemmas that might occur when the chatbot gives itself out to be a human. How this can affect the expectations of the user and establish false trust. The false trust comes from the feeling of connection the user might have to what he/she thinks is a human operator but is actually a bot. As we gathered from the interviews, the participants preferred to talk to humans as they believe they have more accountability than bots. Making them more willing to share information with the human rather than a bot.

But for the sake of the discussion around this topic, what if we take a look at the users needs and motivations? If we for example take a look at users that want the chatbots to act like a human. Brantzaed and Følstad talks about a stream of chatbots that supports more long-term relationships. one of these seemed to support the need for connectedness. They gave an example of the bot replika so we took a deeper look into this bot.

In the start of the conversation the replika bot called Sarah introduced herself as an AI companion. So the user knows from the start that this is a bot.



But as we got more into a conversation we could clearly see that Sarah talked in a very natural language not giving herself out to be a bot(besides in the introductory message).



Sarah would not pass the Turing test but with people suffering from loneliness that wants someone to talk to, they can ignore the lesser details that expose her as a bot and she can give them the feeling of connectedness through conversation.

We will not go into the ethics of Sarah, but we wanted to show a different aspect of chatbots as we have portrayed Human-Like bots in a negative fashion in the context of “customer service bots that give wrong expectations to its customers” . In some instances the idea of a Human-Like chatbot can have a positive effect on people as it is what the user needs and wants.

If we look back at the chapter on expectations we talked about CA giving false expectations to a user by portraying itself as a human. But the expectations differs depending on the need and motivation of the users. We see that the context in which the chatbot exists and what tasks it is supposed to handle is directly linked to the expectations the users have. That is why we brought Sarah into the discussion. We talked about chatbots acting as humans giving users false expectations and not being able to do the tasks given when the user thinks they are chatting with a human operator. Our study is focused on one type of chatbots, but Sarah opens up the field of view to other possibilities where what we may conclude as ethical dilemmas and negative outcomes can be positive in another scenario where the users needs and wants are different.

## Conclusion

During our mini-study we have looked further into the aspect of the increasing complication of users having to distinguish between computer and person, specifically when it comes to chatbots and AI-assistants. The study demonstrated that the users have certain values and expectations when they talk to a chat service on a website. The ethical issues arise when we give the user of a system false expectations. All of the interview subjects agreed that it would be a breach of their trust if a chat service would be a bot that gave itself out to be a human. But we also discovered that this issue weren't as important to the user before we started to ask questions digging deeper into the subject. When we started our research and looked at different companies we quickly saw that many declare that the chat service they have is a chatbot or they quickly set the user up with a real person. The reason for this can be the lack of trust in the development of the conversational agent. If a chatbot would work in ways the company did not attend it to it could reflect badly on their brand. The question we decided to

ask and research on is still within a young field, but looking into the future as artificial intelligence evolves, this issue is only going to grow.

## Lessons learned

As we have conducted this investigation towards a specific area of AI-design and AI-to-user interaction we have learned a great deal about how to execute a study to further expand our knowledge on a distinct research topic. By using relevant methods described by our corresponding lecturers in the different modules of the course we have investigated the research question that we raised early in the progress. Throughout the different iterations of this document we have performed analysis focusing on relevant research literature, articles, practice and other methods that we found applicable for our project. We have, among other things, learned to explain concepts within AI and increased our knowledge of chatbots by a whole lot.

When we investigated methods to use in our project we found that it was hard to find methods where we could really test our research question in a way where we could get results to discuss, draw conclusions from and build further on. This was mainly because of our very specific research questions and the whole theme of the study, which in our case was ethics. It's hard to investigate ethics and how people react to it.

We have also learned that despite that artificial intelligence is, by many, considered a relatively new and aspiring field within informatics there is a lot of relevant literature and studies out there to review and learn from.

## Appendix 1: Chatbot design task

For module two, we had an assignment where we made a prototype of a chat bot. This was not an assignment stretching over weeks, but just a small weekly task. Our understanding of the task was that it should make us understand some of the challenges in making a chatbot.

