

When People and Algorithms Meet: User-reported Problems in Intelligent Everyday Applications

Malin Eiband, Sarah Theres Völkel, Daniel Buschek, Sophia Cook, Heinrich Hussmann

LMU Munich, Munich, Germany

malin.eiband,sarah.voelkel,daniel.buschek,hussmann@ifi.lmu.de

sophia.cook@campus.lmu.de

ABSTRACT

The complex nature of intelligent systems motivates work on supporting users during interaction, for example through explanations. However, there is yet little empirical evidence on specific problems users face in such systems in everyday use. This paper investigates such problems as reported by users: We analysed 35,448 reviews of three apps on the Google Play Store (Facebook, Netflix and Google Maps) with sentiment analysis and topic modelling to reveal problems during interaction that can be attributed to the apps' algorithmic decision-making. We enriched this data with users' coping and support strategies through a follow-up online survey (N=286). In particular, we found problems and strategies related to content, algorithm, user choice, and feedback. We discuss corresponding implications for designing user support, highlighting the importance of user control and explanations of output, not processes. Our work thus contributes empirical evidence to facilitate understanding of users' everyday problems with intelligent systems.

CCS CONCEPTS

• Human-centered computing → Empirical studies in HCI.

KEYWORDS

Algorithm, explanations, transparency, user control

ACM Reference Format:

Malin Eiband, Sarah Theres Völkel, Daniel Buschek, Sophia Cook, Heinrich Hussmann. 2019. When People and Algorithms Meet: User-reported Problems in Intelligent Everyday Applications. In *24th International Conference on Intelligent User Interfaces (IUI '19)*, March 17–20, 2019, Marina del Ray, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3301275.3302262>

1 INTRODUCTION

Algorithmic decision-making has permeated many interactive systems that people use on a daily basis (e.g. film recommendations, social networks, navigation). Based on Jameson and Riedl [25], we define that an intelligent system “embodies one or more capabilities that have traditionally been associated more strongly with humans than with computers, such as the abilities to perceive, interpret,

learn, use language, reason, plan, and decide”. Hence, they select what kind of information is to be considered relevant, influence which content users do or do not see, draw inferences from user behaviour, and shape future behaviour [18].

This algorithmic decision-making poses difficult challenges for human-computer interaction (HCI), for example with regard to established usability principles, such as easy error correction and predictable output [3, 13]. As a result, HCI and related fields have recognised the need to support users during interaction with intelligent systems [4, 11, 22]. Potential problems that users may face in interactions with such systems have been examined on either a general level or tied to evaluations of specific prototypes:

For example, on the general level, people may lack awareness of algorithmic decision-making [16]. Moreover, uncertainty and lack of knowledge about such decision-making may cause what has been called “algorithmic anxiety” [26]. Related, researchers have also observed “algorithmic aversion” [12], when users put less confidence in algorithmic than human predictions, even if the latter are less accurate.

On the other end, possible solutions for supporting users in interactions with intelligent systems focus on specific research prototypes, for example, in work on explanations [30], scrutability [27], interactive machine learning [13], and end-user debugging [29].

We argue that the actual meeting point of humans and intelligent systems *in practice* today can be located somewhere in between these general and prototype levels – in publically available intelligent everyday applications. However, there is yet little empirical evidence on practical problems with such intelligent everyday applications. The research community thus risks that its concepts and prototype solutions remain decoupled from everyday problems and user needs. Moreover, unawareness of current practical problems may hinder addressing the right problems in future research.

In this paper, we aim to assess user problems in intelligent everyday applications, as well as users' coping strategies and wishes for support. In particular, we investigate the following questions:

- (1) *Which problems do users encounter when using intelligent everyday applications?*
- (2) *What kind of support do users want for which problem?*

To address these questions, we analyse problems as reported by users online via app reviews that can be attributed to algorithmic decision-making. We focus on three apps on the Google Play Store that are used in everyday contexts, namely Facebook, Google Maps and Netflix. Using sentiment analysis and topic modelling, we reveal problems which we then use as a basis for a follow-up online survey (N=286). In the survey, we describe the extracted problems as scenarios to collect participants' experiences in these situations and to assess their coping strategies and need for support.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IUI '19, March 17–20, 2019, Marina del Ray, CA, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6272-6/19/03...\$15.00

<https://doi.org/10.1145/3301275.3302262>

Our contribution is threefold: (1) We provide empirical evidence for problems with intelligent everyday applications as reported by users, (2) we extract key problems as well as coping and support strategies as suggested by users, and (3) we present and discuss implications for supporting users in intelligent everyday applications. Our work thus connects and structures practical problems and research directions. We hope to facilitate closer understanding of problems users face when interacting with intelligent systems in everyday contexts.

2 BACKGROUND AND RELATED WORK

We focus on intelligent systems that mediate *everyday* tasks and practices [41].

2.1 Assessing User Problems with Intelligent Systems

The impact of such systems has received considerable attention in the last years. The application context of social media in particular has sparked interest, both among the general public and HCI researchers. For example, Bucher [6] examined and catalogued situations in which users become aware of and experience algorithmic decision-making on Facebook, such as when they note being profiled or “found” by the algorithm. An example is drinking a cup of coffee while seeing an ad for the same coffee brand in the feed.

Eslami et al. [16] found that users’ awareness of the Facebook news feed curation algorithm was generally low: More than half of their participants (62.5 %) did not know that the content they see on Facebook is algorithmically selected, but had rather assumed that friends and family were actively hiding posts from them. In another study the year after [15], Eslami et al. assessed user beliefs about Facebook’s curation algorithm, and how such beliefs affect their behaviour. A similar study was conducted by Rader and Gray [37], who classified a wide range of user beliefs about Facebook’s news feed algorithm.

Other specific application areas in which user-reported problems have been assessed are health and well-being [14] as well as sharing economy platforms [26]: The work by Eiband et al. [14] addressed the problem that users erroneously thought that a human, and not an algorithm, assembled their fitness workouts. Jhaver et al. [26] looked into the effects of algorithmic evaluation on Airbnb hosts. They found that uncertainty about the workings of the algorithm and perceived lack of control caused what the researchers called “algorithmic anxiety”.

2.2 Addressing User Problems in Intelligent Systems

Given the problems and related negative experiences users may encounter in interactions with intelligent systems, supporting users in such situations has been recognised as a crucial challenge for HCI and related fields [4, 11].

A particular concept which has received attention in this context is seamfulness [1, 2, 23]. Seamful (in contrast to seamless) technology design makes system properties apparent to users [8]. In terms of intelligent systems, a seamful approach might expose (part of) the algorithmic decision-making at work in the interface [23]. For example, Eslami et al. [16] developed a prototype to illustrate the

effect of algorithmic decision-making on the content users get to see. In another example, Munson et al. [34] built a tool to make users aware of possible biases introduced by the algorithm.

Another substantial body of work focuses on helping users make sense of intelligent systems, for example through explanations. Such explanations often target the algorithmic decision-making process or a particular output instance [36] and their potential benefits have been investigated in diverse studies (e.g. [9, 10, 28, 30, 31, 35, 36], also cf. [1]).

A third related line of research investigates interactivity of machine learning: Here, users and their feedback and corrections to the system are an integral part of the machine learning process (see [13] for a review of the field). Related, Sarkar [40] introduces the term “metamodels” to describe models about machine learning models as structure for interaction. In particular, these metamodels seek to facilitate user understanding and assessment of the correctness of the system’s workings.

