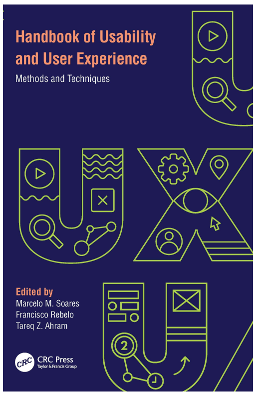

Handbook of Usability and User Experience Methods and Techniques

Marcelo M. Soares, Francisco Rebelo, Tareq Z. Ahram

User Experience and Information Architecture

Publication details
https://test.routledgehandbooks.com/doi/10.1201/9780429343490-18

Luiz Agner, Barbara Jane Necyk, Adriano Renzi

Published online on: 13 May 2022

14 User Experience and Information Architecture

Interaction with Recommendation System on a Music Streaming Platform

Luiz Agner, Barbara Jane Necyk and Adriano Renzi

CONTENTS

14.1 INTRODUCTION

Human-computer interaction has evolved dramatically since the last decade, expanding and revising user experience (UX) concepts. This evolution contributed to shifting trading processes, rethinking companies' culture, changing laws and political scenarios, revisiting privacy concerns and transforming how people interact with each other, with machines in the digital environment. In a world where physical and digital environments merge into one experience involving digital apparatuses, physical ambiences, users, actors, data, algorithms and information in interoperability, information structures are growing dynamically beyond their boundaries. As Pierre

DOI: 10.1201/9780429343490-18

Levy (1993) said, users embrace editors' and protagonists' roles, gaining an active approach toward media. Furthermore, artificial intelligence (AI) has been rapidly adapting to the new technology era in a background layer, becoming a new and significant change agent.

In this scenario, this text attempts to understand streaming music consumption and human-information interaction based on users' mental models about machine learning (ML) algorithms. Through semistructured interviews with Brazilian undergraduate students regarding their experience with the Spotify platform, our research aimed to identify users' perceptions about how algorithms suggest categories and how users feel its relevance in a context where sometimes transparency and explainability are forgotten requirements. Conflicts in the intersection of artificial intelligence with user experience are discussed, as well as the "black box" feeling and other biased user perceptions about algorithms' functioning.

Regarding Spotify, the following research questions inspired us: What is possible to know about the experience quality based on these interactions? Are users satisfied with custom recommendations? Are users able to formulate a mental model on how recommendation systems work? What is the mental model built by users regarding machine learning algorithms? Do users understand how machine learning operates and can interact with these algorithms to improve their outputs and better match them to their goals? Do users worry about privacy issues?

This research also aims to better understand how the information architecture discipline is evolving toward a broader horizon. From classical authors such as Rosenfeld, Morville, Arango, Lima-Marques, Morrogh, Wurman, Resmini and Rosatti (among others), we propose to discuss how the fundamentals of information architecture are being challenged in the era of machine learning. We intend to highlight contemporary tendencies that may reframe the discipline as the reported changes conducted the goals of overcoming information overload and establishing findability as objectives set beyond. These goals are substituted by technological solutions that systematically collect and analyze massive user interaction data to display relevant information. This work describes interconnections between information architecture and recommendation systems, considering its purpose of creating personalized categories and filtering data to present meaningful content.

14.2 INFORMATION ARCHITECTURE: FROM WEB SITES TO ECOLOGIES

If the user experience could be metaphorically compared to a large umbrella, including several correlated disciplines—such as usability, content strategy and user research—as Willis (2020) has done, information architecture would undoubtedly be one of its main strands.

As Rosenfeld, Morville and Arango (2015) have established, information architecture is the discipline for the structural design of digital landscapes, through the synthesis of organization, labeling, navigation and search to build physical, digital and transmedia information ecosystems. Its goal is to develop experiences and products that provide the best usability, findability and comprehensibility.

According to Albuquerque and Lima-Marques, it can also be seen as "a systematic effort to identify standards and create methodologies for defining information spaces" (Oliveira, Vidotti and Bentes, 2015). Its purpose includes the representation and manipulation of information and the generation of relationships between entities to construct information spaces. According to Benyon and Resmini (2017), information architecture relates to structure, access, retrieval and use of the content, focusing on space, navigation and way-finding.