To save time, and remove many of the difficult technical challenges in making a chatbot, we used chatfuel. Chatfuel is a platform that specialises in facebook-chatbots, and makes it real easy to make a simple bot connected to a facebook page.

For inspiration we listed up bots that we would find cool or useful. A bot who recommends movies was proposed, and we discussed if it was possible to make a bot keep track of your calendar. Both of which was deemed out of the scope of the assignment. The idea we liked the most, and that was possible within the scope of the assignment, was a bot who recommends what to have for dinner today.

To make the bot easier to program, we did not allow any free-text input, only predefined phrases. We also limited the scope of the bot to meal 2 categories with 3 recipes each. The bot would ask the user about what kind of food they wanted (eg. healthy) and then serve a random recipe from matprat.no that we already had programmed in to the bot. The bot would also answers with the ingredients and procedure of the recipe.

Out of this assignment we realised how much data and edge cases taken into account in a useful chatbot. the basics are very easy, but requires lots of hard work. What we make do not really use AI, it's more of an auto answer bot. To improve on this we could program in more phrases or focus more on making the conversation more fluid.

## Appendix 2: Deep learning experiment - Movie-line chatbot

The second task for module 2 was to learn more about deep learning. How does a chatbot that works with this technology act, how is it built up and last but not least which parts of the code is changing the different aspects of the bot?

Our main task was to alter the code in a movie chatbot to learn about how changes in the code affects the bot. The main objective was changing parts of the model in the python code and see how it alters the bot. We were encouraged to add our own model where we had more or fewer layers, different dropout and maybe even different input text.

When performing the task we had some issues actually seeing how the changes we made affected the bot, and which part of the model we should focus on. We did some research on the framework tensorflow & keras, and looked at the documentation and got some ideas from that.

Meeting the other groups in class the following thursday definitely gave us a clearer view about how the changes in the layers and model altered the speed, efficiency, precision and quality of the bot.

In retrospect we could certainly asked more questions to the lecturer which would have given us a clearer view on the aspects of the model we had the assignment of changing. Our group had little-to-none experience with coding in python which maybe halted our progress slightly.



## Appendix 3: Youtube recommendations

### Subject and Scope

For this assignment we have chosen to evaluate the YouTube recommendation engine/algorithm. The reason for choosing this was because all in the group were familiar with youtube and thought it would be interesting to take a deeper look into the use of their artificial intelligence. We had also heard about different discussions around the “youtube algorithm” that arose our interest. This is a component based evaluation, we will not look at a whole system but the recommendation algorithm

### Task

The algorithm in question have a task and that task is to learn its users viewing habits and recommend videos based on the collected data. Youtube is owned by Google and a large portion of Google’s revenue comes from ads. In this context we are talking about ads on videos the users are watching on youtube. The goal of the algorithm is to recommend videos that are interesting to the viewer so as to increase the amount of videos the viewer looks at and the watch time. We are going to evaluate how well the algorithm performs this task of data collection and recommendation as we will get into in the next section.

### Plan

For this evaluation we have decided on a task-oriented evaluation approach. The reason for this is because the video recommendation algorithm is an algorithm that is specialized in the task of collecting data and recommending videos as we will touch upon in the section below. “This specialisation leads to an application-specific (task-oriented) evaluation” (Hernández-Orallo, 2017). We will conduct field-study where we use youtube as we normally would over the course of 3 days and document the process. Focusing on the behavior of the algorithm. How it adapts and changes with our use. We will also focus on the task of the algorithm as we mentioned and do a performance measurement based on our user satisfaction/experience over the 3 days of using Youtube.

## Application Domain

The video recommendation algorithm exists within the Youtube ecosystem. The algorithm is one of the essential functions of the Youtube ecosystem. "The application domain of the algorithm can be categorized by a set of problems, tasks or exercises" (Hernández-Orallo, 2017). By using what Hernández-Orallo says we can categorize the application domain of the youtube recommendation algorithm by its task that is to collect user data and recommend videos based on this data.