As an overarching view on interaction with intelligent systems, Alvarado et al. [2] propose the concept of “algorithmic experience”. This concept includes fostering awareness of algorithmic decision-making and how it works and deliberately activating or deactivating algorithmic influence.

2.3 Linking Problems and Solutions

As shown in the previous sections, several lines of research seek to assess and address problems which users have with intelligent systems and algorithms. However, to the best of our knowledge, a large-scale empirical assessment of current everyday problems with intelligent systems is still missing. Moreover, we see the need to consider possible solutions for support *along* with the identified problems since these solutions may be specific to particular situations or contexts. For example, Bunt et al. [7] found that the perceived cost of reading explanations in everyday applications tended to outweigh their benefits. Finding problems and solutions as well as a link between them motivates our research presented in this paper.

3 ANALYSIS SCOPE

We focused on three everyday apps which incorporate intelligent algorithmic decision-making: *Facebook*, *Netflix* and *Google Maps*. Social media, recommenders for entertainment and navigation are investigated and discussed in the literature as interesting application domains for interaction with intelligent systems (e.g. [2, 15, 39, 41]). We found Facebook, Netflix and Google Maps to be the most popular representatives of these domains as per their number of downloads in the Google Play Store in July 2018.

Facebook hosts one of the biggest social networks on the web with more than 2.23 billion active users in 2018 [17]. Facebook users are presented with a news feed that is algorithmically curated by taking into account user behaviour such as *likes* [16]. Facebook does not disclose its algorithm which resulted in the formation of folk theories about its working among users [15].

Netflix is a recommender system for streaming films and TV shows. To facilitate users’ choice of a film or TV show, Netflix uses a variety of supervised and unsupervised machine learning techniques for personalised recommendations [19].

	Review Period	Scraped	≤ 3 Stars	Sentiment	Topic
Facebook	15 Aug – 19 Aug 2018	10,000	2,894	2,216	2,196
Netflix	23 Jul – 19 Aug 2018	10,000	2,887	2,345	2,340
Google Maps	30 Jun – 19 Aug 2018	15,448	4,015	3,019	3,018

Table 1: Overview of dataset sizes throughout analysis.

Google Maps is the most popular navigation app which provides real-time GPS navigation, traffic and transit information as well as points of interests such as restaurants and events. Based on Dijkstra’s shortest path algorithm, Google Maps uses deep learning and supervised machine learning techniques to suggest routes [24].

4 DATA ANALYSIS METHOD

Investigating problems with intelligent everyday applications is challenging. Terminology like *algorithm* or *intelligent system* is often out of a layperson’s terms. Simply asking users about their experience with such applications therefore will likely not yield fruitful answers. To overcome this obstacle, Bucher [6] examined people’s personal stories about the Facebook algorithm through tweets and interviews. We followed a similar approach but chose to analyse *reviews* to increase the sample size. Reviews reflect users’ attitude towards an application and document the experiences they make [21, 32]. We scraped reviews from the Google Play Store and extracted and clustered user-reported problems both with machine learning and manual analysis.

4.1 Data Acquisition

We built a web crawler to scrape the latest 10,000 reviews for *Facebook* and *Netflix* respectively, as well as 15,448 reviews for *Google Maps* from the US Google Play Store. Since we wanted to analyse reviews and problems in detail, we decided to take a *snapshot* of current reviews instead of scraping millions of reviews over a longer time period. This sample size allowed us to include manual analysis and human interpretation in addition to quantitative data analysis. We reflect further on this choice in our discussion. An overview of the dataset sizes throughout analysis can be found in Table 1.

4.2 Preprocessing

In line with findings by Maalej et al. [32], a first manual inspection of the data indicated that most reviews praised the app or simply repeated the star rating. Since these reviews were uninformative for our analysis of *problems* with intelligent apps, we excluded all reviews with higher ratings than three stars (out of five stars). As a consequence, the dataset size was decreased by 71 to 74% respectively (cf. Table 1). We also removed reviews written entirely in non-Roman script.

To obtain a more nuanced impression of users’ attitudes towards the apps, we conducted a sentiment analysis of the remaining reviews as recommended by Maalej et al. [32]. We used the Google Cloud Natural Language API¹ to assign each review a score and a magnitude value. The score indicates a positive or negative emotion of a review, while the magnitude refers to the amount of emotional

content within a review [20]. For our analysis, we selected all reviews with negative emotional content (score ≤ -0.1). Since the magnitude is proportional to a text’s length, we accepted all reviews with a magnitude ≥ 0.1 . The final number of reviews for each app can be found in Table 1.

4.3 Extracting Problems from Reviews

We adopted a two-step data analysis approach with machine learning and manual analysis to identify user-reported problems in the remaining 7,559 reviews. This allowed us to combine the benefits of both data-driven analysis (e.g. covering a large dataset) and human interpretation (e.g. insight into subtle nuances of user problems).

4.3.1 Statistical Topic Modelling. In the first step, we applied topic modelling to break the data down into topics. In particular, we used Mallet’s (MACHINE Learning for Language Toolkit) implementation via Gensim [33] of Blei et al.’s Latent Dirichlet Allocation (LDA) algorithm [5]. We applied this modelling approach to each app’s set of reviews separately.

For this step, the reviews were further preprocessed by removing punctuation, superfluous characters (e.g. emojis), and stopwords (i.e. common words such as “the”), and through lemmatization and tokenization. We incorporated bigram and trigram representations to allow the model to account for phrases that consist of more than one word (e.g. “please bring back”). This removed reviews consisting of stopwords only.

The main hyperparameter for the LDA model is the expected number of topics. To inform this, for each application we tested 25 models with topic numbers ranging from two to hundred. We then chose the model with the highest coherence score. In cases without a peak coherence score, we selected the smallest number of topics that resulted in a score of at least 0.5 to keep the number of topics at a manageable size for our manual analysis step.

4.3.2 Manual Analysis. In the second step, we manually labelled the topics based on the top ten keywords as well as the ten most representative reviews associated with a topic as discovered by the model (900 reviews out of 7,559 in total). That is, we manually extracted and coded the core *problem* underlying each topic. This was done by two researchers independently, who then synthesised their analysis. Finally, we aggregated the resulting problems into more abstract *problem categories* identified across all three apps.

5 RESULTS: USER-REPORTED PROBLEMS

Topic modelling yielded 22 (Facebook) and 34 (Netflix and Google Maps, respectively) topics in total. For 21 (Facebook), 33 (Netflix) and 34 (Google Maps) of these topics we extracted an underlying core problem. Reviews in the remaining two topics (Facebook and Netflix) did not indicate a distinct underlying core problem and were therefore excluded for further analysis. Other topics related to bugs and usability issues were ignored as well. This left three (Facebook), eleven (Netflix) and 13 (Google Maps) topics. In total, these topics covered 12.5% (Facebook, 274 reviews out of 2,196), 29.7% (Netflix, 694 reviews out of 2,340) and 35.5% (Google Maps, 1,072 reviews out of 3,018) of the datasets used for topic modelling. We aggregated our results into problem categories along (up to) four steps of a basic pipeline of algorithmic decision-making and

¹<https://cloud.google.com/natural-language/>

related interactions: (1) knowledge base, (2) algorithm, (3) user choice, and (4) user feedback. In this way, we not only focus on the algorithm itself but also capture the interplay between algorithmic decision-making and users more broadly.

