Kirsi Hantula & Oskar Korkman

Embracing user unpredictability

How to get to the next level in building recommendation systems

$$\prod_{n=1}^{\infty} \frac{1}{1-p_n^{-z}}$$

ALICE LABS × SOLITA

Contents

$\int x 2dx =$

Laura Lauch / Unsplash

“I think that currently algorithms mostly work like this: ‘You like this, take more’. The romance continues for a while but then it withers.”

Juho, 31, Helsinki, Finland

Introduction

Over the past decade or so, people have become increasingly used to interacting with AI systems which give recommendations on what to read, watch, listen or buy next, whom to date, or how to exercise. As these systems have gotten embedded into the fabric of our daily lives, people have also accumulated intricate tacit knowledge about the systems' strengths and weaknesses.

During this process, we have gradually learnt that recommendation algorithms are very useful for many needs but tend to fail in allowing users to do new things with their content. People do not always want to be their normal selves, and like sometimes to be 'unpredictable'. Current recommendation systems struggle to support this need, and people have begun to figure out ways to game or by-pass them, in a quest for more relevance and value for the content these systems act as gatekeepers to.

At the same time, developers of recommendation systems have not been blind to the ways how people try to tinker with their systems. They know that today's recommendation systems sometimes err and cannot easily adjust to changes in users' preferences. Quite often, however, developers seem to trust that these problems are temporary – or, at the very least, will markedly diminish in the long run, once more sophisticated algorithmic systems are designed.

According to this thinking, more data and more powerful algorithms will bring superior predictions, and user problems with recommendation systems will diminish. Technical mastery will make future recommendation systems respond accurately to users' varying needs and aspirations.

A new relationship with the user is needed

In this report, we present another view. We suggest that developing technical mastery and more sophisticated data capabilities will not be enough. Rather, the relationship between the user and recommendation systems will need to change, too: the offering of the next-generation recommendation systems should vary to mirror the changes in the user expectations, and users ought to be allowed to alter their relationshp with the systems.

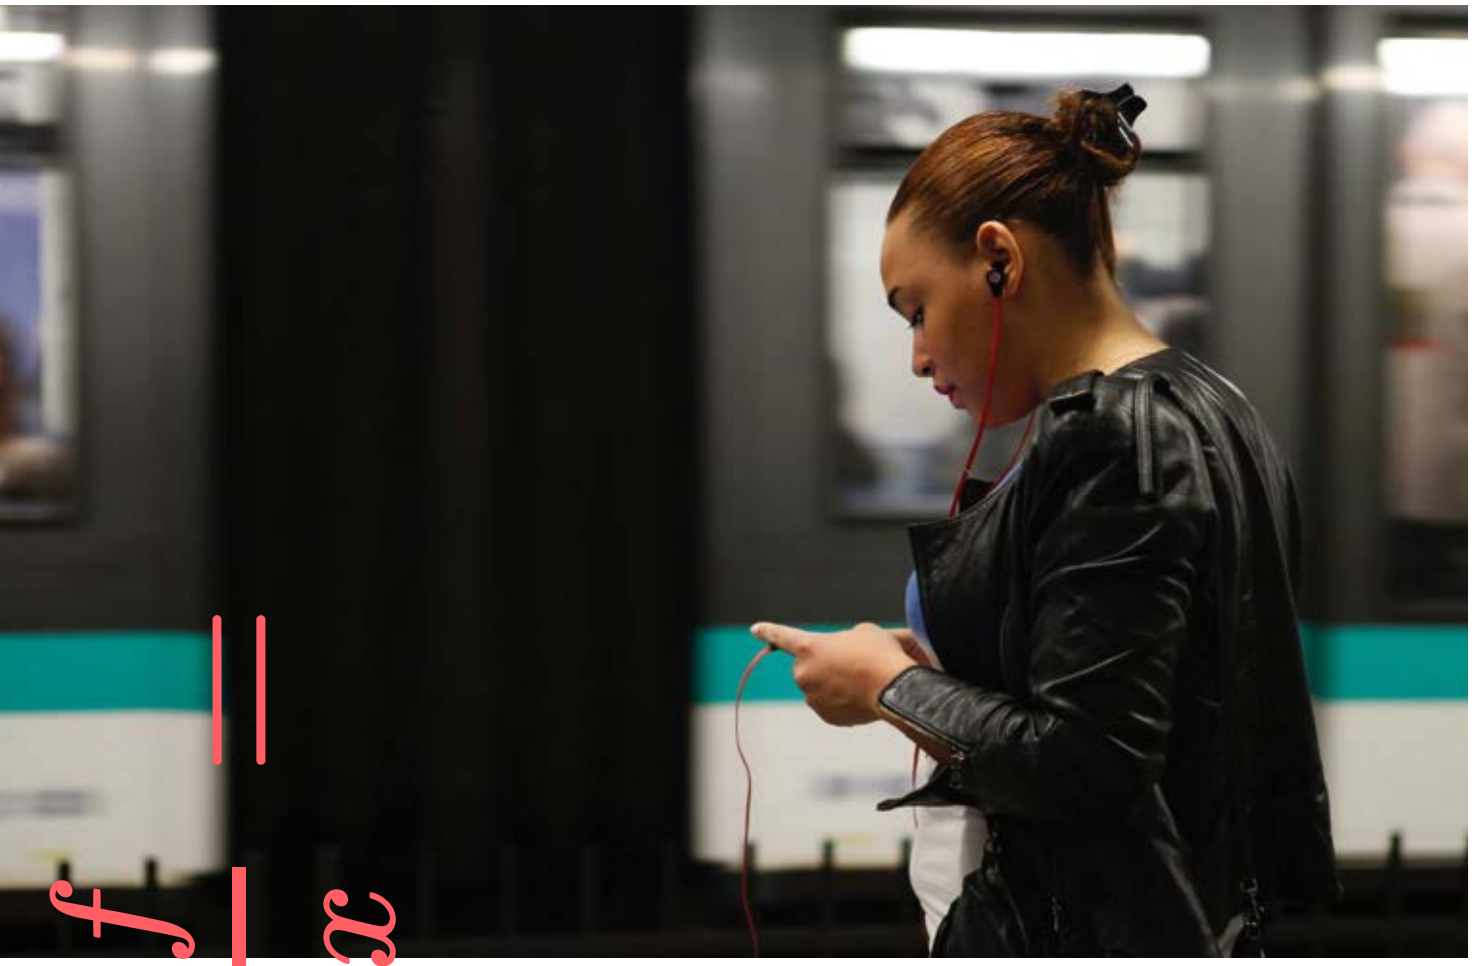

Eran Menashri / Unsplash

On some occasions, this might mean letting people start their relationship with a recommendation system temporarily anew, with no data or data history. On other occasions, that the user can adopt a new persona which creates different traces and consumption vectors for the recommendation system.

For the relationship to alter these ways, interfaces to recommendation systems will need to provide new ways of interacting with the technology. Users will no longer be as passive as they are now.

This is not to say that developers of recommendation systems will not be able to do a lot more with more advanced machine learning systems and data analytics. But this report shows that despite these advancements, the problem with the unpredictable and diverse user needs will remain, unless providers of recommendation systems change their thinking in the ways just suggested.

Interactive recommenders cooperate with the user

That is why our report proposes another approach. To get to the next level in the development of recommendation systems, system providers should start building systems that will allow users to be more in charge of

their recommendations. This report calls such systems *interactive recommenders*. Unlike the present-day systems, interactive recommendation systems will sometimes deliberately rely on user assistance to ease the tensions that are currently formed between users and recommenders in the situations where people's expectations towards their recommendations divert from the normal. And to let users participate, these systems will provide versatile interaction tools which can be used to alter the user's personal recommendations or nudge the system in the right direction.

This way, people will have a possibility to adopt a more active role in the generation of their recommendations when they see a need for it and let recommendation systems do their job on their own at other times. From the user's point of view, the ability to move between a more active and passive role will make these systems seem more flexible.

Derick Anies / Unsplash

No doubt, most often most users will still be happy to have recommenders do their job unassisted. When peoples' expectations remain predictable, there is no reason to interfere. But when tensions between them and a recommendation system do arise, people will now have effective interaction tools at their disposal which they can access to push their recommendations in the right direction.

With this help, recommenders will become more capable of accommodating more unpredictable user needs. They will be able to assist users also when they are feeling adventurous, silly, curious – or simply, in need of a change ■

A few words about the report and the research

This report is based on two years of extensive research by *Alice Labs*, a Finnish consultancy agency specialized in strategic consumer understanding. During the research, done in collaboration with the *Centre for Consumer Society Research, University of Helsinki*, 60 users of recommendation systems from five different countries have been interviewed.

The industry partner *Solita* has offered its expertise for developing the insights during the writing of this report and published the text together with Alice Labs. The main case partners – the music streaming service *Spotify* and the Finnish public service media company *Yle* – have offered invaluable help for finalizing the report findings by commenting on the report texts.

During the research project, each case partner has received a report about the findings of their own case study. For contractual reasons, the present report will not cover the company-specific findings of these studies. Instead, it summarizes other learnings from the case studies and the rest of the research conducted during the project.

The report is divided into two parts. In the first much longer section, called ***Tensions***, we will present four key tensions identified through our research which currently exist between users and recommendation systems. Each of the tensions describes a different situation where a discrepancy between the user and their recommendations is typically formed, due to recommenders' tendency to over-emphasize predictability and consistency.

Throughout the first section we will adopt the user's point of view. We will portray how the tensions are born by providing descriptions about our respondents' everyday interactions with recommendation systems, most of which are derived from the two case studies conducted with Spotify and Yle. Based on this understanding, early design ideas on how each of the identified tensions could be tackled will be offered.

Bryan Pulgar / Unsplash

The four tensions presented in the report are the following:

Chapter 1, titled as ***The downside of stability and continuity,*** outlines the tension which is created by recommenders' incapability to adjust to sudden changes in the user preferences. The chapter details typical everyday situations where this incapability generates a more enduring mismatch between the user's actual preferences and their personal recommendations.

Chapter 2, titled as ***The need for broadening recommendations***, describes the tension that is often formed between recommenders and users who have a long-lasting area of interest. The chapter depicts this process during which these users gradually acquire a deeper understanding of their interest. Sometimes this process eventually leads to a point where the recommendation system can no longer accommodate the user's thirst to learn new things about their interest. To overcome the problem, recommendation systems ought to assist users in finding adjacent or overlapping topics of interest, thus helping people to protect their interest from turning stale.

Chapter 3, titled as ***Limits of conventional categories***, details the tension that exists between recommenders and users, due to the way recommendation systems categorize their content according to a genre, a keyword, or an artist, for example. The main point in the text is that users, on the other hand, do not always want to operate on this level of granularity. Sometimes, they want to embark on a personal inquiry instead to find out about a fascinating detail or to study the larger context of a topic. When they do this, they usually leave behind the normal content categories which recommendation systems work with.

Chapter 4, titled as ***No assistance outside of my bubble***, depicts the tension that is formed when users want to find relatable, yet unfamiliar content outside of their normal preferences. In the chapter we refer to this aspiration as a wish to jump out of one's "bubble". The "bubble" here means the user's normal patterns of preferences which may be partly sustained by personal recommendations. The chapter concludes that

merely increasing the diversity of provided recommendations will not help users to jump beyond their bubbles. Instead, alternative ways to approach the problem are needed, and insights from users' real behaviours during these occasions will help.

The report's second part, called ***Suggestions***, concludes our findings by offering general guidelines that may help to get started with designing interactive recommenders.

The two-year research project by Alice Labs and the Centre for Consumer Society Research has been funded by the Foundation for Economic Education. We want to express our sincere thanks to the foundation for their support. The authors of this report also want to thank warmly **Anna Metsäranta**, **Antti Rannisto**, and **Jani Suvanto** from Solita for their valuable contribution in finalizing and publishing this report.