Rosenfeld, Morville and Arango (2015) pointed out that today information has become more abundant than ever before: with the ubiquity of media, a series of devices and objects connected to the Internet for all kinds of daily and routine activities, in homes, offices or urban spaces, have come to configure new and diverse ways of human beings interacting with information. The process has caused the escalation of a problem known as "information overload" (Wurman, 2001; Morrogh, 2003), which we will discuss later.

It is essential to determine priorities to our attention and interest, create hierarchies, information structures, groups, categories and sequential interaction according to people's affinities and expectations. Information architecture is responsible for creating structures to effectively allow users to transform their informational necessities into actions and reach their goals. In this sense, the traditional information architecture's role is to organize and structure the information to help users discover and consume content as well as facilitate their decisions and actions.

Roselfeld, Morville and Arango (2015) believe that information can be comprehended through four interdependent systems:

(i) *Organization system:* This determines how content should be organized, classified and presented.
(ii) *Label system:* This determines verbal and visual signs for each information element and support navigation.
(iii) *Navigation system:* This specifies routes within the informational space.
(iv) *Search system:* In this system, users call upon typing questions to achieve desired information.

For Rosenfeld, Morville and Arango (2015), findability and comprehension are the key objectives when structuring information—emphasized by the multiplicity of channels and diversity and dynamism of technological apparatuses in systems ecology. Nowadays, content is building up dynamically in new connected hyperlinks. As the emerging participative culture grows, users cite, comment, share, reinterpret, edit, mix, create and recreate information through interconnected channels to build their journey. The possibilities of channels involve not only new digital devices but also physical spaces, services and new actors—involving algorithms. Even if the ecosystem is not prepared for bridging channels, users will blend spaces and actions by themselves (Renzi, 2017), finding their paths to a smoother and trustworthy experience.

Since the 1960s, information architecture concepts have evolved as new possibilities of interaction; and although expanding to new visions, its essence remains the same. Our ubiquitous reality with pervasive experiences urged for new information architecture understanding, not restricted to Web sites or apps, as the information flows out artifacts. Resmini and Rosatti (2011) pointed out a list of transformations from the classic information architecture to a new kind of experience:

(i) *Information architecture as an ecosystem:* Different media and different contexts are integrated; therefore, no artifact can be considered an isolated device. Artifacts now belong to a large ecosystem.
(ii) *The users' new role:* Users contribute and actively produce new information or edit/add to something already published, inserting links, comments or critics.
(iii) *Static versus dynamic:* The users' active role makes the structure forever unfinished, permanently in change and open to constant refinement and manipulation.
(iv) *Hybridization:* Boundaries that separated different media, genres, entities and domains are becoming blurred and spanning different environments.
(v) *Horizontality:* The correlation of elements overrules traditional top-down hierarchies. Users push the system to spontaneous, ephemeral or temporary semantic structures.
(vi) *Focus on experience:* Design changes its focus from individual devices to planning and developing experiences and processes.
(vii) *Experiences are now cross-media:* According to Resmini and Rosatti (2011), experiences connect different environments and media in ubiquitous ecologies, a process where all parts contribute to building one whole experience journey.

User interactions in ecosystems are now collected and interpreted by algorithms to respond to desires and train the continuous learning system. Through machine learning, ecosystems can now collect data, classify subjects, cluster related actions, display ads, filter items and recommend personalized content, generating additional complexity to the functioning of information architecture—as we will detail in the next sections.

14.2.1 Information Architecture: From Ecologies to Machine Learning

Given the impressive achievements of recent advances in intelligent machines, several platforms somehow incorporate AI, or at least machine learning. These technologies have significantly impacted the user experience and how content is organized, discovered and accessed in places replete with information (Wallach, Flohr and Kaltenhauser, 2020).

Notorious examples of "AI summer" (the ongoing promising period of accelerated development and funding) are content recommendation systems, online marketing and advertising, voice recognition, speech-based systems, face recognition and stand-alone cars, among many other specific applications. This chapter is particularly interested in discussing user interaction with content recommendation systems due to its natural relationship with information architecture, an essential part of our theoretical framework.

ML-based recommendation systems are answers to one of our age's central problems—traditionally addressed by information architecture—the information overload. As we know, information is being produced at a rate that exceeds the human ability to find, review or understand it. In this sense, information management has become one of our main challenges, on both individual and societal levels. It may be the cause of information anxiety (Wurman, 2001). When faced with the extraordinary profusion of online content, traditional information management strategies are no longer useful, Morrogh (2003) warned us. Bawden and Robinson (2020) stated that although the problem has been there for centuries, it has become more urgent since the arrival of ubiquitous digital information. Among the best ways of avoiding information overload and guarantee findability, there are strategies such as filtering or better design of information systems.