All identified problems, their assignment to the problem categories and exemplary reviews can be found in Table 2. All user quotes have been reproduced with original spelling and emphasis.

5.1 Facebook

5.1.1 Problems with Curation Algorithm. The overall content curation of the Facebook news feed seemed biased to users and they accused Facebook of deliberately concealing posts. Furthermore, users reported problems with the content composition and ranking of the feed: They complained that they were not able to arrange the news feed in chronological order and that they missed posts by their friends and family. Instead, they had the feeling of being “spammed with posts from pages [they] follow” and advertisements.

5.1.2 Problems with User Choice. Users expressed the wish for more control over their feed. For example, they asked for the option to remove intelligent components, such as the *People You May Know* section or to turn off *Marketplace* notifications.

5.2 Netflix

5.2.1 Problems with Knowledge Base. Users criticised Netflix’ basis for recommendations, namely the available films and shows. First, they complained about the uninteresting selection of films in general (“[...] for a film lover myself, it’s grown stale. [...]”). Furthermore, users claimed that the available films were outdated and their favourite films had been removed (“When you see that they have added new movies that you are interested in, make sure to watch them quickly because they’ll probably remove it within a week.”). They also accused Netflix of a biased content selection by favouring their own original content over external content (“[...] now I’ve noticed that Netflix is starting to kind of put on a lot of their own stuff and they’re taking out the movies and TV shows that I want to watch [...]”). Finally, users disapproved of incomplete TV shows for which not all seasons are available.

5.2.2 Problems with Recommender Algorithm. Users complained about a mismatch between Netflix’ recommendations and their interests (“You guys don’t even give good recommendation anymore”). In particular, they noted that notifications about new releases and top picks were not to their taste (“Just got one [notification] telling me *Better Call Saul 3* is out. I don’t care and don’t want that clutter in the way of my workflow”).

5.2.3 Problems with User Choice. According to the reviews, users would like to have more options for selecting interesting content themselves. They claimed that Netflix’ current interface makes manual selection difficult due to missing information about films and TV shows. For example, users requested a synopsis and description about an item since they “[needed] more than a still photo and search keywords to decide if [they wanted] to watch something”. Moreover, the interface misled users and caused confusion about the language of an item: While all film titles were shown in their system’s language, users did not know whether an item was actually available in this language until they started watching.

Finally, the reviews indicated an incomplete search coverage since the search functionality only returned “irrelevant results”.

5.2.4 Problems with User Feedback. Users indicated that the binary rating system (thumbs up/down) Netflix currently uses is not sufficient to provide meaningful feedback, which in turn might be reflected in the relevance of recommendations. Users noted that the provided certainty of a recommendation *always* seemed to be very high. Hence, a user asked: “How am I supposed to know if I should actually watch the movie if every movie is a match”.

5.3 Google Maps

5.3.1 Problems with Knowledge Base. The reviews indicated problems with the app’s basis for routing suggestions. For example, users complained about inaccurate location sensors when Google Maps “is unable to pin point [their] location”. In particular, users experienced these problems off-road. For example, a user described that her location “appears to shift kilometre or more in seconds” when she goes hiking. In addition, users encountered missing or incorrect map information when the application sent them “off to a different location than you expect”.

5.3.2 Problems with Routing Algorithm. Users experienced several problems with the routing algorithm. First, the time estimate for a route seemed inaccurate and inconsistent to them. For example, one user recounted: “Keeps adding completely random routes, as well as 10 minutes extra travel time to the same route that I take every day”. Other problems were unstable route selection and incorrect directions which led to users getting lost. Furthermore, users reported that Google Maps announced last minute turns and surprised them with sudden route changes or suggestions of dangerous and illegal manoeuvres such as a U-turn on the highway. Users also strongly criticised that the routing algorithm ignored their settings. Although they had explicitly excluded tolls, the app guided them via toll roads in the end.

5.3.3 Problems with User Choice. Users experienced problems when they tried to override algorithmic decision-making by manually selecting a route or location. They indicated the need for more and more specific routing options to match their needs. For example, they recommended a mode for trucks or big cars which avoids narrow streets. Users would also appreciate alternative routing criteria, such as a scenic route. They further found that their manually selected routes were suddenly overwritten by Google Maps “without telling” them. Finally, users were frustrated when their knowledge was ignored but felt that their routing decision would have been better than Google Maps’. In particular, they complained that Google Maps repeatedly led them through a construction site.

5.3.4 Problems with User Feedback. Users encountered problems when trying to give feedback or corrections to the app, such as missing or incorrect data points. They claimed that Google Maps “has failed to respond to fix the flaw”. For example, a user indicated that Google Maps “consistently thinks [her] home is on another continent”. Similarly, users reported that Google Maps did not allow them to easily report road construction sites. The reviews suggested that these reports were not taken into account “even if reported by many users for 4 weeks”.