Interested to know more about the research project? Check the back of the report for a longer explanation about what we have done and what were our main findings during the first year of the research.

Focus of the study	Studied systems	First-hand data	Case studies in 2021
The use of AI systems at home /outside of work contexts	Recommendation systems in music streaming Recommendation systems in video streaming Recommendation systems in mobile and online news	60 in-depth user interviews with users from USA, China, the Netherlands, India, and Finland during 2020–2021. Interviews were conducted mostly online.	**Yle:** User modes in news consumption and the effects of these modes on peoples' expectations towards news recommendations **Spotify:** Users' relationship with recommendation systems

In addition to the data specified above, the research team has also studied other AI consumer systems than recommendation systems during the project and conducted one more case study with *Kela*, the *Social Insurance Institution in Finland*, amongst their chatbot users. As this report focuses on recommendation systems, this data is not shown in the table.

Tensions

In this section, we will present four persistent tensions that currently exist between users and recommendation systems, due to the systems' incapability to accommodate user unpredictability. More tensions likely exist, so the ones described here should be seen as examples of the problems that people commonly encounter with the present-day recommenders.

Each tension is described from the user perspective.

1. The downside of stability and continuity

As people's daily life circumstances and states of mind change, so do their content preferences – and quite often, personal recommendations cannot accommodate this change. Sometimes, however, this mismatch between the person's preferences and their personal recommendations endures longer, which pushes the user to temporarily abandon their recommendations.

"So, I don't ever watch the Dream Weddings** [a Finnish reality tv show] **at home. Somehow, we** [the user and his friends who have developed a habit of watching reality TV together] **need to be together, and then we can watch an episode from the Dream Weddings that could be four weeks old."

Teemu, 31, Espoo, Finland

As users of recommendation systems, people sometimes consume content that they do not normally engage with, simply because their mood or situation temporarily changes. As long as this consumption is a fleeting disruption to the normal preferences, people often don't mind even if the recommender cannot accommodate these changes. During these occasions, most are happy that their personal recommendations remain stable and coherent, even if the person temporarily strays away from their normal preferences.

In these situations, recommendation systems' reliance on aggregated user data is a blessing. It prevents personal recommendations from swinging from one end to another and retains the sense of continuity. For users, the experience is almost like sitting in a train which is safely moving along the main track, without ever turning to the smaller and less used side tracks which could direct the train off the course.

But repeatedly, there are also everyday situations where aggregated user data becomes a burden. During these occasions, people would like to stop the train or make it turn to another direction, as the same qualities – the consistency and stability of personal recommendations – which have guaranteed a great user experience in other situations, now constitute a problem. This happens because the mismatch between the user's altered preferences and their personal recommendations is bigger and more enduring.

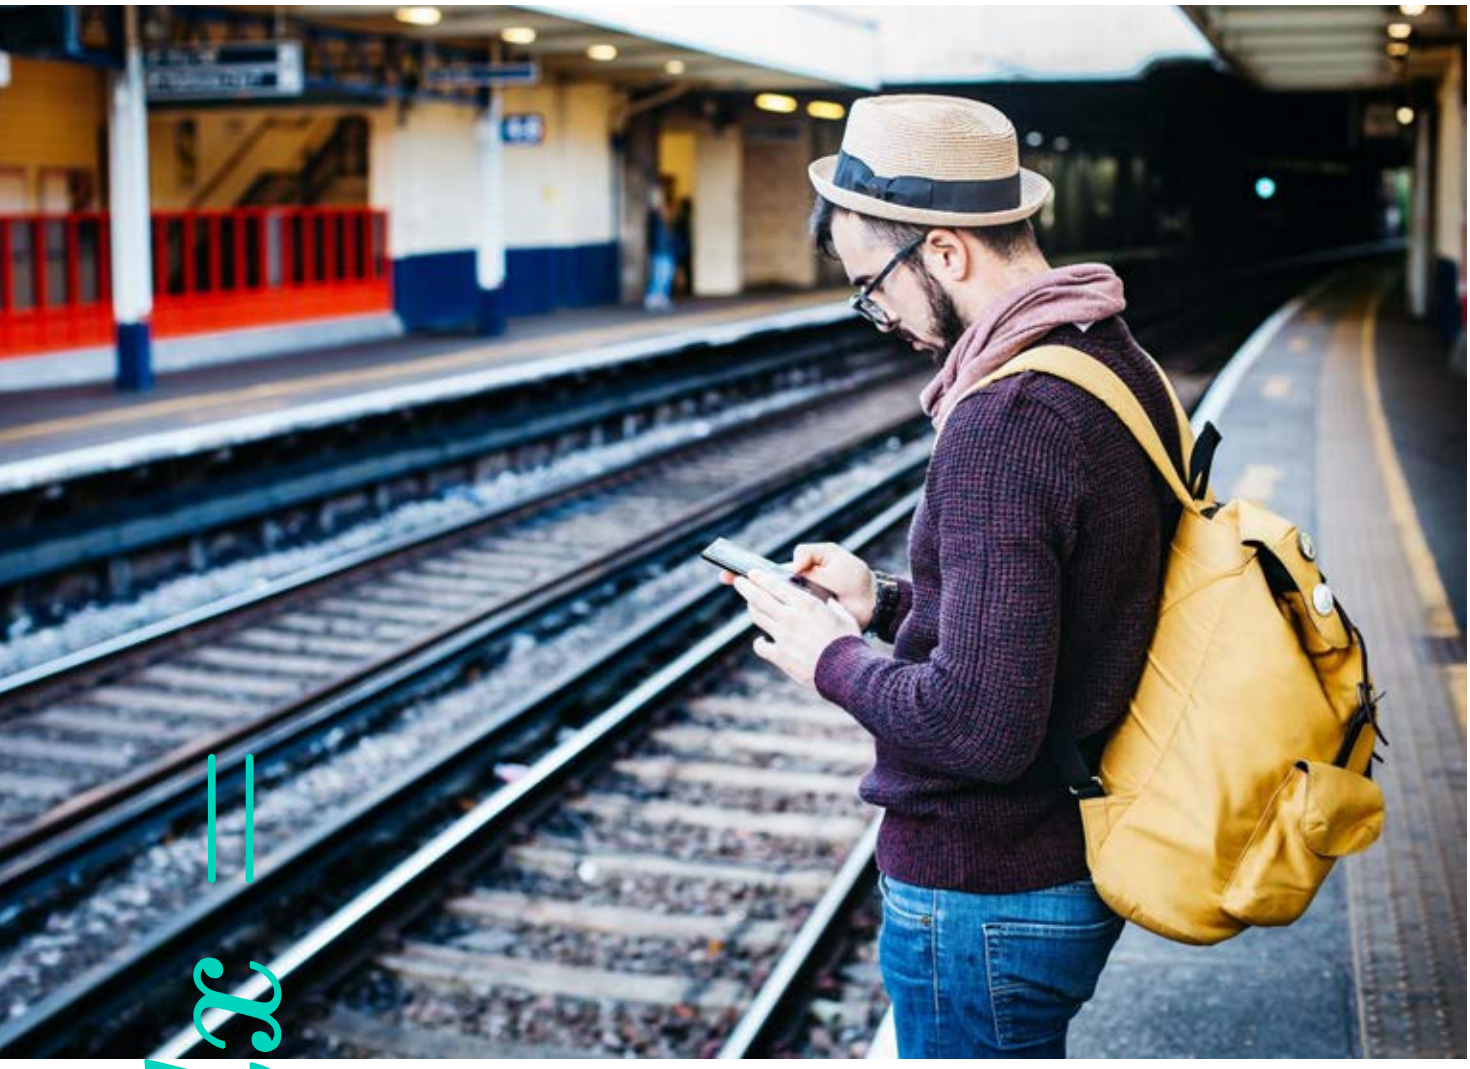

Clem Onojeghuo / Unsplash

$$x^2\,dx =$$

Imagine, for example, that a person has been traveling abroad on holiday, listening to a travel-related playlist for two weeks "non-stop", as one of our respondents did. Once they get back, they do not want to receive recommendations based on these music choices anymore. But for a recommender, which relies on historical aggregate data, making such a quick turn-around is impossible. In this situation, the user feels that their recommendations keep dragging in the past for too long – a past which now seems too distant from their present.

Or imagine that a person has a steady area of interest which they follow more or less continuously, but that they also have a few more occasional or lesser interests. Here, a mismatch between the user's preferences and personal recommendations is sometimes created because the user's most frequent content consumption dominates their personal recommendations, also during the occasions when the person wants to engage with the more sporadic interests.

To return to the earlier analogy about a passenger train, in a situation like this, the user would like the train to temporarily switch from the main track to a sidetrack and then back to the main track, which, of course, it will not do. Instead, the train continues its steady course along the main track, as in any other day.

> *"[Be]**cause...I've listened to so much pop and rap, that stuff kind of filters out any chance of seeing any rock or indie rock in my recommended** [songs]."*
>
> Daniel, 23, Bonaparte, USA

Melanie Pongratz / Unsplash

Similar mismatches between people's actual preferences and personal recommendations happen for most users every now and then, but currently there isn't much that they can do about them. People have no way of indicating to the recommender when and how a mismatch occurs, and most recommendation systems do not provide any tools that could be deployed to nudge the system in the right direction, either.

All that users can do is to temporarily turn away from their personal recommendations and find other means to get the content that they really want.

> *"**With the functions** [which the user presently has at their disposal for changing their recommendations] **that I know, there aren't much I can do.**"*
>
> Cheung, 33, Shanghai, China

Some common mismatches between people's preferences and their personal recommendations which create tensions:

- The user occasionally wishes to take a complete break from a topic that they are normally highly engaged with. This creates a more enduring mismatch between their preferences and their recommendations.
- Sometimes people want to alternate between their core and fringe interests.
- Occasionally in these situations, people would wish that the recommendation system could acknowledge their outlier consumption and other times that it would ignore it.