Faced with the current profusion of environments and services with vast amounts of content, we must keep in mind one of the central concepts of the discipline of information architecture—the idea of findability. According to the information architect Morville (2005), findability is an essential notion—that connects to the degree to which a specific object can be discovered or located; and alternatively, to the degree to which an environment can support user navigation and item retrieval.

In this context, machine learning opened a new era for information architecture—considering findability in information-rich ambience. Its approaches are widely employed nowadays to analyze large bodies of data in cloud-based computing, extract patterns and support users' decision-making. As examples, we can cite the choice of a movie to watch on Netflix, a product to buy on Amazon or a song to listen to on Spotify. ML-based systems have significantly increased and lead us to question how these new techniques can meet users' information search needs while respecting human-centered design principles. For Wallach, Flohr and Kaltenhauser (2020), the UX discipline's traditional evaluation methods could underpin ML-based products' best development.

14.2.2 Discussing AI and ML

Artificial intelligence seeks to study the mathematical principles of learning applied to computers. The European Commission's High-Level Expert Group on Artificial Intelligence stated that AI systems are software or hardware that act by perceiving their environment through data acquisition. These systems interpret the collected data, process the information derived from this data and decide the best action(s) to take to achieve a given goal. AI systems can also "adapt their behaviour by analyzing how the environment is affected by their previous actions."

For Bengio (2019), human learning is not limited to reading books or accumulating facts or data; it adapts to the environment's stimuli. Learning means integrating information that we obtain through experience into abstractions that allow us to make better decisions, understand the connections between the things we see and predict what will happen next. In artificial intelligence, we work with the notion of generalization: the machine can generalize from things it has already seen and learned to new situations, in a slow and gradual process, totally based on experience, through contact with a large volume of data. It is what is called machine learning or statistical learning. Machine learning has emerged as a topic of great interest within the field of artificial intelligence research.

Based on Lovejoy and Holbrook, we can consider machine learning as the technique of making computers unveil patterns and relationships in a set of data, dispensing with their manual programming. Lately, digital product developers have had more access to ML tools and techniques and have been more dedicated to creating better user experiences. These experiences are now more personalized and customized; they are engaging and present in various services, with examples from streaming music to autonomous cars.

Machine learning is the study of algorithms that improve automatically through experience (Hechler, Oberhofer, and Schaeck, 2020). The following are the three main types of machine learning: supervised learning, unsupervised learning and reinforcement learning.

(i) *Supervised learning:* According to Honda, Facure and Yaohao (2019), the primary characteristic of supervised learning systems is that the data used to train the model contains the desired response. The training data is previously annotated with the answers or categories to be predicted. We can mention linear regression, logistic regression, artificial neural networks (ANNs), decision trees and others among their best-known techniques. Examples of supervised learning include predicting products that customers are likely to buy or classifying images based on content (Hechler, Oberhofer, and Schaeck, 2020).

(ii) *Unsupervised learning:* It does not require a previously labeled sample of training data: its purpose is to cluster data points into different groups (Hechler, Oberhofer, and Schaeck, 2020). Unsupervised learning problems are considerably complex; thus, this model is at the frontier of machine learning knowledge. Examples of unsupervised learning are film or music recommendation systems, anomaly detection and data visualization. Among the known techniques are artificial neural networks, expectation-maximization, hierarchical clustering, essential components analysis, etc. (Honda, Facure and Yaohao, 2019).

(iii) *Reinforcement learning (RL):* It is based on learning through trial and error. In this approach, an agent interacts with the environment to learn how to perform tasks, improving actions based on rewards and punishments (Hechler, Oberhofer, and Schaeck, 2020). RL considers uncertainty and incorporates changes in the environment into the process of making the best decision. It is based on Skinner's behavioral psychology (Honda, Facure and Yaohao, 2019). As the model learns from the experiments, it is expected that it selects the outputs that generate the most significant rewards for each situation and avoids actions that bring the poorest rewards. The process is repeated until the computer can choose the best action for each scenario.

14.3 USER EXPERIENCE AND MACHINE LEARNING

14.3.1 Problems in UX-AI Intersection

Cramer and Kim (2019) denounce that the field where the user experience meets artificial intelligence is full of tensions. Addressing the intersection where UX has intertwined with AI, the two authors listed a range of problems.