Problem Category	Problem	# Rev.	Example Review
<i>Facebook</i>			
Algorithm	Biased content curation	117	"Cowards, censoring videos made by Republicans! [...] Why not let people post what they want?! I want to see Will Witt's videos, but you don't let me because you don't agree with him. [...]"
	Selective content composition and ranking	81	"[...] You know what i miss, seeing actual post from friends. Not ads, not what my friends liked or commented on. What they post and share. Not spammed with post from pages I follow. Things ACTUAL FRIENDS POST. [...]"
User choice	Insufficient options for intelligent components	76	"I get irritated with that People You May Know. I tried that I don't want to see this option but it doesn't work anymore. Please kindly give us any option to remove this thing."
<i>Netflix</i>			
Knowledge base	Uninteresting content	81	"It used to be good, but now for a film lover myself, its grown stale. [...] Unfortunately, once you watch many of whats on there, it just gets bland and you wind up scrolling through things you've already been through."
	Outdated content	64	"[...] I want to say that Netflix is getting really old, that and it picks some of the worst movies I've ever seen especially when it comes to horror movies. [...]"
	Biased content selection	44	"I used to love Netflix a lot cuz they had all the shows I wanted and movies but now I've noticed that Netflix is starting to kind of put on a lot of their own stuff and they're taking out the movies and TV shows that I want to watch [...]"
	Incomplete content	74	"At first i thought its really worth to pay BUT when i'm on the end season of series, its incomplete [...]"
Algorithm	Mismatch between recommendations and user interest	59	"You guys dont even give good recommendation anymore. [...] For example how in anyway is Godzilla 2 related to midnight run with Robert de niro"
	Mismatch between notifications and user interest	39	"I'm still receiving notifications from the app. Just got one telling me Better Caul Saul 3 is out. I don't care and don't want that clutter in the way of my workflow"
User choice	Missing information	96	"They just removed the synopses for all their shows... What gives Netflix? I need more than a still photo and search keywords to decide if I want to watch something"
	Missing information	26	"They've now removed all descriptions of the show or movie. [...] remember exactly what every single thing you ever put in your watchlist is about"
	Misleading information	59	"Its just I speak English and it don't even tell you its in another language until you start watching it. Please add more English/American movies or at least state their not in English."
	Incomplete search coverage	94	"The search system is poor. A search returns with irrelevant results. Can you please make it basic search?"
User feedback	Insufficient feedback options	59	"The rating system is still horrible, every movie I look at says 98% match like how am I supposed to know if I should actually watch the movie if every movie is a match. Bring back the star system."
<i>Google Maps</i>			
Knowledge base	Inaccurate user location	206	"Maps is unable to pin point my current location."
	Inaccurate user location (off-road)	85	"This works well when you are on a road. However, take it off road, and its useless. Its laughable when I go hiking, and share my location with my wife. If I'm not on a road, my location appears to shift a kilometre or more in seconds: I even appear to walk on water, crossing lakes and rivers without any effort!"
	Incorrect map information	79	"BEWARE: This app could send you off to a different location than you expect. Check the address multiple times and make sure your address is correct first [...]"
Algorithm	Inaccurate time estimate	80	"[...] Keeps adding completely random routes, as well as 10 minutes extra travel time to the same route that I take every day. Also it screws up distances by several miles and adds about 15 minutes to known estimated travel times for zero reason."
	Unstable route selection	53	"Even though we have told them of a major error on a heavily traveled tourist route numerous times, they refuse to listen. Trucks are getting ticketed and tourists are getting lost at a rate of at least 15-20 per hour. [...]"
	Incorrect directions	81	"[...] gives incorrect directions or dramatic unnecessary detours [...]"
	Unsuitable turn advice	67	"It will tell you it has a faster route mid drive and automatically change routes on you unless you hit an option on your phone which is not only illegal but dangerous! [...] Then when I finally got closer to my destination it told me to go down a two way street then through a one way no entry section, if a car had come around the corner while I was trying to reverse and turn around I could have been killed!"
	Routing ignores user settings	87	"[...] A month after my trip I received a toll violation of \$26 for being in the carpool lane. I had tolls turned off on the app but I didn't see the option on Android Auto. [...] (by the way you owe me \$26)"
	User choice	Lack of routing options	65
Overwriting user-selected route		82	"It offers me a route and I choose it because I want to take the scenic route. Then, without telling me just puts me back on the quickest route. Which drives me insane - not everyone is trying to get places fast some of us like to see the world while do it."
Ignoring user knowledge		59	"For the road closures that need satnav help it goes out of its way to Always route through them no matter what even if it takes longer."
User feedback	Ignoring user map corrections	69	"Consistently insists my home is on another continent. Google have failed to respond to requests to fix the flow, indeed they wont even acknowledge the problem. If it cant get my address right, how can we know other addresses are right Well, actually there are many corrections I've suggested, but I'm ignored, even though I'm right and Maps is wrong. Untrustworthy!"
	Ignoring user route corrections	59	"Google maps also doesn't have a way to report raod construction closures [...]"

Table 2: Results of the review analysis: User-reported problems related to intelligent everyday applications that we further aggregated to more abstract problem categories. The columns on the right show the number of reviews assigned to each problem category and an example review to illustrate the underlying problem (original spelling and emphasis).

6 ONLINE SURVEY

We conducted an online survey to extend our understanding of the extracted problems with user feedback beyond the Google Play Store reviews. This also allowed us to collect statements on these problems from people who did *not* write an app review about it. Moreover, we used the survey to assess users' own suggestions for solutions to the extracted problems.

6.1 Survey Design

In the survey, we described each extracted problem as a scenario, similar to how they were expressed in the original reviews. For example, we illustrated the *Google Maps* problem *Overwriting user-selected route* in the following scenario:

Tom and Paula go for an outing. They selected the most scenic route on Google Maps. While on the road, Google Maps automatically switches to the shortest route.

For each scenario participants were asked to (1) indicate whether and how often they had encountered a similar situation (Likert scale), (2) describe how they coped with this situation (open question), and (3) describe how the app could support them in this situation (open question). Each participant was randomly assigned ten scenarios to keep the survey short, especially considering the open questions.

Finally, participants were asked to indicate their average usage of the three apps, their smartphone or tablet's operating systems, as well as demographic information.

6.2 Participants

We distributed the questionnaire via university mailing lists and social media. It was completed by 287 participants. Before analysis, we checked and cleaned all collected data. We excluded one participant from the final dataset due to invalid answers. The remaining sample consisted of 286 participants (68.5% female; mean age 24.5 years, range 17-67 years). 175 participants used an Android-based smartphone, 112 participants an iPhone, and nine participants a Windows-based smartphone. 64.0% used Facebook at least weekly, 48.6% used Netflix at least weekly, and 78.7% used Google Maps at least weekly. Participants had a chance to win €50 in cash.

6.3 Data Analysis

The first two authors reviewed all qualitative data to derive a coding scheme for the two open questions, namely participants' coping strategies and their wishes for support. Codes were extracted independently and then discussed together to construct a code book. A random sample of 10% of the answers of each scenario used in the survey was then coded by both coders independently using the code book. The results were compared and discussed to eliminate any discrepancies until a consensus was reached. The dataset was then split between both coders for the final coding. Table 3 shows the number of participants who were presented with at least one scenario in a problem category and those who indicated to have experienced one of these scenarios at least once. The last column shows how often the problems within one problem category had been experienced as median (scale: once, monthly, weekly, daily).

	Problem Category	# Presented	# Experienced	Median
Facebook	Algorithm	220	138	weekly
	User choice	91	71	weekly
Netflix	Knowledge base	167	92	monthly
	Algorithm	87	44	monthly
	User choice	200	93	monthly
	User feedback	104	37	monthly
Google Maps	Knowledge base	208	114	monthly
	Algorithm	235	99	once
	User choice	207	85	once
	User feedback	177	41	once

Table 3: Number of participants who were presented with a scenario relating to a problem category of the respective apps. Out of those, number of participants who had experienced one of these scenarios at least once. The last column shows the median of how often the problems within one problem category had been experienced (scale: once, monthly, weekly, daily).

Algorithm	User Choice		
	<i>Coping Strategies</i>		
Control content directly	36	Control content directly	15
Actively search for content information	27		
Check settings	6		
Influence algorithm indirectly	3		
Search for explanations	3		
	<i>Support Strategies</i>		
Options to control news feed	43	Choice of intell. components	18
Prioritise in a user-centered way	21	Options to control news feed	12
Default chronol. order of news feed	20	Prioritise in user-centered way	5
Show non-personalised content	13		
Let users filter content, not algorithm	13		
Give explanations	18		

Table 4: Facebook: Selected user-reported coping and support strategies for each problem category (with number of mentions).

7 RESULTS: COPING STRATEGIES & SYSTEM SUPPORT

Participants reported a great variety of coping strategies and ways in which they would like to be supported by the system. Within the scope of this paper, we restrict our account to those strategies which we deem most relevant in terms of our research questions and for HCI research in general. In the following sections, participants' quotes have been translated to English where necessary or have been reproduced with original spelling and emphasis. Since our final sample size only includes participants who completed the survey, participant IDs may be bigger than 286.