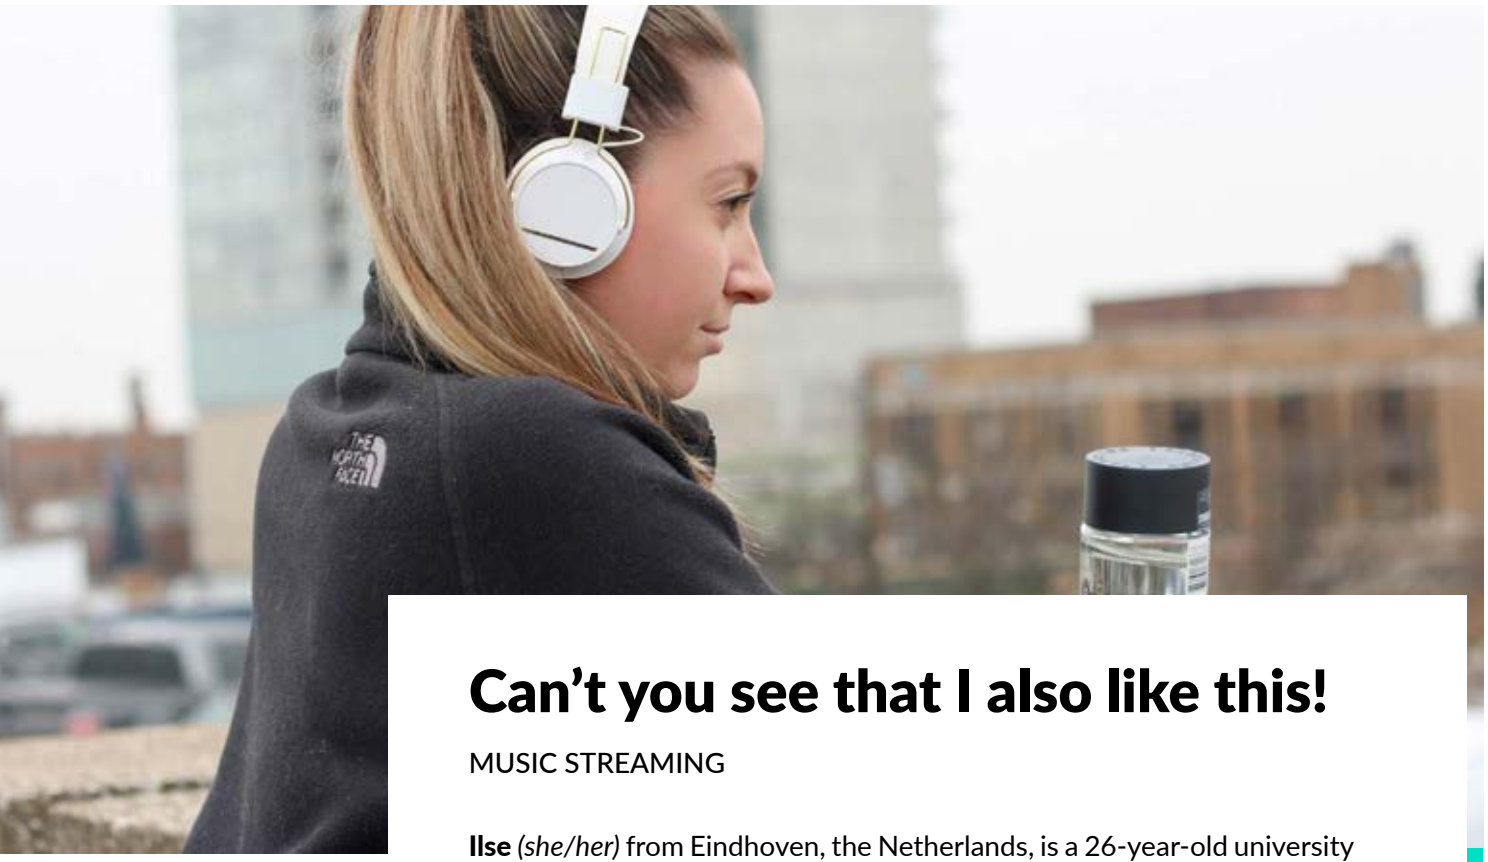

Element5 Digital / Unsplash

Can’t you see that I also like this!

MUSIC STREAMING

CASE

Ilse *(she/her)* from Eindhoven, the Netherlands, is a 26-year-old university student who listens to a lot of music and enjoys new musical discoveries. So, nearly every week she takes time to listen to the latest personalized recommendations for new music which her streaming service delivers.

By and large, Ilse is quite happy with these recommendations. She thinks that most often they reflect quite accurately what type of music she has been listening to and help her to discover new music that is close to it. The problem is, however, that Ilse does not always want to discover music which is similar to the music that she has been listening to the week before – maybe not even yesterday. Or, at least, not just that music.

Ilse gives an example. Though she primarily listens to other types of music, she has also developed an interest in Russian music. In the personal music library of her music streaming service, Ilse has saved albums by one or two Russian bands. The bands sound a bit like the *Cure*, a popular rock band from Britain, she explains, except that they sing in Russian. Every now and then, Ilse goes to the bands’ albums in the music streaming service and puts their music on, but there are also periods when she does not listen to their music at all.

And here comes the problem. Despite that Ilse only listens to the bands from time to time, she would still like to receive recommendations for similar-sounding Russian bands on a rather regular basis. This way, recommendations for new Russian music would be readily available any time when Ilse is in the mood for it. In other words, Ilse would hope that the recommendation system could largely ignore the pauses in her listening

history and keep delivering a few recommendations about new Russian music, also during the times when she has not been listening to her Russian favorites at all.

Now, Ilse is quite tech-savvy, so she does understand why the recommendation system of her music streaming service won't do this. She knows that the data from her more regular listening preferences overrides the more sporadic data that the system receives from the rarer occasions when she listens to the Russian bands. She is also wondering if the music recommendation algorithms in the service are more governed by user 'likes', and not so much by the data that is gathered when users listen to entire albums. If that is the case, she says, it would explain why the occasions when she listens to one of her favorite Russian bands don't seem to have much of an effect on her recommendations, either.

"[M]aybe it [the recommendation algorithm] just looks into the liked songs, whereas I think I am signaling [that I like the music] by listening actively to it [the album] a lot."

In the end, though, Ilse's algorithmic understanding does not help her that much. Despite her know-how, she has no way of overcoming the problem that she has with her personal recommendations. If only there was a way to make the recommendation system to understand her better!

From a top priority to give me a break

NEWS RECOMMENDATIONS

Many people have a long-term professional interest in certain news topics. This often has an effect on their expectations towards news recommendations: people are more likely to want to take a sudden break from the content that they follow for professional reasons, compared to the topics that they follow due to a personal interest.

Reeta, 37 *(she/her)* from Valkeakoski, Finland, is one of these people. In her job as a manager in the public service sector, she usually meets customers face to face. During the COVID-19 pandemic, however, this hasn't often been possible, and she has had to carefully follow the news about restrictions concerning the use of face masks or in-person meetings, for example.

To advise her colleagues and be able to react to the coming changes in time, Reeta has started her days by checking the latest pandemic news from her news app. In the app, users can set topical preferences to direct the app's news recommendations. Reeta has set "corona" as one of her preferences to find all COVID-19 related news more easily.

When Reeta has a holiday, her expectations towards news recommendations completely change, however. During the vacation, she wants to distance herself from the constant stream of COVID-19 related news and only maintain a general picture of the situation. To do that, she grabs her phone and adjusts the topical preferences in the news app, turning off all pandemic-related preferences. Reeta is pleased that the application does allow her to do that. Otherwise, she says, it would take a lot longer for the news recommender to notice how her news priorities have suddenly changed.

"I can change the topical preferences...but an automatic [AI-powered recommendation] system cannot do this, or it would take time before I get the preferences away from the system."

Reeta still uses the news app during her vacations also, but her primary source for COVID-19 related news during her time off is TV news – which she watches once a day to keep up with the general corona developments.

The problem of exceptional user preferences

Reeta's and Ilse's cases reveal the tension that is formed by a more enduring mismatch between personal recommendations and the person's actual preferences. The two examples show that recommendation systems' reliance on aggregated user data regularly produces dissonance with the alternating expectations that people have towards their personal recommendations in different everyday situations.

There are times when people want to take a complete break from their normal recommendations. There are occasions when they would like to reduce the volume of recommendations that relate to their prime interest in order to make more room for their more sporadic interests, as we saw

in Ilse's case. And there are also periods when people intensively consume content that they would not normally engage with – after which they may or may not wish that this consumption changes their personal recommendations.

In all these situations, users' expectations towards their personal recommendations are often exceptional, irregular, or maybe even one-of-the-kind. That is why historic aggregate data does not help recommenders to accommodate them. Quite the contrary, the aggregate data ends up feeding and sustaining the most discernible content consumption patterns, as it does in any situation. This makes personal recommendations skew on consistency and stability, even if it is not what users would wish to experience during these moments.

Implications for designing interactive recommenders

For this reason, the early design ideas that we offer below are all based on the idea that the tensions between users and recommendation systems described here can only be solved if users are allowed to participate in the generation of their recommendations.

During these occasions, the user is needed to modify the conclusions which the recommender has made based on aggregated user data because these conclusions strongly deviate from the person's actual preferences. And to correct the system, the user must have better ways to communicate with the system which will allow them to indicate explicitly enough what the present problems with their personal recommendations are.

To put it briefly, one-way communication in these situations must be replaced by channels of two-way communication.

"I feel like it's a one-way communication... It [a recommendation system] kind of analyzes your information and predicts what you might like... Other than that, you can't really communicate with your apps directly."

Justine, 30, Shanghai, China

Underneath, we will present some preliminary design ideas that might help users to better communicate with recommendation systems about when and how, they would like their recommendations to be altered in order to make personal recommendations better adjusted to the daily flux of their lives.

Adjusting the weight of different data streams

Providers of recommendation systems could develop tools which allow people to signal when they want to temporarily change how the recommendation system weighs different data streams. We think that these kinds of tools would make it easier for recommendation systems to cater for users' more sporadic or fringe interests.

Currently, personal recommendations are often heavily skewed towards certain topics or genres that users primarily consume content from. For people whose content consumption is heavily concentrated (meaning that they mostly follow one or two topics or genres), this sometimes becomes a problem. The data that a recommendation system obtains from their service use is centered on the core topics, and the sheer magnitude of this data tends to easily override any other data that has been obtained during the user's less frequent content engagements. This way, the bulk content consumption dominates the person's recommendations, thinning or squeezing out recommendations for the other topics or genres that they are also interested in.

Tools that allow people to easily modify how the recommender weighs their different data streams might help to correct such imbalances. For example, **Daniel**, our 23-year-old respondent from Bonaparte, California, whose recommendations are presently dominated by pop and rap music, would probably appreciate a tool that lets him increase the weight of indie music in his personal recommendations.

Easy-to-use tools for steering the system in real time

System providers should develop easy-to-use tools that people could resort to when the offered recommendations fail to acknowledge important changes in their preferences. These tools should give users a chance to nudge the recommendation system to making immediately discernible changes in their recommendations.

It should not be forgotten, however, that different types of services will probably require different types of tools. As an example, users of news

Nalai Nobel / Unsplash

applications might appreciate tools which will allow them to temporarily block certain recommended topics, as seen in Reeta's case. Or they might prefer to make quick changes to the volume of recommended news or the amount of news notifications. This hope was most vocally expressed by 56-year-old **Markus** from Finland who receives dozens of notifications daily from different news outlets and follows them partly for professional reasons also.

"It would be really cool if, for example, I could change [in the recommendation system]...that now, only notifications about the news that are above this level of criticality are allowed to come through... Like, now I will change the level of criticality to 5, I will only want the most important news."

Markus, 56, Turku, Finland

21 Possibilities to downplay or highlight specific periods in the user history

For a number of reasons, people sometimes have longer time periods when they consume unusual content. They may be sick for two weeks and prefer to doze off in front of the television while watching soap operas, even if they would normally never watch soaps. This is what happened to one of our respondents. Or they may have a road trip with a friend, during which they allow the friend to choose the music from their music streaming service for the entire journey, as another respondent did.

Afterwards, people usually wish to erase the effects of the unusual content consumption from their personal recommendations – but not always. Based on our research, we know that there are also occasions when people would like to amplify and prolong these effects, often because they perceive this as a chance to diversify their own recommendations.

To accommodate both of these (quite opposite) needs, providers of recommendation systems could develop tools which make it possible for users to specify how they would like the recommender to treat the data that has been gathered from the exceptional use period. One way of doing this could be to turn the user history into a cluster visualization which the user could study. With its help, people might, for example, indicate a spot in their user history where the recommendation system should start a new recommendation thread. Or maybe users could change their personal recommendations on their own by directly amplifying or "stifling" data in a specific cluster ■

2. The need for broadening recommendations

Many people have long-term interests which they gradually become more knowledgeable about. Over time, this learning process sometimes leads to a stage where personal recommendations can no longer help them to get deeper into the topic. To offer users a way of maintaining a sense of progress, recommendation systems should be capable of introducing adjacent or overlapping areas of interest.

Let's imagine that a user of a music streaming system becomes interested in alternative rock music, after hearing it at a friend's house who grew up with such music. At first, she only knows the bands that the friend has been playing, so she keeps listening to them on her own music streaming service. But soon enough, her knowledge about the genre starts increasing, as the streaming service recommends more of similar music in response to her changed listening habits. A virtuous circle reinforces itself, as the system feeds the user's appetite for more alt rock music and her increasing consumption pushes the recommender to deliver more recommendations along the same lines.