Information architects and UX professionals well know the first reported UX-AI tension: there are conflicting goals for the various actors involved: users, business and company stakeholders. It implies that turning machine learning into a positive user experience will require proper integration between algorithm programming, business models and UX. Machine learning developers are trained to predominantly value features and quantitative aspects, such as accuracy metrics or click rate. However, they should be based on holistic criteria that consider the results for humans, ruling out possible harm to them—and this is not an issue yet solved.

The authors highlighted that algorithms are difficult to explain, which creates imbalances of power and understanding about how to influence their results. Cramer and Kim (2019) also warned us that if today we are facing a growing awareness of the idea of algorithmic responsibility, we must direct our attention both to human and social tensions as well as to ethical and environmental tensions. The spread of machine learning algorithms has encouraged the proliferation of principles to guide development. Fjeld et al. (2020) have pointed to a growing consensus around eight trends that should guide AI: privacy, accountability, security and safety, transparency and explainability, fairness and non-discrimination, human control of technology, professional responsibility and promotion of human values.

So, how to put algorithms under human control? It is a question already addressed by other researchers, such as Lovejoy and Holbrook (2019). They postulate that UX professionals need to learn a lot more about machine learning to help users feel in control of this technology. As the authors state, an in-depth study of users' mental models is essential because in machine learning there is a connection between the algorithms and these models. As people interact, they change the outputs they will see in the future, and this will influence and change the way people interact. That is a feedback loop.

Researchers observed that "conspiracy theories" regarding algorithms could arise from the feedback loop and influence user-algorithm interaction (Lovejoy and Holbrook). These "conspiracy theories" are generally bad for both users and the algorithms' results because users may create incorrect mental models and manipulate the results according to imaginary rules.

The usability researcher Budiu (2019) has profoundly studied the interaction between users and ML-based systems and has found that people produce weak mental models regarding algorithms; therefore, they lose control over their results. Budiu's research observed that recommender algorithms are not transparent to users. People cannot associate their interactions and results based on an adequate understanding of the process logic.

Users often consider the list of suggestions and recommendations as meaningless or random. Concerning information structure, algorithms often create categories according to obscure and not mutually exclusive criteria. From an exclusively statistical or mathematical perspective, grouped categories may even make sense, but they do not function as intuitive information architecture for humans. What we do have, in most cases, is a taxonomy centered on the logic of the algorithmic model.

One of the main problems signalized by Budiu is the so-called black-box model. The black-box model's issue has already been addressed by Herlocker et al. (2000) in their original article. According to the authors:

> *Often, there is not the opportunity or possibly the desire to convey the conceptual model of the system to each user of the system. In such cases, the ACF [automated collaborative filtering] system becomes a black box recommender, producing recommendations like those of an oracle.* (Herlocker, Konstan and Riedl, 2000)

So, when analyzing collaborative filters for recommendation systems based on ML algorithms, researchers point out that:

> *It is inevitable that some users will form incorrect conceptual models of the ACF [automated collaborative filtering] systems that they are using to filter information.* (Herlocker, Konstan and Riedl, 2000).

The above situation is unsatisfactory because users need to develop a robust mental model of how the system works to interact correctly and change or correct its results. The system must be clear about how it processes data and how people can change outputs, thus putting users in control. However, Budiu observed that users perceive cryptic and out of control inputs and outputs (Agner, Necyk and Renzi, 2020).

Budiu also noted that imperfect calculations might cause interesting items to be hidden and the presentation of low-relevance items during user-algorithm interaction. Similarly, the order in which the items appear may not make any sense. On the other hand, too much time between user interactions and the production of outputs that reflect those interactions makes understanding even more complicated.

14.3.2 Basic Guidelines for Human-Algorithm Interaction

There are recommendations for algorithms to increase the quality of the user experience. Harley is a researcher who suggested a list of guidelines to guarantee a satisfactory user experience in ML-based systems. These are just a few primary guidelines as proposed by Budiu and Harley (Agner, Necyk and Renzi, 2020):

(i) *Transparency:* The system must be specific and clear to the user about which people's data are tracked and processed to generate their personalized lists. A clear explanation helps users decide whether recommendations are relevant or not and will add credibility to the system. Working with transparency means informing people about which interactions are considered by the algorithm to help the user build a relevant mental model of interaction. Information and explanations about the model can remove the black-box feeling and benefit users in multiple ways, increasing acceptance (Herlocker et al., 2000).