7.1 Facebook (Table 4)

7.1.1 Addressing Problems with Curation Algorithm. Coping Strategies: Participants tried to influence Facebook's news feed mostly by directly controlling which posts are shown: They disliked pages

Knowledge Base	Algorithm	User Choice	User Feedback
<i>Coping Strategies</i>			
Search for content manually	19 Search for content manually	14 Get recommendations somewhere else	44 Get recommendations somewhere else 6
Get recommendations somewhere else	6 Ignore recommendation	10 Search for content manually	27
Try to “reset” algorithm	3 Change settings	3	
<i>Support Strategies</i>			
Make content selection transparent	23 Allow users to give feedback	12 Provide more information about items	50 Improve rating system 26
Provide more information about items	6 Use more fine-grained input	4 Include user reviews	10 Include user reviews 6
Let users teach the system	2 Let users teach the system	4 More filtering options	6

Table 5: Netflix: Selected user-reported coping and support strategies for each problem category (with number of mentions).

they followed, unfriended people, or marked advertisements as spam. Many participants felt that the best way to not miss out on their friends’ news was to visit their page directly instead of relying on the feed. P246 stated that she “got into the habit of visiting profiles of people whose news [she’s] interested in”. A few participants tried to indirectly nudge the algorithm to show specific content. For example, they visited their friends’ pages or liked their posts and wrote personal messages more often. Some participants also indicated that they had changed their settings to influence the content of the news feed, or had searched online for explanations of the algorithm’s workings.

System Support: Although many participants already used options to organise their news feed, more settings were their most prominent wish for support. For example, P521 would like to be able to configure the feed in a way that posts by friends are displayed at the very top. This was also reflected in the claim that the algorithm should prioritise posts in a more user-centred way. P518 suggested that the frequency of chatting with people as well as the frequency of visiting their pages should be considered. Interestingly, many participants would like the news feed to be ordered chronologically by default. For example, P148 claimed to “just want a news feed in which all friends and subscriptions appear in chronological order”. Some participants wanted to be able to switch algorithmic curation on and off, or even stated that content should only be filtered by the users, not by an algorithm. Participants were also interested in non-personalised content even “contrary” (P48) to their opinion, or demanded support via explanations to better comprehend why a particular post was shown.

7.1.2 Addressing Problems with User Choice. Coping Strategies: Problems with user choice were similarly addressed using available options to control the feed content.

System Support: Participants called for choice regarding the (intelligent) features that are included in the news feed, such as the “people you may know” feature presented in one of the scenarios. Moreover, the demand for more options or settings to control the feed was repeated in this problem category, as well as the claim that the algorithm should more strongly take into account users’ needs and wishes.

7.2 Netflix (Table 5)

7.2.1 Addressing Problems with Knowledge Base. Coping Strategies: Participants experienced problems related to Netflix’ knowledge base, such as the lack of interesting content. They then tried to

search the database manually for relevant films and shows to watch or looked elsewhere for recommendations, often online. Others tried to get access to new content by creating a new additional account, to “discover new films without the pre-built algorithm” (P253).

System Support: Participants suggested to make content selection more transparent to users, for example by displaying newest releases as the very first film category in the interface or by adding release dates. Participants also wished for more detailed information about films and shows. Some even mentioned that the selection of content should be based more on users teaching the system about their preferences.

7.2.2 Addressing Problems with Recommender Algorithm. Coping Strategies: Recommendations that participants perceived to be irrelevant were often ignored in favour of manual search. Some participants also stated that they had changed their profile settings to influence the algorithm’s suggestions.

System Support: In this problem category, participants generally called for expressive feedback options. They suggested that the algorithm should take their preferences into account on a more fine-grained level, such as through more detailed film categories and genre types. For example, P393 suggested “mood-dependent filtering”. Others stated that the system could explicitly support discovery with “films that do *not* match [a users’] preferences” (P253).

7.2.3 Addressing Problems with User Choice. Coping Strategies: Selecting films and shows on Netflix was often accompanied by using other sources for recommendations – mostly asking friends or searching on popular rating websites. Interestingly, only one participant explicitly stated that she had followed Netflix’ matching score (P352).

System Support: The majority of people presented with this problem called for more detailed descriptions of films, shows, and actors as a basis for accepting a recommendation. For example, they asked to include reviews and ratings from other platforms, and to bring back user reviews. Such reviews sometimes even seem to be more relevant to users’ choices than system recommendations. For example, P374 said that she would also like to see “films that match my films only by 50 percent, but have been high-rated (by other users)”. Similarly, P415 stated that the system “should show me more shows that all people like [...], not only those that I will probably like”.

Knowledge Base	Algorithm	User Choice	User Feedback
<i>Coping Strategies</i>			
Manually without Google Maps	42 Deal with uncertainty	20 Manually without Google Maps	33 Manual without Google Maps 17
Reload	21 Manually with Google Maps	17 Manually with Google Maps	24 Manually with Google Maps 7
Recalculation	19 Manually without Google Maps	10 Deal with uncertainty	9 Feedback to Google Maps 5
<i>Support Strategies</i>			
Allow users to give feedback	13 Adapt to users	15 More options for routing	16 More routing alternatives 8
Provide information about uncertainty	7 More routing alternatives	13 Inform early about critical parts	16 Allow user to give feedback 5
	Inform early about changes	11 Ask for permission before changing route	8

Table 6: Google Maps: Selected user-reported coping and support strategies for each problem category (with number of mentions).

7.2.4 Addressing Problems with User Feedback. Coping Strategies: The current way of giving feedback in Netflix via “thumbs up/down” was often criticised by participants for not being fine-grained enough. For example, P224 said that “the old rating system was better, the new one does not provide a neutral expression or a mediocre one, but you have to choose one of two options”. P99 even said that she “just got annoyed. I have not rated any films since out of frustration”.

System Support: Some participants used the current rating option nevertheless in want of alternatives (P118), but wished for a general improvement of the rating system. Many called for bringing back the star rating that Netflix had once used in order to give more nuanced feedback. Other suggestions included an option to write reviews, or the possibility to answer a short questionnaire at the end of a film.

7.3 Google Maps (Table 6)

7.3.1 Addressing Problems with Knowledge Base. Coping Strategies: To cope with problems regarding the app’s knowledge base, participants indicated that they tried to find a location or route manually without Google Maps, for example by asking pedestrians or getting their bearings in the surroundings. Some participants also tried to reload the app or restarted GPS to trigger a recalculation.

System Support: Regarding support strategies, most participants simply suggested to improve location sensing. Moreover, participants asked for an option to give feedback to the system. For example, P288 wrote: “Since I knew where I am, it would have been helpful to simply tell the app, this is my position, trust me”. To deal with situations in which location accuracy cannot be guaranteed participants suggested that the app should “inform about these uncertainties”.

7.3.2 Addressing Problems with Routing Algorithm. Coping Strategies: To prevent problems with routing, users pointed out that they tried to deal with uncertainty, e.g. by “always adding an additional 20 % on Google Maps’ travel time” (P273). Furthermore, participants reported that they selected the route manually in Google Maps. For example, P495 wrote that she “compared the different routes regarding their length in Google Maps” when she thought that Google Maps suggested a detour. Other participants indicated that they abandoned the app when facing routing problems. When confronted with unexpected routing suggestions, participants explained that they drove more carefully, but this also negatively influenced their feeling of the routing reliability.

System Support: To overcome these problems, participants indicated that Google Maps should adapt to their needs. For example, they suggested that the algorithm should include a user’s usual driving or walking speed and adapt the travel time accordingly. Additionally, participants suggested that the app should provide more routing alternatives, allowing users to choose between them. Finally, participants emphasised that Google Maps should inform early about changes of the route so that users “have the opportunity to react in time” (P15).