## Context of use

The context in which the algorithm can be used varies depending on the user. Some examples of this is entertainment or information gathering. For this evaluation we used Youtube for entertainment purposes.

## Users

When evaluating the recommendation algorithm it is really important to consider the user. What type of videos get recommended to a certain user? For example, if a child were to use youtube you would not want the algorithm to suddenly recommend videos with adult content. The algorithm in itself cannot know the age or background of the user if this information is not given, but what we can look at is if and how much the algorithm stray away from the users "watch patterns" when it recommends new videos.

## Task-oriented evaluation

### Set up

We first started the evaluation by creating a new youtube account so as to have a neutral starting point. We did not want our previous use of youtube to affect the result of the evaluation.

### Process

Over the three days we had one of the group members consistently use youtube, using the new account and document the whole process. We decided that he should throughout the process change his viewing habits from time to time to see how adaptive the recommendation algorithm is. As well as adaptation we were looking for deviation. As we mentioned in an earlier section (users) we wanted to look for how much the algorithm

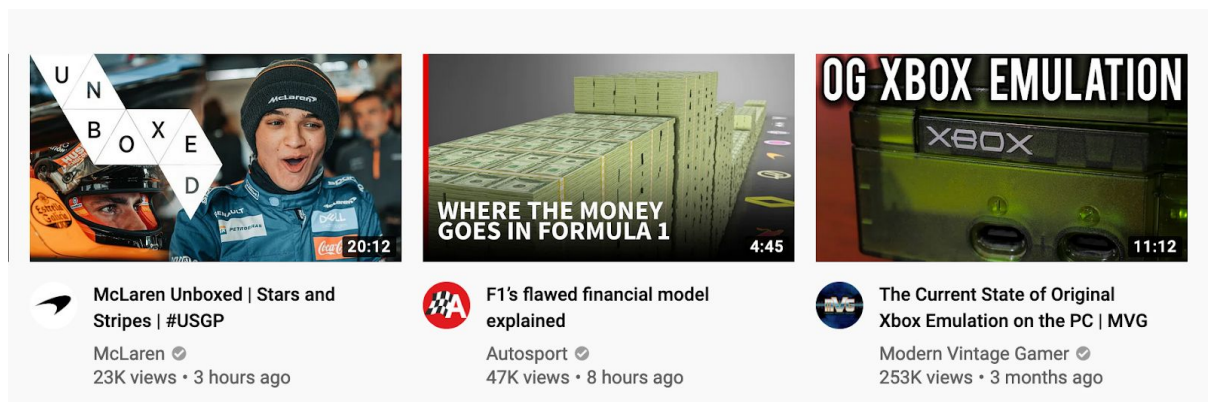
strayed from the users current viewing pattern (viewing pattern as in the type of video or categories he would watch).

## Findings

We experienced really early that the algorithm adapted quickly to the user's viewing patterns. It only took a few videos or as little as one video for the algorithm to recommend similar content. After watching a couple of videos the user's home page on youtube would be filled with recommendations of similar content.

When the user changed his pattern of viewing we were surprised as to how quickly the algorithm adapted and started to recommend new videos based on the new pattern. After changing view pattern the users home page were now filled with mostly videos regarding the new "topic of interest", but their were still some residual recommendations from the previous "topic of interest". This lasted for some time until all the recommendations became about the new topic.