At some point, however – let's say after two years – this circle starts breaking apart. By now, the user has heard so much of alternative rock music that she already knows all the big names, and most of the smaller ones, too. She is starting to feel that her recommendations cannot really introduce any interesting new bands. At best, they can direct her attention to unheard songs by the bands that she already knows. And on a rare occasion, when the personal recommendations do point her to a completely new band, the experience is often lukewarm. Somehow, even the new band ends up sounding too familiar and too similar to the

ones that she has been listening to for the past couple of years. There is no sense of progress.

This is not what she wants, however. Instead, she hopes that her personal recommendations could still occasionally help her to experience a similar thrill as what she experienced when she first discovered alternative rock music. This does not seem to be possible anymore, though. She feels that her personal recommendations have gotten stuck at a certain level, incapable of progressing beyond that, though she would still like to keep moving on.

To solve this tension, the user starts to increasingly ignore the system's recommendations for alternative rock music and tries to find other ways to discover new bands which draw from the alternative rock tradition. This change in the user's relation to the personal recommendations is an indication that she has stepped into a new phase in her "user career", and her recommendations ought to reflect this change. Currently, recommendation systems cannot accommodate this transition very well.

In these situations, users have common expectations towards their recommendations which present-day recommenders fail to accommodate.

- Users would like to discover new things that will broaden their understanding on the topic of their interest.
- Users wish to find adjacent or overlapping areas of interest in order to retain a sense of progress.
- Users expect that the recommendation system will understand how their "career" as a user gradually advances and how these changes impact their expectations towards the personal recommendations.

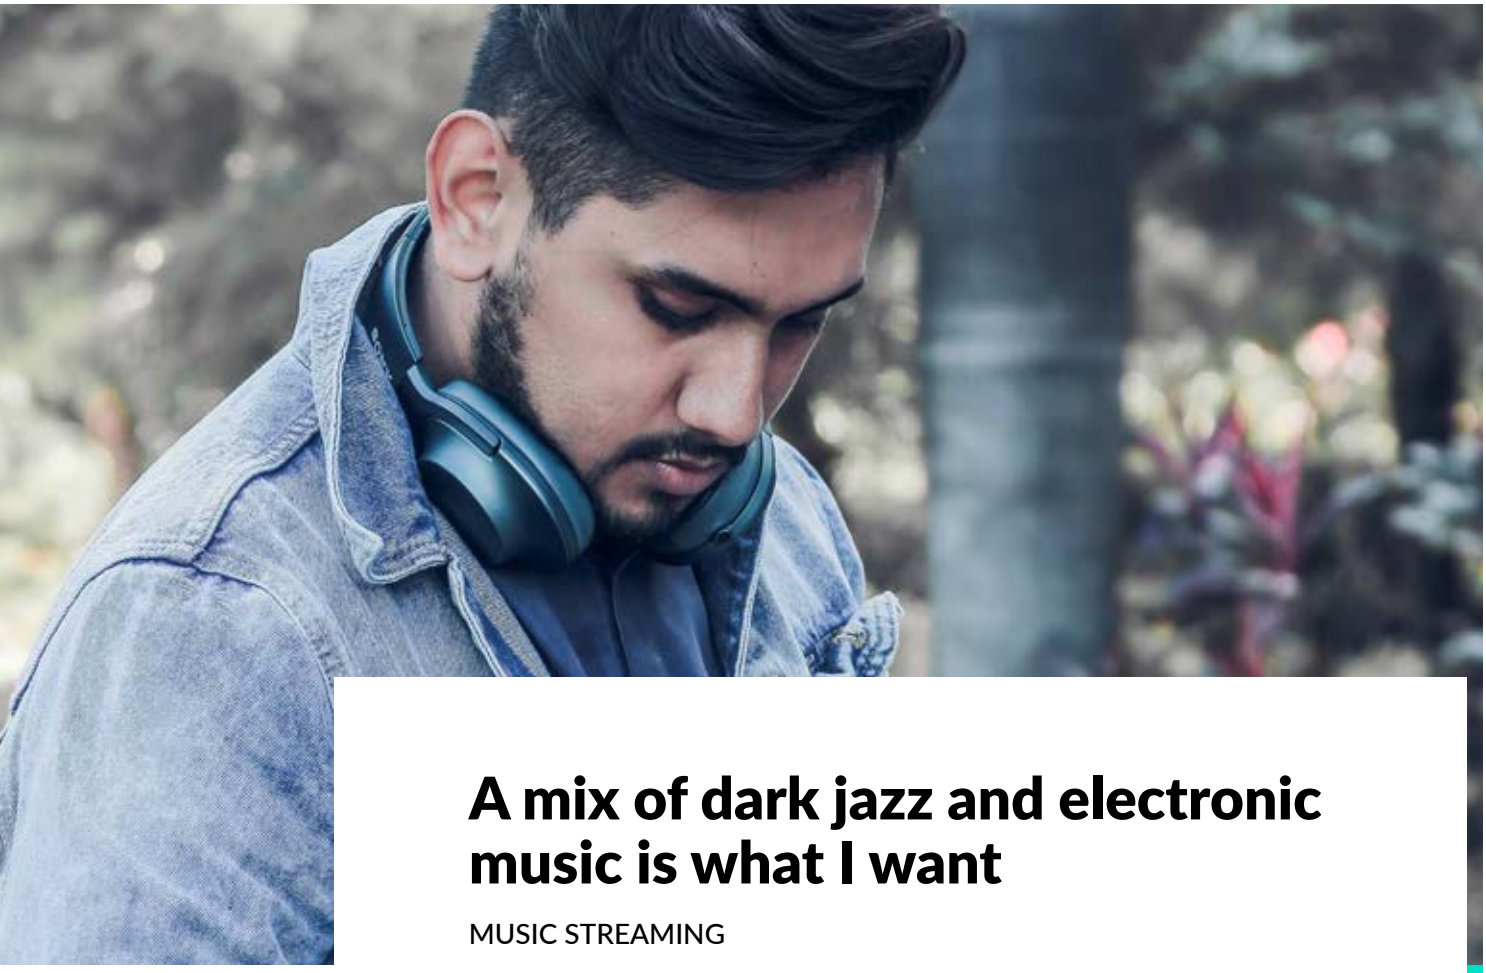

Angello Pro / Unsplash

A mix of dark jazz and electronic music is what I want

MUSIC STREAMING

CASE

When users' understanding of a topic of their interest is growing but recommendations cannot keep up with this change, people often feel that their recommendations have started dragging in the past. If the feeling persists longer, users typically start suspecting that they may have "emptied" the topic: maybe there simply isn't much more to be discovered within the subject, and that is why the recommendations seem repetitive?

These thoughts are familiar to **Rahul** *(he/him)*, a 28-year-old music fan from Indore, India. He has noticed that recommendations of his music streaming service have hit the wall in a couple of his favorite genres. These are niche genres where the amount of new music released is negligible, compared to more mainstream music, and the personal recommendations by his streaming service reflect this scarcity. For a couple of years now, the service's music recommendations have only managed to bring forward artists and bands that Rahul already knows.

"They [the artists that Rahul already listens to within these genres] keep on releasing new albums, but new artists are very difficult to come by."

Even though his personal recommendations have not managed to introduce any new artists from his favorite genres for a long time now, Rahul does not hold it against the system. He is aware that new artists and bands rarely enter his favorite genres.

To overcome the recommendation system's limitations and maintain the sense of personal evolution, Rahul would like to find new bands from other genres which are musically close enough to his favorite bands. He has noticed, however, that the recommender in his music streaming service cannot do this type of matching very well, so he has found another way to accomplish this goal.

When Rahul wants to find new music that is compatible with his long-lasting musical interests, he starts searching for playlists that are made by users who seem to like the same genres as he does. When he browses through these playlists, he usually knows most of the songs, as they are by bands which he already listens to. But sometimes these playlists also include songs which feel like true gems. These are songs by bands from distant genres that still share a similar musical language with Rahul's favorite bands. These discoveries have introduced Rahul with new bands from other genres, allowing him to retain a sense that his musical taste is still evolving.

"I don't usually listen to electronic music...but there are a couple of bands in electronic music whose music I actually prefer. But not because it's electronic, but because it's very close to [dark] jazz. So, it is something similar and it's the same vibe. But technically those bands are very different. One is in jazz, and the other is in electronic. So, they're poles apart", Rahul explains.

I need tools that help me to see wider

NEWS RECOMMENDATIONS

When people wish to expand their interests, a common challenge that they face as users of recommendation systems is their ignorance about possible directions where their interests could expand to. **Elli** *(she/her)*, a 17-year-old rail buff – and an avid follower or news – from Leppävirta, Finland is a living example of this. In order not to miss any interesting news about railways and rail traffic, she has set several topical preferences in her news app. There is "railway history", "railway accidents", "metro", and others.

While Elli is quite happy with the rail traffic-related news recommendations which she receives from her news app, she wishes that she could expand its range of recommendations further. This, however, has turned out to be

easier said than done. By now, Elli has already added several railway-related topical preferences in the news application and feels that she has already used the most apparent keywords.

Earlier, at the end of each news article, the news application offered suggestions for keywords that she might want to add as topical preferences. Then the feature was taken away. Ever since then, the effort of finding new keywords that could broaden the spectrum of Elli's rail traffic-related news has become more laborious.

"It [the earlier system which highlighted keywords that users could add as their topical preferences] could help you to understand that, without even realizing it, you were actually searching for particular topics. If, for example, you read several news articles within a day or a week, which all had some niche keyword...it helped you to notice what kinds of news you were really looking for."

"In a way, it [the earlier system] made [recommendations] wider and more focused at the same time."

Without the app's suggestions for topical preferences, Elli now finds the underlying content categorization in the service less transparent and perceives it more difficult to understand how the rail traffic-related content in the application is grouped and categorized.

Users need help for recognizing overlaps

When users have a long-lasting area of interest, many eventually feel the need to broaden it one way or another in order to protect their interest from turning stale. This process pushes the user into a new phase in their "career". At this stage, the user no longer only wants to receive recommendations from their long-term interest, they also wish that their recommendations could help them to expand their interest to connected areas. When the recommender cannot satisfactorily accommodate this expectation, a tension gets formed.

To make personal recommendations more attentive to people's exceptions at this stage of people's "user career", providers of recommendation systems should give users better possibilities to work with the systems.

Here, it is good to remember, though, that people who wish to get a wider range of recommendations, don't often know exactly how they would like their recommendations to be expanded, either. Once they have acquired a certain level of understanding, the most obvious answers to the question of where to look next will no longer satisfy them, but getting past these answers is not easy, as Rahul's and Elli's examples show.

Daria Nepriakhina / Unsplash

$$\int_a^b x^2 dx =$$

Thanks to a user-generated playlist, Rahul had learnt about musical similarities that exist between some bands that technically belong to very distant genres but finding the playlist which revealed the connections had probably taken him some time and required luck. Before the feature that suggested keywords underneath news articles was taken away, Elli had received similar help from the news application.

In both cases, our respondents needed to become aware of the existence of overlapping or adjacent areas of interest which could enrich their own long-lasting interests. Surely, there would be many more similarities which Rahul and Elli could uncover if there was a more effective and systematic way to do this! This way, they could feel that their interests are still evolving and that they are still able to discover new things within their area of interest.

Our respondents' experiences underline an important point that we want to make. In the world of abundant content, there are always overlapping interests to be explored and fascinating parallels to be discovered but finding them is often challenging. While interactive recommendation systems will not fully remove this pain either, they can considerably alleviate it if they become better at highlighting overlaps and similarities between users' interests across their entire content universe.

Implications for designing interactive recommenders

Below, we will suggest some design ideas on how recommenders might better help users to spot these similarities. Our suggestions are largely based on the idea of cooperation: in the vast content universe, neither the recommendation system, nor the user alone may be able to consistently keep finding the hidden gems that will successfully widen the range of user's recommendations. But if the recommender and the user work together, benefitting from each other's strengths, this becomes easier.