7.3.3 Addressing Problems with User Choice. Coping Strategies: Several people experienced difficulties with Google Maps when they wanted to manually choose a route and override the algorithm’s suggestions. As a result, most participants abandoned the app and tried to find a route manually. Alternatively, participants described that they deactivated the routing function and simply used the map to manually select a route.

System Support: To overcome these problems, participants suggested more use-case specific options for routing, e.g. for different car/truck sizes or scenic versus shortest route. Participants also expressed the wish to be informed early about critical parts of the route, such as “highlighting difficult parts of the road in advance” (P466), allowing to be prepared or to manually change the route. Finally, participants stressed that Google Maps should “not simply re-route but ask for permission before changing the route” (P17). Several participants also suggested to prevent re-routing by default.

7.3.4 Addressing Problems with User Feedback. Coping Strategies: Participants indicated that they had repeatedly tried to give feedback to Google, but that this feedback was not being taken into account. When having the impression that Google Maps’ suggestions led to delays or detours, most participants ignored the suggested routing in want of options for correction. For example, they manually tried to avoid a construction site by following a diversion, or used the map to manually find the address they were looking for.

System Support: Most participants suggested faster updates of the database to incorporate users’ feedback. Participants also expressed the wish for more routing alternatives they could use to manually select a different route, for example to avoid construction sites. Moreover, participants stressed the need for feedback and correction options. For example, P352 suggested that users should be able to “mark a construction site for them personally, which is detected by the system to provide an alternative route”.

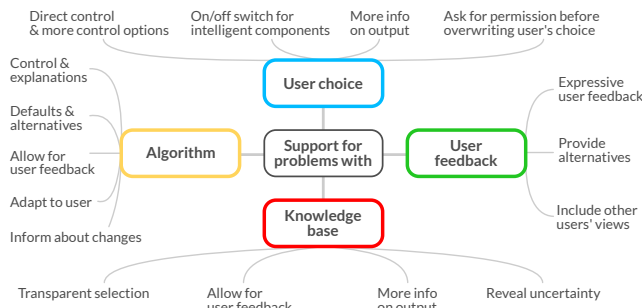


Figure 1: Summary of users' suggested support strategies for their reported problems, aggregated across the studied apps.

8 DISCUSSION

8.1 Implications for Design of User Support

The identified problems and coping and support strategies revealed insights that span the three apps. Figure 1 provides a visual summary. We next discuss design implications.

8.1.1 Let Users Stay in Control. Control issues – and the wish for more control – reoccurred across all applications: Users appreciate suggestions but want to make the final decision themselves. Systems should thus provide means for users to make such decisions on an informed basis: System suggestions should be presented with sufficient information to allow users themselves to assess the value of a suggested item (e.g. film, route, post in news feed).

Our analysis revealed that this information should primarily describe the suggested item *itself* in an accurate way, instead of referring to how and why the suggestion was made. For example, participants called for more information about Netflix recommendations in order to assess their interest in a film or show rather than relying on the information about their personal matching score. In the context of intelligent everyday applications, this introduces a new perspective in contrast to the literature about explainability of intelligent systems, which so far has focused more on supporting users in understanding *why* a system decision has been made (e.g. [27, 29, 31]).

Moreover, the system should not alter a decision that has been made by a user: Users want to have the last say. For example, both our analysis of app reviews and our survey revealed that people repeatedly complain about Google Maps automatically switching to a different route on the way.

8.1.2 Provide more Fine-grained Control Options. Our participants stressed the need for more fine-grained options to control algorithmic decision-making. While the analysis of Facebook revealed that many participants already have assumptions on how to make use of different settings to determine (part of) the content of their news feed, Google Maps users wanted to have more options to control route suggestions. For example, they would have liked to be able to indicate their vehicle type. Netflix users also called for more setting possibilities, such as to state their preferences about genres or mood of film. Supporting these user needs likely requires the integration of control elements into the UI, or reconsidering design and scope of existing such elements.

8.1.3 Explain Interactively. We found an overall need for explanation of system workings in the analysis of reviews and the online survey, even though it was not as prevalent as the desire for more control. Nevertheless, our results hint at the potential practical value of a more interactive approach to explanation, in contrast to the predominantly static one presented in the literature (e.g. [36, 38]): Our participants expressed a desire to *try out* different settings and observe effects on the algorithmic output, possibly “live”. For example, people suggested to include different “modes” for the Facebook news feed, namely chronological order or algorithmically curated. This echoes a recent call by Abdul et al. [1] who suggest to allow users to “[freely] explore the system’s behaviour through interactive explanations”, for example, using interactive visualisations.

8.1.4 Allow Users to Turn Intelligence On and Off. We found that a large number of participants used the systems without their intelligent features: For example, users navigated themselves without Google Maps’ intelligent algorithm, they tried to find film recommendations by asking friends or searching online, or they directly visited friends’ profiles they were interested in on Facebook. This behaviour was often used as a coping strategy for problems with the respective system, but could also be a feature of intelligent systems to stress user control or foster “algorithmic awareness” (cf. [26]).

8.1.5 Honour Expressive User Feedback and Corrections. Giving feedback and corrections to intelligent systems has been recognised as an integral part of interaction with such systems, in particular in the area of interactive machine learning [3]. Yet, options for doing so in the practical systems we analysed in this paper were sparse or difficult to find for users, or were not seen as helpful in their current state. For example, Netflix offers a possibility for giving feedback, but the binary approach of “thumbs up” and “thumbs down” was heavily criticised both in the reviews and our online study. Users appreciate a more fine-grained, meaningful feedback system – a 5-level star rating was often mentioned as a preferred approach. Moreover, a seemingly obvious, but crucial follow-up issue is to actually *take feedback into account* and provide a means of confirming this to the user.

8.2 When in Doubt, Trust in Humans Seems to Outweigh AI

Although our work did not focus on trust in intelligent systems in particular, our results indicate a tendency to favour human over algorithmic decision-making, not only in sensitive domains but also in comparably low-risk intelligent everyday applications. Decisions for films and shows on Netflix were mostly accompanied by drawing on *human* recommendations elsewhere, either on rating platforms online or by asking friends directly. Moreover, some users of Google Maps estimated the arrival time themselves: From their experience, they had found that their own estimates better matched the time they actually required. These findings suggest that intelligent everyday applications are used as support for dealing with and organising daily issues and thus for *augmenting* users’ abilities, but not *replacing* or readily exceeding them. This picture might be inverted for intelligent systems with abilities that humans lack,

such as high-speed reaction times (e.g. self-driving cars), introducing randomness (e.g. a drawing program) or performing repetitive tasks (e.g. industrial robotics), which might present an interesting issue to explore in future work.

8.3 Locating and Motivating Designs for User Support

Reflecting on the identified problems, we highlight that they occur at different stages of algorithmic decision-making: First, there are problems with *input* and *knowledge base* of said systems (e.g. GPS data, available maps and films). Second, problems may occur related to the *inference* process (e.g. content prioritisation and recommendation). Third, problems can appear when the system presents or acts according to the *output* of said inference (e.g. missing content and information, overwriting user choice). On a high level, these stages and associated problems indicate *points of action* with which researchers and practitioners may explicitly frame and motivate a specific support concept or system. For example, a UI concept that adds a “revise” button next to algorithmic decision outputs could be explicitly motivated as addressing the output stage and the corresponding set of identified practical problems. In this way, a researcher or practitioner working on the concept could link her proposed solution to a problem category and specific practically occurring example problems.