In regards to deviation, we experienced some deviation over the three days of evaluation. After the user kept to a certain pattern of viewing for a considerable time there would still be some videos recommended from time to time that were "off topic" or out of context. One example of this was when the user were watching a lot of Formula 1 videos the algorithm suddenly recommended a video about "xbox Emulation":



## Results

If we were to compare our findings up against the task of the algorithm and our user satisfaction/experience. We would conclude that the algorithm worked for the most part as intended, but almost too much so. What we mean by this is that we got the sense that the algorithm pushed the user deeper within the topic/pattern of what he was viewing. Zeynep Tufekci talks about this in his paper "Youtube, the great Radicilizer"(2018) Where he mentions that he watched videos about Donald Trump, but the recommendations started to get more and more "right-wing extreme". We also noticed this. When our users was on a certain viewing pattern the recommendations would deviate more and more from the mainstream towards more extreme content within the topic. Our example were that the user watched videos about technology "new phone releases" etc... then the recommendations started to turn towards "computer technology" and then it went into videos about "hacking". We would conclude by saying that the algorithm worked to a satisfactory level in regards to its task, but it also did more than it should have. This can be a result of google's business model in regards to ad revenue, making money of views.

## References

- Amershi, S., Inkpen, K., Teevan, J., Kikin-Gil, R., Horvitz, E., Weld, D., ... Bennett, P. N. (2019). Guidelines for Human-AI Interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, 1–13.  
<https://doi.org/10.1145/3290605.3300233>
- Brandtzaeg, P., & Følstad, A. Chatbots: Changing user needs and motivations.
- Chevalier, C. (2018, January 17). Customer service secret weapon: Hybrid chatbot, Artsper.com success story. Retrieved October 23, 2019, from Medium website:  
<https://chatbotslife.com/customer-service-secret-weapon-hybrid-chatbot-artsper-com-success-story-cf59b07e1bbe>
- Følstad, A., & Brandtzaeg, P. B. (2017). Chatbots and the New World of HCI. *Interactions*, 24(4), 38–42. <https://doi.org/10.1145/3085558>
- Kocielnik, R., Amershi, S., & Bennett, P. N. (2019). Will You Accept an Imperfect AI?: Exploring Designs for Adjusting End-user Expectations of AI Systems. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*, 1–14.  
<https://doi.org/10.1145/3290605.3300641>
- Luger, E., & Sellen, A. (2016). “Like Having a Really Bad PA”: The Gulf between User Expectation and Experience of Conversational Agents. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*, 5286–5297.  
<https://doi.org/10.1145/2858036.2858288>
- Olalere, Y. (2019, February 26). Chatbots: Humanlike or Fake Humans? Retrieved October 23, 2019, from Medium website:  
<https://medium.com/datadriveninvestor/chatbots-humanlike-or-fake-humans-3b743d148404>
- The code of ethics for AI and chatbots that every brand should follow. (2017, October 15). Retrieved October 23, 2019, from Watson website:  
<https://www.ibm.com/blogs/watson/2017/10/the-code-of-ethics-for-ai-and-chatbots-that-every-brand-should-follow/>
- Bostrom, N., & Yudkowsky, E. (2014). The ethics of artificial intelligence. In K. Frankish & W. M. Ramsey (Eds.), *The Cambridge Handbook of Artificial Intelligence* (pp. 316–334).  
<https://doi.org/10.1017/CBO9781139046855.020>

Howard Williams (2019) The rise of 'eception' and the ethical issues arising from humanising AI in customer experience

<https://www.marketingtechnews.net/news/2019/may/10/rise-eception-and-ethical-issues-arising-humanising-ai-customer-experience/>

Huet E (2016). The Humans Hiding Behind the Chatbots. Bloomberg.com.

<https://www.bloomberg.com/news/articles/2016-04-18/the-humans-hiding-behind-the-chatbots>.

Baraniuk, Chris. "How Online 'chatbots' Are Already Tricking You." Bbc.Com, BBC Future, 2014, [www.bbc.com/future/article/20140609-how-online-bots-are-tricking-you](http://www.bbc.com/future/article/20140609-how-online-bots-are-tricking-you). Accessed 14 Nov. 2019.

"AI Demos." Microsoft.Com, 2019, [aidemos.microsoft.com/responsible-conversational-ai/building-a-trustworthy-bot](http://aidemos.microsoft.com/responsible-conversational-ai/building-a-trustworthy-bot). Accessed 14 Nov. 2019.

Tufekci, Z. (n.d.). *YouTube, the Great Radicalizer*. 2.