*"So, you know, as the tool [the recommendation algorithms] **shapes us, we shape the tool. The more we shape the tool, the more efficient, the better it gets in what it's doing, and then, in return,** [this] **shapes our behavior.**"*

Chris, 42, Millersville, USA

Longer-term use patterns reveal users in risk

Recommenders will not be able to assist users who need broader recommendations without the ability to distinguish them from the other users who are still in an earlier phase of their "career". To be able to do that, providers of recommendation systems should develop their capability to discern longer-term use patterns.

Based on our research, we understand that the challenge of locating the correct users within a large pool of users is not easy to tackle with the help of aggregated user behavior data. Though we know that the users whose primary topic(s) of interest have not changed for a long time and whose

content consumption is heavily focused on these interests, comprise a risk group, we also know that not all these people need broader recommendations. To put it differently, amongst these users, too, there will be many who have engaged with the same topic(s) in the service for a long time, but still don't find their recommendations repetitive or restrictive.

Andrew Leu / Unsplash

$b3-a3$

To make it easier to find the users who have hit the wall with their recommendations, we suggest that providers of recommendation systems might explore whether they are able to complement their current analysis methods by utilizing longer-term usage data more effectively. For example, if the service provider looks at the larger group of users whose interests in the service have remained unchanged for a long time, can they find a sub-group whose satisfaction with their recommendations has diminished unproportionally, in comparison with the rest of the group? Or can they see other longer-term changes in the use of the service which indicate that certain users may be in need of broader recommendations?

Suggesting overlapping or closest interests

Identifying longer-term use patterns from the user data will only help system providers to roughly gauge users who might need help in keeping their long-standing interest alive, and there will still be many who are mistakenly placed in this group.

For this reason, it will be important to develop sensitivity about the least intrusive ways to expand recommendations and find interaction points where users tend to be most open to new content. Services that offer personal recommendations often include features which users tend to access when they are feeling more explorative. They will be a good place to start.

Based on this type of service-specific understanding, system providers will be able to carefully experiment how to broaden recommendations for the users most likely in need of change. One rather obvious way to do this could be to suggest selected content from subject categories that are closest to, or overlapping with, the user's present interests. This is an explicit wish by many users also.

"For example, if I have chosen the Champions League and the Finnish Football League as my topic preferences, it is pretty likely that I might be interested in the English Premier League, Serie A, or some of these other top football leagues in the world, as well."

Tomi, 43, Turku, Finland, a user of a news application

In addition to broadening the user's recommendations by including recommendations from adjacent topics, providers of recommendation systems might also notice other less apparent opportunities which they could tap into, when they want to broaden the user's recommendations. Noticing them will require a deeper and service-specific understanding of users' aspirations, though..

"I wish that it** [the news recommender]**...was more apt to recognize what types of news I am interested in, rather than what kinds of topics... Like, if I am interested in human interest -type stories and like to read them,** [the system could understand that] **a human-interest story** [that the user is interested about] **can be about sports, history, urban culture, or anything."

Miikka, 38, Helsinki, Finland

Regardless of the way users' personal recommendations are broadened, the shared ambition of system providers should be to develop solutions that offer a gradual path away from the users' original interests. A visual representation of this could be a Venn diagram where the overlapping or parallel topics form circles that connect with each other, steadily leading further away.

Finding people with similar interests

Besides developing their capability of broadening users' recommendations, providers of recommenders could also study possibilities to develop new tools that will let the user cooperate with their system to expand the personal recommendations.

These tools could mine from the understanding of users who have similar interests to locate the more surprising and hidden relations between different topics. As Rahul's example shows, people who have matching interests are usually the best channels for finding adjacent areas of interest.

In practice, a solution like this could be achieved if the recommender was able to detect users with somewhat similar interests, determine which content these users are highly engaged with, and pick from this content pool the most interesting recommendations that it will show to the user. Here, an effective, real-time feedback loop between the user and the recommender would be essential for success. The feedback loop should allow the user to easily communicate to the system which recommendations seem relevant, and which do not. Ideally, the user might also be able to adjust the degree of overlap that exists between their present interests and the interests of the users which the recommender uses to produce their personal recommendations ■

3. Limits of conventional categories

Recommendation systems cluster different items based on the same, and often rather conventional attributes. They use genres, topial keywords, and names, for example. Users, on the other hand, do not always want to operate on this level of granularity. There are situations when they want to zoom in or out of these categories to study a small detail or to see the larger context.

Basic building blocks of any recommenders are different categories that the system uses to place each item of content somewhere within its universe and to find similarities and differences between items. Modern, advanced recommendation systems do this filtering by using a variety of different attributes and techniques, but they still largely rely on traditional categories which group content by a genre, a topic, or an artist, for example.

A lot of times, these traditional content categories correspond with how users, too, perceive content. They reflect the way people have learnt to define their interests and taste. Most of our interviewees, as well, could easily name their favorite bands or define their top film genres. They could talk about the specific categories of news that they followed in depth. While this was true, for them, genres or news categories were still just one way of perceiving content. When we poked the subject deeper, we noticed that behind the conventions of how to talk about films, music, books, or news, there was more variation in how our interviewees approached content.

This became particularly obvious when people talked about the occasions where they had more time at their hands to engage with content. During

such moments, people often became more inquisitive and playful in their engagements with content.

They enjoyed using the extra time that they had to find out about a specific detail that had caught their eye, either on- or off-platform.

"Even when** [I'm reading] **a book, I notice that I often use** [Google] **Maps to follow where the events take place if the locations change a lot...I remember when I was reading a historical novel which included a lot of references to art, it took a lot longer than usual to read the book because I had to check from Wikipedia how a specific statue looks like or find out where the Uffizzi Gallery is located...It was multitasking-reading, in a way."

Mika, 44, Kaarina, Finland, an avid user of a Finnish news application

Or they took joy in following an interesting trail of information – a trail that often had many twists and turns along the way. A New Jersey -based 39-year-old music lover (who in normal circumstances prefers to stick with R&B music) gives an example. She explains what happened when she heard about a recent break up of a long-lasting American electronic band and learned through her inquiries that the band's style had evolved from punk-ish lo-fi to something akin to Italo-disco.

"I was like, 'Who are the Chromatics? Why did they break up?' Because it was like, the BBC announced that they're breaking up...and I haven't really listened to them intentionally. So, I went through this** [play]**list [in her music streaming service] for a while, and then** [after hearing the band's more recent disco-sounding songs] **I was just like, 'I just wanna listen to disco'."

Latyah, 39, New Jersey, USA

The above examples highlight how people typically act when they are in an inquisitive mode and not pressed by time. They show how people's inquiries during these moments regularly take them across and beyond the genre-, topic- or artist-based categories which normally structure their experience of using recommendation systems. And they reveal how these inquiries repeatedly push people away from the level of granularity which the normal content categories offer.

Hitesh Choudhary / Unsplash

Let's say that a user of a news recommendation system becomes curious about a detail in a news story, for example. In this mode, the user does not want to read another article that belongs to the same keyword category as the article they have just read, even if the recommendations underneath the article offer these articles. Instead, they want to leave the first keyword category behind and move on to a more specific category that can help them to learn about the detail which caught their attention.

In a situation like this, the user is still eager to make sense of content, but not by moving within the conventional genre- or keyword-based

categories but rather by moving up, down or across them to find out about the aspect that has caught their attention.

In these situations, people often have certain aspirations which recommenders fail to accommodate:

- Users wish to drill into an intriguing detail.
- They want to understand how different aspects impact each other or form a whole together.
- They want to study a topic from multiple angles to establish a larger context or to understand different viewpoints.

As recommendation systems cannot properly accommodate these expectations, people's efforts are presently often cut short. Or, alternatively, their inquiries are largely conducted outside of the original services – often with the help of search engines or other handy sources of information, such as the online encyclopedia *Wikipedia*.

"I love Wikipedia...When it talks about a subject, it [the text about the subject] has links all the time. And by clicking on a link, you get deeper and deeper into the topic, and then you can return to the main story again. It would be wonderful if something similar existed in the news stream as well. For example, if...a news story included a building...you could expand the informational content through the link and read about its history."

Mika, 44, Kaarina, Finland

Guilherme Stecanella / Unsplash

“Should glyphosate even be used, that is my question!”

NEWS RECOMMENDATIONS

CASE

31-year-old **Juho** *(he/him)* from Helsinki, Finland browses different news outlets multiple times a day and considers news articles as an important way to make sense of life and the world. After talking with him, it becomes clear that this sense making happens in at least two ways. Browsing the news helps Juho to stay informed about current issues – simply to know what is going on.

But equally importantly, news articles also regularly push Juho to form his own perceptions about the topics that they cover. This process is much more active. Juho tries to explain this by describing how news articles “raise questions” in his mind. Often, these questions go beyond the scope of the original article, so the process of searching and finding answers to these questions becomes Juho’s way to form his thoughts about the issue, on his own terms.

“[A news story about] ***Japanese roses is a good example, and how their poisoning impacts nature. On the Metsähallitus*** *[a state-owned enterprise in Finland which produces environmental services] [web]* ***page, there is a comparison that they use 60 kilos of glyphosate per year for poisoning Japanese roses in the Archipelago Sea*** *[a sea area in Finland].* ***That’s nothing since farming uses 600 tons a year! So, should glyphosate be used at all, that is my question.”***

The questions that news articles lead Juho to ask can be all kinds. Sometimes, they relate to the wider context of the original news topic, and other times they are about a small detail. No matter what the question is, the process of an active personal inquiry that the question sparks is important for Juho. It helps him to better grasp the topic.

To explain how his personal inquiries sometimes change his thoughts, Juho draws an analogy to the act of reading a theater critique versus seeing a play. He describes how reading a theater critique will only give the reader a partial understanding of what the play is really about, whereas seeing a play will allow you to reflect on what you have seen and form your own thoughts. With this analogy Juho wants to point out how important the news-related inquiries often are for him as a way to develop a deeper understanding on current issues.

Presently, these inquiries usually take Juho away from the original news sites to other news outlets or social media, for example. According to him, this largely happens because news sites tend to only recommend their former articles about the same topic, but answers to the questions in his mind can rarely be found in them.

A personal learning tour about Trent Reznor

MUSIC STREAMING

35-year-old **Rob** *(he/him)* from Orlando, USA is a music aficionado who listens to music almost constantly and has a large appetite for music. On a normal day, he easily spends several hours listening to music.

There are times when music mostly plays in the background. It accompanies what Rob is doing, without him paying that much attention to it. But there are also times when he can concentrate more on the music.

During these occasions, Rob sometimes embarks on a rather similar inquiry as Juho does, except that Rob's inquiries have to do with music. And similar to what happens to Juho, also Rob's inquiries usually lead outside of the service, or alternatively back and forth. That is what happened recently, when Rob was listening to an American 90's rock band and his music streaming service suggested that he'd continue by listening to another band, *Nine Inch Nail*. Rob had listened to the band years ago, but not recently. So, in

a spur of a moment he clicked on the suggested song and arrived at the Nine Inch Nail's page on the music streaming service.

At this point, things took another turn. As Rob came to the band's page, his attention was caught by the profile picture. Why was there just one man in the photo? After glancing at the introductory text on the band's page, Rob learnt that there was indeed just one official member in the band, and he was called **Trent Reznor**. This information made Rob curious, and he wanted to learn if the man had done something else musically, as well. So, he left the service and searched for information on Trent Reznor's music on Wikipedia.

Through his inquiries, Rob learnt that Trent Reznor had also composed music for several video games and films. After learning about this, he wanted to hear what these soundtracks sounded like. He returned to the music streaming service to see whether he could find some of them in the service.

According to Rob, he regularly conducts similar inquiries about all kinds of music-related facts that spark his interest. Most often, these inquiries are done with the help of Wikipedia, not by his streaming service, but Rob wonders if the music streaming service could better alert him of the existence of interesting connections or fascinating musical detours that he might like to explore when he is in an inquisitive mood. While Rob would not want his streaming service to turn into a text-heavy online encyclopedia, he thinks that the service could find its own ways for pointing him to directions that he might like to probe.