8.4 Reflection on Methodology

We combined sentiment analysis and topic modelling with manual analysis to extract and aggregate user-reported problems. A similar combination of topic modelling and manual refinement has been used in recent data-driven literature surveys (e.g. see [1]). This approach allowed us to explore a large dataset in a structured way (via topic modelling) and still look into user problems in detail via human interpretation. We believe that this approach is valuable for analysing experiences with intelligent systems beyond this paper: For example, it could also prove useful to extract reported problems on social networks, blogs, or services like twitter.

As a basis for our analyses, we took a *snapshot* dataset of recent reviews. This could be repeated at a later point in time, or in regular intervals, to see how problems develop. A short-term sample might be influenced by events such as recent update releases. While our dataset contains reviews that relate to such events, the extracted main problems are clearly more long-lived and tied to fundamental app features and characteristics. We thus conclude that our snapshot was not overshadowed by a particular update. On the other hand, we cannot assess if we missed certain problems, especially rare ones. Nevertheless, any reoccurring dominant problem should also be reflected in our snapshot.

In future work, a long-term dataset could be collected and analysed as well, potentially also covering further applications. However, it might then become more difficult to analyse the problems in nuanced detail, since manual analysis with human interpretation is challenging to scale up to millions of reviews. To address this, future analyses could use our work combining topic modelling and manual interpretation as a starting point (e.g. by training a problem classifier model on our extracted problems and corresponding reviews).

Note that our work is based on users’ experiences with the apps and their underlying intelligent algorithms. As our work was conducted without the involvement of Facebook, Netflix and Google, we have no insight into the details of the respective algorithm’s workings at the time of data collection and analysis. Thus, we cannot verify if the problems that users reported can be attributed to actual problems in the respective algorithm. Different experiences might also be caused by different versions of the apps or algorithm. However, to inform support for users (e.g. explanations) it is the very problems that users *experience* that are of interest to us, be they caused by actual malfunction of the system or not.

9 CONCLUSION

The complex nature of intelligent systems with algorithmic decision-making poses difficult challenges for human-computer interaction (HCI) and often violates established usability principles, such as easy error correction and predictable output [3, 13]. This motivates work on supporting users during interaction, for example through explanations [36].

However, there is still little empirical evidence on practical everyday problems which people face when using intelligent applications. As a result, it often remains difficult for researchers to clearly link prototype solutions to empirically well-founded practical user problems.

To address this, we assessed and revealed such problems, as reported by users, by analysing 35,448 public app reviews with sentiment analysis and topic modelling. We then enriched this data with information about users’ coping strategies and support ideas through a follow-up online survey (N=286). Based on this information, we extracted problem types and corresponding implications for designing user support.

In particular, our implications point towards supporting user control, feedback, and corrections, and reconsidering how and what to explain. For example, our results suggest that explanations for recommendations should mainly describe the recommended content itself – the how and why of the recommendation process is less interesting to users in our studied everyday practical context. This stands in contrast to the existing focus in the literature (e.g. [27, 29, 31]),

More generally, our work thus contributes empirical evidence for problem situations and facilitates a closer understanding of practical user problems in intelligent everyday applications.

10 PROJECT RESOURCES

The reviews with annotated problems and categories as well as the coded survey results are available via the project website:

<http://www.medien.ifl.lmu.de/userproblems/>

REFERENCES

- [1] Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y. Lim, and Mohan Kankanhalli. 2018. Trends and Trajectories for Explainable, Accountable and Intelligible Systems: An HCI Research Agenda. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 582, 18 pages. <https://doi.org/10.1145/3173574.3174156>
- [2] Oscar Alvarado and Annika Waern. 2018. Towards Algorithmic Experience: Initial Efforts for Social Media Contexts. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 286, 12 pages. <https://doi.org/10.1145/3173574.3173860>