"Maybe if I was listening to** [the rapper] **Jay-Z, and if there was something on the service that was like, 'Oh, if you like Jay-Z, you'll like this collaboration album that he did with** [the rock band] **Linkin Park,' or something like that. Whereas…that piece of information** [about the unusual cross-genre collaboration] **is something that I would** [now] **find out by going on Wikipedia and reading about Jay-Z."

Aid for cross-category inquiries

Users of recommendation systems often get curious about tiny details, when engaging with content. Or they become interested in the larger context which a particular topic is part of or wish to understand how

different things form a chain together. Juho's and Rob's cases are descriptions of such occasions.

These attention shifts regularly get users to abandon the normal content categories which are based on genres, artists, or other similar keywords, and instead opt for other ways of making sense of content. As current recommendation systems are largely incapable of accommodating these shifts, a tension is formed.

To better accommodate people's needs in these situations, recommendation systems ought to be capable of helping users to temporarily travel between and across the usual genre- or keyword-based categories. Below, we will sketch some ideas about how this might be done.

Two simple observations offer a starting point for our ideas. Firstly, moments when users want to engage in personal inquiries to educate themselves on something, typically take place when users can leisurely browse content. Secondly, people usually find such occasions pleasurable, and feel this way partly because they can freely decide where the inquiries will lead.

"I think it's more of my curiosity that sends me off to other places to look for more information."

Rob, 35, Orlando, USA

To conclude, users in these situations often enjoy putting some personal effort on discovering answers to the open questions in their mind. They would not want this joy to be deprived completely.

Implications for designing interactive recommenders

We think that the above observations provide guidance for the design of interactive recommendation systems at least in a couple of ways.

We know that users tend to conduct their personal inquiries when they are not pressed by time. This means that providers of recommendation systems will not have to offer possibilities to crisscross or abort the usual keyword- or -genre-based categories every time users interact with their system. It is enough to be able to offer them when users are browsing

Imad Alassiry / Unsplash

content leisurely. To better serve users during these occasions, providers of different recommendation systems should have an exact understanding of which features in the service people prefer using more leisurely and use this knowledge as the basis of their design efforts.

We have also established that part of the pleasure from users' personal inquiries often comes from the joy of hunting. People enjoy when they have succeeded in uncovering an intriguing detail or when they have found an interesting trail of information which they can follow. For us this tells that the developers of recommendation systems should be careful not to take away this thrill by creating features that are too enclosed. In other words, system developers should avoid pushing all users along the same routes, and instead make sure that people are offered enough freedom to decide what they will become curious about.

The two design ideas that we present below are grounded on these observations.

Providing openings for personal inquiries

Providers of recommendation systems could investigate possibilities to offer deliberate openings – or "snags of thread", as a user of a Finnish news application described them – that will take users across the usual genre- or keyword-based categories. Such possibilities should be predominantly offered in the features that users mainly engage with when they are browsing content more leisurely.

These experiments should acknowledge that people's inquisitive curiosity will probably receive slightly different forms in different services. Users of news recommendation systems, for example, often wish to dig into specific details in and around a story. Knowing this, providers of news recommendation systems could experiment with creating different kinds of openings that will accommodate these user aspirations. Imagine that one such opening could, for example, reveal the article that had spurred the reporter to make the story in the first place, as one of our respondents pictured.

41 *"[It would be]* ***sort of like 'inspired by'."***

Miikka, 38, Helsinki, Finland

Users of film or music recommendation systems, on the other hand, might appreciate "snags of thread" that help them to grasp the entire breadth of an artist's work, as we saw in Rob's case.

To keep these openings diverse and engaging enough, providers of recommendation systems could also consider offering recommendations more boldly across their entire offering and even between different recommendation systems, instead of just looking at the service's own content. As a user experience, receiving a recommendation like this would thus become something which is close to encountering a hyperlink that leads to other related content. These "hyper-recommmendations" could highlight interesting content across content categories or from different recommender systems. **Elli**, a 17-year-old user of a Finnish news application gives an example.

"If I was reading news about space, for example... it *[the recommender in her news application]* ***could suggest tv or talk radio shows*** *[by other services than the news from the same media company]* ***which relate to the topic."***

Elli, 17, Leppävirta, Finland

Developing the search function

The efforts of conducting personal inquiries might also be eased without making extensive changes to the recommendation system. We believe that in many recommendation systems this could be achieved by developing the search function.

For example, quite a few popular streaming services and news applications currently allow users to only conduct searches by using a title, a keyword, or the name of an artist, but making more elaborate searches by combining two different search terms is not possible. And even the services that do allow for simple combinations of search terms are not able to help users to understand how the search results fit into a wider

context or connect to other things. Let's say that our respondent Rob could have conducted a search on his streaming service by using '*Jay-Z*' and '*Linkin Park*' as his search terms, and had found out that the two had released a collaboration album together. What if Jay-Z would have had a period in his career during which he collaborated with other (rock) bands also? Wouldn't Rob have likely been eager to learn about that too, but the search function had not definitely been capable of delivering that information!

Here, the problem partly comes from how services present the search results. Currently, there is usually no way for the user to see things that closely connect to the topic that they have searched for. The results are technically flawless, but they do not offer additional "snags of thread" which the user could follow, should they want to do that.

We suggest that this problem might be partly tackled by presenting search results differently. For example, if the results were presented visually as part of a broader word or tag cloud, this might help users to see new connections between different things ■

"If you could, for example, take an actor in the middle and see that she has done this and that. Or, you could take two actors and see what they have done together... [The search results would visually form] almost like a neural network. I think it is a very interesting way to present things, because you see where things are most concentrated and least concentrated. Like, over here there's this, and this is connected to that."

Tiina, 55, Helsinki, Finland, a user of a video streaming service

4. No assistance outside of my bubble

People sometimes want to expose themselves to completely new ideas or viewpoints. They want to challenge their thinking or develop their taste – to grow. During these moments, users of recommendation systems typically try to search for unfamiliar content that they can still recognize and relate to on some level. The challenge is how to find this content, without having to shift through a lot of unengaging, cryptic, or even abhorrent, material.

Regular musical explorations by **Gerardo**, 32, from Portland, USA, may provide an example. Being what he calls "socially progressive", he considers it important to listen to music from many places and sources. He describes music as a way to educate himself about issues that are removed from his daily experiences but still important for other people.

"We live automatically in a social bubble of people around us and the setting that we are in. And so, I think the best that we could really do, it'd be to just change the bubble for a little bit and then bring the elements of that back into our own [bubble]."

To expose himself to the elements outside of his bubble, Gerardo occasionally sets off to learn about music from other cultures. Recently, he has been exploring Kazakhstan music, for example.

When he embarks on his cultural explorations Gerardo often encounters the same problem, however. On one hand, he cannot rely on his personal recommendations because there is nothing in his listening history that could make the recommender suggest the music he is looking for. On the other hand, he also knows that to discover unfamiliar music that is approachable enough, he will need help, as he has no prior knowledge

of the music that he wants to learn about, and he usually cannot understand the language, either.

To get past this problem, Gerardo usually moves his explorations from his music streaming service to another service that has a wealth of largely user-generated content. In this service, he is unrestrained by his personal recommendations and can search for user-generated playlists which, according to him, will provide the best and shortest way to learn how people from a different culture think and feel. In a way, user-generated playlists thus work as Gerardo's compasses in an unfamiliar musical terrain, providing access points to a music that he has not prior experience of.

In these situations, people have common content aspirations which present-day recommenders cannot help to fulfill:

- Users wish to discover new content outside of their normal bubble which will challenge their thinking and introduce new ideas.
- Users are curious to learn what kinds of things and ideas attract people who occupy a social bubble that is different from theirs.

The problems that Gerardo encounters during his out-of-the-bubble musical explorations highlight the typical challenges that users of recommenders often face, when they want to expose themselves to new ideas and viewpoints. To put it briefly, users' problem in these situations is the following: people require assistance to find relatable content within the sea of unknown content, but their personal recommendations do not help.

To solve this problem, most users end up making similar conclusions as Gerardo. They abort their personal recommendations and try to find the content that they are looking for by using other means.

This is not to say that users would be against receiving personal recommendations during their out-of the-bubble explorations. On the contrary, many people remain quite positive about the recommendations and would gladly resort to them if they could be of more help. After all, they know that finding relatable content outside of their own bubble usually requires both time and patience. They also realize that their explorations outside of their bubble are repeatedly hindered by their own lack of understanding. To rephrase Gerardo, if the user does not even know what they are looking for, how could they possibly find it?

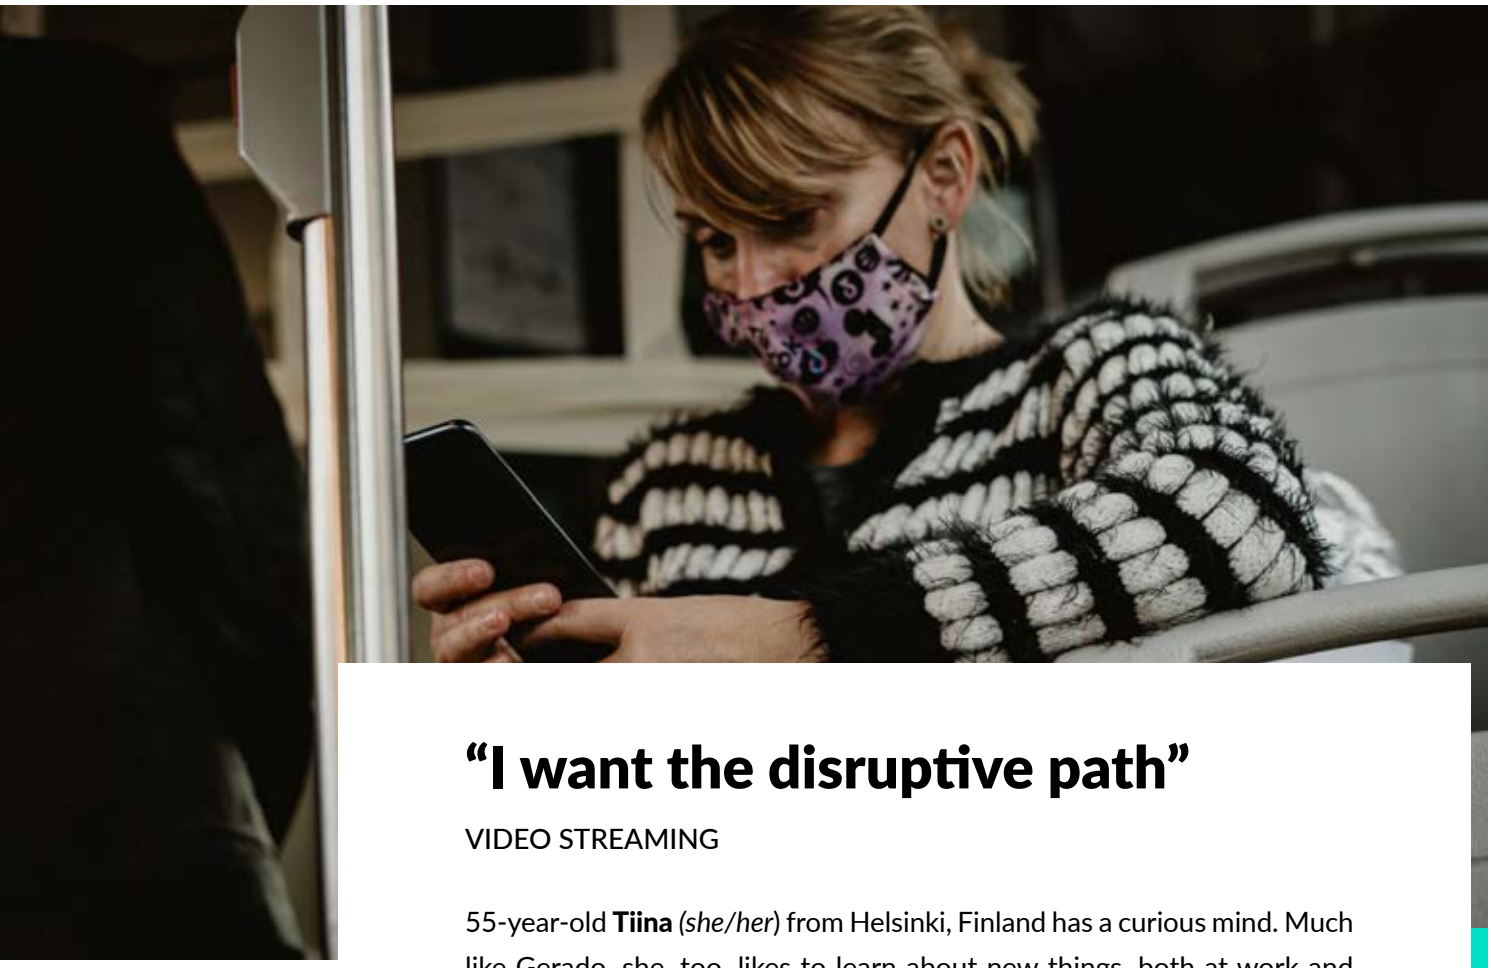

Gabriella Clare Marino / Unsplash

CASE

“I want the disruptive path”

VIDEO STREAMING

55-year-old **Tiina** *(she/her)* from Helsinki, Finland has a curious mind. Much like Gerado, she, too, likes to learn about new things, both at work and in her leisure time. This tendency also affects the way she uses her video streaming service. Of course, there are many days when she is quite happy to receive recommendations about TV shows or films that belong to her favorite genres. But there are also recurring days when she wants to step out of her bubble to learn about completely new ideas and perspectives.

During these occasions, Tiina’s personal recommendations become a problem, and she finds it very difficult to escape them.

“I don’t want it** [the recommendation system] **to recommend things that I watch anyway.”

To tackle the challenge, Tiina has created two separate user profiles in her video streaming service, one of which she only uses when she wants to exit her normal bubble. But this has not been enough for her, either. Tiina knows that in this profile, too, her recommendations will soon start to circle around the same genres or topics if she keeps watching too much content that is thematically aligned. To avoid this, she tries to systematically outwit the recommendation system by not offering any coherent user data, whenever she is using the profile.

“I think I’m pestering the algorithm, or it cannot understand why I never accept anything it offers. But it** [this user profile] **would also turn into a full bubble if I took what it offered.”

***"There are times when I've put on something that I would never ever want to watch** [while using this profile], **just to see what happens next** [to the recommendations]."*

Despite her efforts, Tiina must stay vigilant. She has noticed that she only needs to watch a couple of TV shows or films which focus on a similar topic for the video streaming service to start recommending related content. So, she keeps systematically deceiving the recommendation system – and thinks that it is worthwhile.

"It's almost like on Facebook. If you change your gender and age there, you'll get a totally different view. You are being offered scooters and all kinds of fun things. But [say that] you are 50 – they will start promoting incontinence products."

Shortcuts for adventures in the mind

NEWS RECOMMENDATIONS

38-year-old **Miikka** *(he/him)* from Helsinki describes the relationship that he has to online news as something akin to the relationship that many people have to social media. For him, browsing news is often a way to entertain himself. It offers a break from whatever activity he is in the middle of.

At its best, news also exposes Miikka to completely new perspectives or ideas that he would not normally encounter – an experience which he enjoys a lot.

"I'm interested in having adventures in my thoughts. Because that takes me to a new place."

Inspired by his thought adventures, Miikka wonders if news sites could help these mental collisions to occur more often. To make his point, he draws an analogy between news sites, a global crowdfunding platform that he sometimes visits, and his *Facebook* account. First, he talks about his visits to the crowdfunding platform. Whenever he goes there, he is confronted with a dilemma of how to choose which projects he might be interested in studying amongst the dozens and dozens of projects which the service is displaying. In this situation, the information that crowdfunding service

offers about the projects his acquaintances have backed up comes in handy. Miikka uses it as an aid to decide which projects he could look at.

Next, Miikka moves on to talking about Facebook. He describes how his Facebook contacts often share links on the service which he encounters when he is browsing the service. Here, again, Miikka applies a bit similar method as in the crowdfunding site to quickly evaluate whether the shared link is worth his attention. He uses his previous knowledge about the person who has done the sharing to quickly assess the offered link.

"Let's say that Heikki here has shared this story. He has a religious background…These days he is politically heavily leaning to the left and has worked in an advertisement agency…So, *[based on this background information]* ***you can conclude that he must have been particularly interested in this thing in this article."***

With the two examples Miikka wants to underline how, in both cases, he uses his acquaintances as signposts – or "stop-overs", as he says –, to potentially interesting content that could open doors for new thought adventures. Choosing where to go is hard when you are outside of your normal bubble. To make the task more manageable, Miikka has learned to use other people as his shortcuts.

Inspired by his experiences, Miikka wonders if news sites, too, could provide some kinds of shortcuts for potentially relatable content to help users when they want to step out of their bubble. If they could, Miikka thinks, users might eventually be exposed to a greater variety of perspectives and ideas.

Familiarity cues direct users' out-of-the-bubble explorations

In the situations where people want to make a jump outside of their normal bubble, they often encounter a practical problem: when all unfamiliar content is, per definition, equally unknown or foreign to them, how can they decide where to look? Though people enjoy exploring the unknown and like encountering quirky surprises, they are also conscious of their time and effort. They do not wish to spend endless amounts of their time wading through content that is too distant, completely unrelatable, or maybe even off-putting. Instead, they hope to find ways that will make

it easier to locate hidden jewels in the midst of abundant – and unfamiliar – content.

"I don't tend to just scroll and look and click on something which I don't know, because I think I won't like it."

Lina, 28, Amsterdam, the Netherlands, user of a music streaming service

Through practice, many people have developed methods for making the effort easier. One way is to use other users as shortcuts to potentially interesting content outside of the person's normal bubble. Both Miikka and Gerardo have learnt to do this, though both have adopted a slightly different way of doing it.

Beeline Navigation / Unsplash

Another common method is to use emotional states as a starting point for the user's personal exploration, which we noticed especially when we studied users of music streaming services.

Here, the method is largely the same as the idea of using other users as shortcuts. People who want to find music outside of their normal bubble know that their personal recommendations will not help, tied as they are with their user history. To move away from their personal recommendations, many turn to mood-based playlists. They conclude that browsing these playlists will provide a handy way to put distance between themselves and their personal recommendations, even if the mood-based playlists will only help up to a point. Thus, in an attempt to find unfamiliar but relatable content, many people often still venture further away by, for example, starting with an artist that they've chosen from a mood-based playlist but moving elsewhere by exploring the artist's connections to other artists – and these artists' connections to other artists.

In these situations, users' perceptions of different moods, and their prior experiences about their emotional responses to music, provide proxies which

help users to evaluate unfamiliar playlists. In other words, the user's understanding of moods offers a set of easy attributes that they can resort to when they are comparing different playlists and make decisions about which playlist might open the most promising avenues to interesting content outside of their bubble. **Manesh**, 18, from Chennai, India, explains:

"That mood-based playlist focuses less on me and more on what I need...Sometimes, when you want a transition, you want to let go of the past a little bit. So, it's like a refresh or a reset that you can find"

"It is what you need at the moment: do I want to be anchored in nostalgia or do I want to fly into another world? Those two options, very different from each other."

Implications for designing interactive recommenders

We believe that embracing the methods which users commonly deploy to ease the "transition" from their normal bubble to the territory of unfamiliar content will be a key for designing recommendation systems that can better assist users to explore the abundance of unfamiliar content. For this reason, the design ideas highlighted here are grounded on the two common methods which we have described above.

We want to underline one key similarity between the two methods. In both cases, people look for familiar measuring sticks that they can use to assess, manage, and rate the content that they have no previous experience of. In the absence of any other meaningful frames of reference, other users' preferences or the user's own perceptions of emotional states will act as these measuring sticks. Or, to put it more theoretically, people's perceptions of emotional states or other users' preferences provide what we refer to as familiarity cues which people can rely on when they attempt to make the unfamiliar content seem less unfamiliar.

Having said that, it is important to stress that other equally relevant familiarity cues probably also exist, besides the two mentioned here.

Currently, however, many services already showcase other users' preferences, and video and music streaming services, for example, routinely present content based on the emotional reactions which different kinds of content evoke. Because of this, other people's preferences or users' own perceptions of emotional states often provide natural and the most easily available familiarity cues in a situation where the user is trying to assess unfamiliar content.

User archetypes as shortcuts

Users' tendency to make use of other people as their shortcuts to potentially interesting content has been already acknowledged by many services. Nowadays, various services from social media to crowdfunding platforms routinely notify users about the content that their personal contacts most engage with.

Such a mechanism may often be of limited help, however, when users wish to leave their bubble. Just think of Gerardo's explorations of Kazakhstan music! It is not likely that he would get valuable tips about contemporary Kazakhstan music from his normal contacts.

Thus, to make users' out-of-the-bubble explorations easier, we suggest that providers of recommendation systems could study the possibility of building tools that will allow users to tap into the wisdom of complete foreigners also. We imagine that tools that provide a chance to study, compare, and use such "human shortcuts" might be appreciated.

To develop these tools, system providers could use the data that they already have about typical content preferences which match with specific tastes or interests. Based on this information, they might, for example, build distinctive user archetypes with different preferences, and let users select from them. This way, people could use the constructed archetypes almost as if they were different kinds of login identities. During the out-of-the-bubble explorations, they could choose an archetype as their shortcut to potentially engaging, unfamiliar content in a somewhat similar way as Miikka has learnt to use his acquaintances. Through the chosen archetype, they could thus temporarily leave their own recommendations behind and view content from a completely new angle.

To make a feature like this work, though, system providers should ensure that people can understand the key differences between the different user archetypes without having to devote much time in studying them. In other words, the archetypes ought to provide distinctive characteris-

tics which users can latch onto, when they are deciding which archetype they will want to pick.

A mix of attributes as familiarity cues

Another way to help users to find relatable content outside of their bubble might be to provide users with tools that make their out-of-the-bubble explorations more imaginative and thorough. Presently, as we have seen, people in these situations often use other users' preferences or their own perceptions of emotional states as their familiarity cues and gateways to unfamiliar content. While these methods do help people to exit their bubble and make unfamiliar content more accessible, using them makes explorations quite arbitrary and may leave a lot of valuable content undiscovered.

We think that providers of recommendation systems could see whether there is a way to help users to conduct their out-of-the-bubble explorations more extensively and systematically. One way to do this might be to develop additional or alternative familiarity cues which users could resort to in an effort to sort and assess content outside of their own bubble.

For such solutions, the key for success would be to offer familiarity cues in a highly accessible way. Imagine if a recommendation system could, for example, display a broad set of vastly different attributes that users could easily mix and combine under an "umbrella", as **Layan**, a 26-year-old user of a music streaming service from the Netherlands, envisions. The offered attributes might represent both the normal genre- and keyword-based categories, as well as the less apparent meanings which users may assign to the service's content ■

Suggestions

In the first part of the report, we have described the tensions that currently exist between users and recommenders, due to the systems' incapability to anticipate and accommodate people's changing preferences and shifting everyday circumstances.

In the few remaining pages, we will summarize the findings, hoping to inspire system providers to start experimenting with building interactive recommendation systems.

5. Towards the next generation of recommendation systems

The four tensions between users and recommendation systems described in this report show that some of the current problems with personal recommendations can only be solved by letting people participate in the generation of their recommendations. New, more interactive recommendation systems will be needed.