- [3] Saleema Amershi, Maya Cakmak, William Bradley Knox, and Todd Kulesza. 2014. Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine* 35, 4 (2014), 105–120. <https://doi.org/10.1609/aimag.v35i4.2513>
- [4] Association for Computing Machinery US Public Policy Council (US-ACM). 2017. Statement on Algorithmic Transparency and Accountability. https://www.acm.org/binaries/content/assets/public-policy/2017_usacm_statement_algorithms.pdf, last visited: 18/12/23.
- [5] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3, Jan (2003), 993–1022.
- [6] Taina Bucher. 2017. The Algorithmic Imaginary: Exploring the Ordinary Affects of Facebook Algorithms. *Information, Communication & Society* 20, 1 (2017), 30–44. <https://doi.org/10.1080/1369118X.2016.1154086>
- [7] Andrea Bunt, Matthew Lount, and Catherine Lauzon. 2012. Are Explanations Always Important?: A Study of Deployed, Low-cost Intelligent Interactive Systems. In *Proceedings of the 2012 ACM International Conference on Intelligent User Interfaces (IUI '12)*. ACM, New York, NY, USA, 169–178. <https://doi.org/10.1145/2166966.2166996>
- [8] Matthew Chalmers and Ian MacColl. 2003. Seamful and Seamless Design in Ubiquitous Computing. *Workshop At the Crossroads: The Interaction of HCI and Systems Issues in UbiComp*. 8 (2003).
- [9] Henriette Cramer, Vanessa Evers, Satyan Ramlal, Maarten Van Someren, Lloyd Rutledge, Natalia Stash, Lora Aroyo, and Bob Wielinga. 2008. The Effects of Transparency on Trust in and Acceptance of a Content-based Art Recommender. *User Modeling and User-Adapted Interaction* 18, 5 (2008), 455–496. <https://doi.org/10.1007/s11257-008-9051-3>
- [10] Henriette Cramer, Bob Wielinga, Satyan Ramlal, Vanessa Evers, Lloyd Rutledge, and Natalia Stash. 2009. The Effects of Transparency on Perceived and Actual Competence of a Content-based Recommender. *CEUR Workshop Proceedings* 543 (2009), 1–10.
- [11] Nicholas Diakopoulos. 2016. Accountability in Algorithmic Decision Making. *Commun. ACM* 59, 2 (2016), 56–62. <https://doi.org/10.1145/2844110>
- [12] Berkeley J. Dietvorst, Joseph P. Simmons, and Cade Massey. 2015. Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err. *Journal of Experimental Psychology: General* 144, 1 (2015), 114–126. <https://doi.org/10.1037/xge0000033>
- [13] John J. Dudley and Per Ola Kristensson. 2018. A Review of User Interface Design for Interactive Machine Learning. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 8, 2 (2018), 1–27. <https://doi.org/10.1145/3185517>
- [14] Malin Eiband, Hanna Schneider, Mark Bilandzic, Julian Fazekas-Con, Mareike Haug, and Heinrich Hussmann. 2018. Bringing Transparency Design into Practice. In *Proceedings of the 2018 Conference on Intelligent User Interfaces (IUI '18)*. ACM, New York, NY, USA, 211–223. <https://doi.org/10.1145/3172944.3172961>
- [15] Motahhare Eslami, Karrie Karahalios, Christian Sandvig, Kristen Vaccaro, Aimee Rickman, Kevin Hamilton, and Alex Kirlik. 2016. First I “Like” It, then I Hide It: Folk Theories of Social Feeds. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 2371–2382. <https://doi.org/10.1145/2858036.2858494>
- [16] Motahhare Eslami, Aimee Rickman, Kristen Vaccaro, Amirhossein Aleyasen, Andy Vuong, Karrie Karahalios, Kevin Hamilton, and Christian Sandvig. 2015. “I always assumed that I wasn’t really that close to [her]”: Reasoning about Invisible Algorithms in News Feeds. In *Proceedings of the 2015 CHI Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 153–162. <https://doi.org/10.1145/2702123.2702556>
- [17] Facebook, Inc. 2018. Facebook Reports Second Quarter 2018 Results. <https://investor.fb.com/investor-news/press-release-details/2018/facebook-reports-second-quarter-2018-results/default.aspx>, last visited: 18/12/23.
- [18] Tarleton Gillespie. 2014. The Relevance of Algorithms. In *Media Technologies: Essays on Communication, Materiality, and Society*, Tarleton Gillespie, Pablo J. Boczkowski, and Kirsten A. Foot (Eds.). The MIT Press, Cambridge, MA, USA, 167–194. <https://doi.org/10.7551/mitpress/9780262525374.003.0009>
- [19] Carlos A. Gomez-Urbe and Neil Hunt. 2015. The Netflix Recommender System: Algorithms, Business Value, and Innovation. *ACM Transactions on Management Information Systems (TMIS)* 6, 4, Article 13 (2015), 19 pages. <https://doi.org/10.1145/2843948>
- [20] Google. 2018. Natural Language API Basics. <https://cloud.google.com/natural-language/docs/basics#sentiment-analysis-values>, last visited: 18/10/03.
- [21] E. Guzman and W. Maalej. 2014. How Do Users Like This Feature? A Fine Grained Sentiment Analysis of App Reviews. In *2014 IEEE 22nd International Requirements Engineering Conference (RE)*. IEEE, New York, NY, USA, 153–162. <https://doi.org/10.1109/RE.2014.6912257>
- [22] Gregory D. Hager, Randal Bryant, Eric Horvitz, Maja Mataric, and Vasant Honavar. 2017. Advances in Artificial Intelligence Require Progress Across all of Computer Science. <http://cra.org/ccc/resources/ccc-led-whitepapers/>
- [23] Kevin Hamilton, Karrie Karahalios, Christian Sandvig, and Motahhare Eslami. 2014. A Path to Understanding the Effects of Algorithm Awareness. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems (CHI EA '14)*. ACM, New York, NY, USA, 631–642. <https://doi.org/10.1145/2559206.2578883>
- [24] Julian Ibarz. 2017. Updating Google Maps with Deep Learning and Street View. <https://ai.googleblog.com/2017/05/updates-google-maps-with-deep-learning.html>, last visited: 18/10/03.
- [25] Anthony Jameson and John Riedl. 2011. Introduction to the Transactions on Interactive Intelligent Systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 1, 1, Article 1 (2011), 6 pages. <https://doi.org/10.1145/2030365.2030366>
- [26] Shagun Jhaver, Yoni Karpfen, and Judd Antin. 2018. Algorithmic Anxiety and Coping Strategies of Airbnb Hosts. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 421, 12 pages. <https://doi.org/10.1145/3173574.3173995>
- [27] Judy Kay and Bob Kummerfeld. 2013. Creating Personalized Systems That People Can Scrutinize and Control: Drivers, Principles and Experience. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 2, 4, Article 24 (2013), 42 pages. <https://doi.org/10.1145/2395123.2395129>
- [28] René F. Kizilcec. 2016. How Much Information?: Effects of Transparency on Trust in an Algorithmic Interface. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 2390–2395. <https://doi.org/10.1145/2858036.2858402>
- [29] Todd Kulesza, Margaret Burnett, Weng-Keen Wong, and Simone Stumpf. 2015. Principles of Explanatory Debugging to Personalize Interactive Machine Learning. In *Proceedings of the 2015 International Conference on Intelligent User Interfaces (IUI '15)*. ACM, New York, NY, USA, 126–137. <https://doi.org/10.1145/2678025.2701399>
- [30] Todd Kulesza, Simone Stumpf, Margaret Burnett, Sherry Yang, Irwin Kwan, and Weng-Keen Wong. 2013. Too Much, Too Little, or Just Right? Ways Explanations Impact End Users’ Mental Models. In *2013 IEEE Symposium on Visual Languages and Human Centric Computing*. IEEE, New York, NY, USA, 3–10. <https://doi.org/10.1109/VLHCC.2013.6645235>
- [31] Brian Y. Lim and Anind K. Dey. 2011. Design of an Intelligible Mobile Context-aware Application. In *Proceedings of the 2011 International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '11)*. ACM, New York, NY, USA, 157–166. <https://doi.org/10.1145/2037373.2037399>
- [32] Walid Maalej, Zijad Kurtanović, Hadeer Nabil, and Christoph Stanik. 2016. On the Automatic Classification of App Reviews. *Requirements Engineering* 21, 3 (2016), 311–331. <https://doi.org/10.1007/s00766-016-0251-9>
- [33] Andrew Kachites McCallum. 2002. MALLET: A Machine Learning for Language Toolkit. (2002). <http://mallet.cs.umass.edu>
- [34] Sean A. Munson and Paul Resnick. 2010. Presenting Diverse Political Opinions: How and How Much. In *Proceedings of the 2010 CHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1457–1466. <https://doi.org/10.1145/1753326.1753543>
- [35] Pearl Pu and Li Chen. 2006. Trust Building with Explanation Interfaces. In *Proceedings of the 2006 International Conference on Intelligent User Interfaces (IUI '06)*. ACM, New York, NY, USA, 93–100. <https://doi.org/10.1145/1111449.1111475>
- [36] Emilee Rader, Kelley Cotter, and Janghee Cho. 2018. Explanations as Mechanisms for Supporting Algorithmic Transparency. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, 1–13. <https://doi.org/10.1145/3173574.3173677>
- [37] Emilee Rader and Rebecca Gray. 2015. Understanding User Beliefs About Algorithmic Curation in the Facebook News Feed. In *Proceedings of the 2015 CHI Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 173–182. <https://doi.org/10.1145/2702123.2702174>
- [38] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. “Why Should I Trust You?” Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '16)*. ACM, New York, NY, USA, 1135–1144. <https://doi.org/10.1145/2939672.2939778> arXiv:1602.04938
- [39] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2015. Recommender Systems: Introduction and Challenges. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, and Bracha Shapira (Eds.). Springer, Boston, MA, USA, 1–34. https://doi.org/10.1007/978-1-4899-7637-6_1
- [40] Advait Sarkar. 2015. Confidence, Command, Complexity: Metamodels for Structured Interaction with Machine Intelligence. In *Proceedings of the 26th Annual Workshop of the Psychology of Programming Interest Group*. Psychology of Programming Interest Group, Bournemouth, UK, 23–36.
- [41] Michele Willson. 2017. Algorithms (and the) Everyday. *Information, Communication & Society* 20, 1 (2017), 137–150. <https://doi.org/10.1080/1369118X.2016.1200645>