The key reasoning which drives the development of personal recommendation systems today goes something like this: The current age of abundance – coupled with consumers' increasing access to this abundance through the Internet – makes choosing between available options often a very difficult task for individuals. Recommendation systems should make this decision-making convenient and easy by offering accurate personal recommendations, thus removing the need for users to involve themselves extensively in searching or comparing between options.

According to this view, recommenders must do their work unassisted because people always prioritize efficiency and won't waste their time tweaking their recommendations. A smart system will manage on its own.

We have spent the early part of this report describing occasions when recommenders are not able to accomplish the above stated goal because they fail to anticipate contingencies in people's preferences and daily lives. Through the examples we have shown the limits of recommendation systems to accommodate human change and detailed how the discrepancies between people's expectations and personal recommendations regularly create tensions between users and the systems, thus pushing people away from their recommendations.

In these situations, recommendation systems' inadequate contextual capabilities, reliance on aggregate user data, and the use of conventional content categories cause tensions since they inevitably lead to a reduction and an oversimplification of human complexity. During these occasions, users would want the recommendation system to see them as what they are: unique individuals with varying needs. When recommenders are not capable of doing this, people often deem the generated recommendations as too "dumb", as one of our respondents said, or complain that their recommendations are too generic, or misplaced to offer any real value.

And over time, if similar frustrations continue, they sometimes start showing symptoms of what we have called 'algorithmic fatigue' – a weary disenchantment over recommenders' conventional categorizations and algorithmic sorting [1].

A new paradigm of the user relationship is needed

Against this backdrop, we suggest that a more nuanced view of the relationship between users and recommendation systems should replace the prevailing paradigm of maximal user convenience. This view should be based on an understanding that people do not always want to be passive recipients of automated recommendations. There are occasions when they are prepared to participate in the generation of recommendations in exchange for more relevant personal recommendations – as long as the required effort remains reasonable.

Looking from this standpoint, the relationship between users and recommendation systems becomes one that can, and should, change, depending on the nature of the situation. It is a relationship that will allow people to easily move from letting the recommendation system work on its own to helping the system to make sense of their preferences, or vice versa.

As a logical next step, this report has thus put forward an idea that system providers should start building interactive recommenders which will allow users to smoothly alter the level of their involvement according to their needs. We see this mindset change as the best way to channel the different strengths that recommendation systems and people have into tackling the tensions which currently exist between the two parties.

On the next couple of pages, we will briefly summarize the key design suggestions that have been presented throughout the first part of the

[1] See **Hantula** & **Korkman** & **Laakso**, *"Engaging with EverydAI. How to move AI from imperfect algorithms to perfect user interactions"*; **Ruckenstein, Minna**, *"The Feel of Algorithms - Data, power, emotions, and the existential threat of the unknown"* [manuscript submitted for publication].

Alex Nemo Hanse / Unsplash

report. Before doing that, we want to conclude the text by asserting that people are ready for interactive systems. After engaging with different kinds of recommendation systems for years, they are now savvy enough and possess the needed skills to move to the next level in their relationship with recommenders.

To move forward, the next step should therefore be taken by system providers. They should realize that time is ripe for starting to build interactive recommendation systems which will embrace user savviness, turning it into an integral part of a better user experience ■

6. Key suggestions on designing interactive recommenders

Interactive recommendation systems will allow people to get involved in the generation of their personal recommendations when the user wishes to do so, and they will do this by providing possibilities for new user interactions. Designing these interactions will require service-specific understanding, but the report's main findings, summarized below in the form of four design principles, may help system providers to focus their attention on vital issues.

KEY FINDING 1

Let users modify their data to accommodate relevant changes in their preferences

Users' preferences fluctuate, but sometimes these changes seem so important to the user that they wish that their personal recommendations could swiftly reflect the change. Interactive recommendation systems can assist people in these situations by providing easy ways to modify the data that is used to generate their personal recommendations.

This report has suggested that system providers could explore ways to allow users to temporarily "mute", increase or decrease the impact of certain data sequences in their user data. Providers of interactive recommendation systems could also develop user tools to let people easily make temporary changes in the overall volume of recommended content or allow users to periodically block specific topics from their personal recommendations.

Relevant design questions for system providers to consider:

- If people wish to temporarily block a specific topic or change the number of recommendations that they receive from a certain topic, what kinds of new user interactions are needed?
- Is there a way to let users easily specify how they would like the recommender to treat data from an exceptional use period?
- Is it enough if our users can modify their data afterwards to correct a mismatch that has occurred between personal recommendations and their actual preferences? Or do the users also regularly encounter situations where they ought to be able to make pre-emptive changes to their future recommendations?

KEY FINDING 2

Assist users to broaden their most enduring interests in meaningful ways

People who have a long-lasting topic of interest often proceed along similar paths, regardless of what they are interested in: they gradually move from knowing very little about the topic to a stage where they struggle to find content that they are not already familiar with. Currently, this evolution is poorly accommodated by recommenders because they do not properly acknowledge how this transition changes people's expectations towards their personal recommendations.

This report has concluded that interactive recommenders could serve users in these situations by providing ways to broaden the user's personal recommendations to overlapping or adjacent topics of interest. An interesting way to do this might be to explore how users with similar interests could be utilized to locate the more subtle and surprising connections between different topics.

Relevant design questions for system providers to consider:

- How can we better recognize users who need broader personal recommendations?
- How can we improve the recommendation system's ability to recognize subtle connections between different subjects across the entire content universe?
- How can we help the user to better recognize subtle connections between different subjects across the entire content universe?

KEY FINDING 3

Provide ways to exit the normal content categories

Recommendation systems largely categorize and group their content based on genres, artists, or keywords, but users are more unruly. Sometimes, they get curious about a subject and start following trails of information which will push them to leave, or crisscross between, the conventional content categories that recommenders operate with.

This report has suggested that providers of interactive recommendation systems could accommodate users' personal inquiries if they developed recommendations which cross recommenders' normal content categories in similar ways as users already do. Another way to assist people's cross-category inquiries might be to allow users to make more complicated searches (i.e., searching with a combination of two different terms, for example) in the service.

Relevant design questions for system providers to consider:

- What types of things do users typically search information on when they conduct their personal inquiries on the service? On which level of granularity do people prefer to operate?
- Which features do people most often use in the service when they embark on personal inquiries? If these inquiries largely lead outside from the service, why is this?
- How could we develop the search function to better accommodate users' personal inquiries, also across the conventional content categories?

KEY FINDING 4

Develop familiarity cues for assessing unfamiliar content

People sometimes want to deliberately expose themselves to unfamiliar content, and when they do that, they often struggle to find content that is relatable and engaging enough. In these situations, users commonly resort to two methods to ease their effort: they use other users' preferences or their own perceptions of emotional states as their familiarity cues to quickly assess where they might plausibly find relatable – but still refreshingly unfamiliar – content.

This report has proposed that providers of interactive recommenders should adopt these user methods as the basis of developing shortcuts to relevant content. One way to do this might be to build distinctive user types that work as human gateways to a broad variety of contrasting content preferences.

Relevant design questions for system providers to consider:

- Which features do people most often use in the service when they look for unfamiliar content? If users in these situations regularly move away from the platform to another service, why is this?
- What kinds of familiarity cues do our users, and users of similar services (i.e., music or video streaming, news applications etc.), typically use when assessing unfamiliar content?
- How do our users, and users of similar services (i.e., music or video streaming, news applications etc.), typically utilize their acquaintances or complete strangers as their human shortcuts to unfamiliar content?

Appendix

The outline of the Everyday AI research project

This report has been produced as part of a two-year-long research collaboration between the Finnish customer insights consultancy Alice Labs and the Centre for Consumer Society Research, University of Helsinki. The research project, titled as *Everyday AI,* started in February 2020 and ended in the summer of 2022. The project has been financed by the *Foundation for Economic Education.*

The purpose of the project has been to explore how AI-powered consumer services, which have become a normal part of people's daily lives, could better accommodate users' aspirations and expectations in all kinds of everyday situations. To do this, the two collaborators have studied user interactions with different AI systems by conducting in-depth qualitative user interviews.

- All in all, 75 in-depth interviews were done with respondents from five different countries: China, India, Finland, the Netherlands, and USA.
- In addition, 10+ shorter interviews with Chinese respondents, and a handful of expert interviews, were conducted.
- Three case studies were conducted in partnership with Spotify, Yle (the Finnish public service media company), and Kela (the Social Insurance Institution in Finland).

Due to the long-lasting global COVID-19 pandemic, nearly all interviews were done online, and not in the locations, as was originally thought.

Based on the interviews and the secondary material gathered, the researchers identified recurring tensions that are formed between users and contemporary AI systems and explored the impact of cultural context on the meanings that people assign to different AI systems. In the latter topic, the Centre for Consumer Society Research has been the active party, and its focus has been on China.

In addition to the research ambitions, an important goal of the project has been to engage members of the AI industry in a dialogue. Alice

Labs has done this by conducting the three case studies outlined above. Two case studies – those conducted with Spotify and Yle – had to do with users' relationship with recommendation systems. This report is an outcome of these collaborations. The third study explored how AI-powered systems could assist public service customers during their benefit application processes. This case study was done in cooperation with Kela by interviewing their customers who had experience of using customer service chatbots.

During the research project, Alice Labs has also had the pleasure to engage with two highly professional Finnish digital service consultancies, *Reaktor* (in 2020-2021) and Solita (2021-2022). Their contributions have been essential in ensuring that the project has become a genuine dialogue between the researchers and the industry. The two industry partners have helped to define the practical implications of the research and published the public reports that have been produced during the project. The reports detail the project's practical findings.

- The first-year report, titled as *Engaging with EverydAI. How to move AI from imperfect algorithms to perfect user interactions*, defined three different roles – passive, guiding, and collaborative – which users of AI systems would like to alternate between, and showed that this currently isn't possible. Instead, AI systems often place users in a passive role when they would rather guide the system or collaborate with it. Or they "force" users to abandon their passivity, even if the users wouldn't want to do this, because the systems make mistakes that people cannot let pass. Alice Labs published this report in May 2021 in cooperation with Reaktor.
- During the second year of research, Alice Labs conducted the three case studies described above and completed this report. The report, published in cooperation with Solita, concludes the project's non-academic research publications.

How our research links with discussion on Ethics of AI

Throughout the Everyday AI project, Alice Labs' research focus has been solidly on the level of user interactions. Despite this, we feel that our findings have a connection to the ongoing public discussions about ethical AI systems also. Through our research we have seen what concepts such as 'transparency' and 'agency' really mean to users, when people interact with different kinds of consumer AI systems in their daily lives.

Based on our research, we know that during these interactions, questions about the system transparency or the user's own agency most often arise when an AI system is unable to produce outcomes that match the user's present needs or aspirations, and the user cannot understand what is causing these problems.

When incidents like that do take place, the most important thing that a user typically wants to know is if and how they could make the system work better next time. Today, there often isn't much that users in these situations can do – or at least nothing that could rapidly change the system's outcomes. We explored these themes in more detail in the first public report, published in 2021.

We think that the more versatile user tools that this report has talked about could be a step in the right direction. They could provide the user with better possibilities to adjust AI system's outcomes to their needs. And by doing this, they could simultaneously diminish the number of transparency- and agency-related questions which people have about consumer AI systems.

Of course, explanations about how algorithms and AI systems work would still be needed. But more effective and versatile interaction tools could provide users with more agency, and simultaneously make the working of these systems seem a bit less opaque.

The more AI systems there are, the more crucial it will become to make sure that users retain enough understanding of how these systems change the activities they are incorporated in, and how they modify people's actions or thinking. For consumer AI systems, building more versatile user tools could provide one way to tackle this problem ■

$$\zeta(z) = \prod_{n=1}^{\infty} \frac{1}{1 - p_n^{-z}}$$

ALICE LABS × SOLITA

Embracing user unpredictability

How to get to the next level in building recommendation systems

Kirsi Hantula & Oskar Korkman
Alice Labs
2022