(ii) *User control:* The UX designers need to provide the user with easy tools to rearrange the output list in a more relevant or familiar way. Categories should be formed more intuitively and close to traditional information architecture; therefore, it is essential to allow users to improve recommendations. Useful features must be provided to enter feedback or edit data to create recommendations and make it more relevant. An example is a tool to edit past actions: deleting items from a user's browsing history or previous purchases. Therefore, the algorithm would be instructed to forget atypical behavior—a way users have to train it.

(iii) *Categories and subcategories:* Personalized recommendation lists must have reasonable findability. The users view them as a valuable navigation resource amid information overload. Besides, regarding ML-generated taxonomy, users prefer to search for content in specific subcategories when browsing extensive inventories, such as Netflix, Amazon or Spotify. It is also essential not to repeat the content within various categories to lower the cost of interaction.

(iv) *Response time:* If users choose to optimize recommendations (by rating, adding items to a favorite list or updating their profile), the expectation is that the result will be fast, mainly when the feedback is negative (Harley, 2019).

(v) *Session-specific customization:* It must be avoided because session-specific art miniatures, descriptions and titles often increase interaction's cognitive costs. The practice of homepage customization (according to a session or device) restricts the learning of the layout, reducing usability (Budiu, 2019).

In many circumstances, researchers identify usability issues with AI systems and suggest solutions. Amershi et al. (2019) have proposed development guidelines, consolidating a set of 18 recommendations. With experts' participation, their work synthesized a guidance for designing human-algorithm interaction. Here we present just a few of their guidelines, selected according to their relevance to the present research:

(i) The user must understand what the system can do and how frequent its errors are.

(ii) The system must display information relevant to the completion of the user's task.

(iii) The user must receive an explanation regarding the system's behavior.

(iv) The experience must be delivered respecting the user's mental model.

(v) When the system makes a mistake, the user should be able to correct it easily.

(vi) The user should be able to provide feedback and indicate his or her preferences.

(vii) The user should be able to customize what the system monitors and how it behaves globally.

(viii) The system must present immediate information on how users' actions will change future behavior.

14.3.3 The Importance of Mental Models

In UX and information architecture, the notion of a mental model is a fundamental concept that has already been defined by Jakob Nielsen as follows: "A mental model is based on belief, not facts: that is, it is a model of what users know (or think they know) about a system." As stated by Lovejoy and Holbrook, a mental model means understanding how a product works and how users' actions affect it.

Considering that the mental model is a user's understanding of how something works (products, places or people), the mental models that do not correspond to reality or maladjusted can cause frustration. However, it is not uncommon for designers or developers to communicate incorrect mental models to users, failing to explain clearly how the product works.

"In case of a mental-model mismatch, you basically have two different options: Make the system conform to users' mental models or improve users' mental models" (Nielsen, 2020). The designer can achieve it by adding transparency features to the interface and better explaining the models. According to Nielsen, "it is a prime goal for designers to make the user interface communicate the system's basic nature well enough that users form reasonably accurate (and thus useful) mental models."

It is also essential to keep in mind that machine learning systems and products adapt, optimize and customize themselves according to their interaction with users. Consequently, it is crucial to communicate correct expectations about how the adaptation will occur, as mentioned above. Through user-algorithm interactions, there will be a co-learning process, as the machine learning models will be transformed as a function of these interactions, while the user's mental model will also be altered. Therefore, it is of utmost importance to communicate the non-human nature and limits of these products to build pragmatic expectations for users and avoid disappointments, preparing them for the changes that will occur after their interactions (People+AI, 2020).

In this context, explainability is one of the most critical AI qualities when considering its cooperation with user experience. Since explainability and confidence are inherently correlated characteristics, UX professionals help users understand and trust the ML system in the correct dosage by helping them build precise mental models.

Meanwhile, there is a recurring problem in the midst of all this. It occurs that explaining to optimize the user's understanding can become a challenge: even AI developers may not fully discern how the technology works to generate their outputs in some instances.

14.4 UNDERSTANDING RECOMMENDATION SYSTEMS

We are living in an era of the preponderance of the Internet and information technology. This scenario has been very challenging in terms of the amount of information and data that ordinary citizens must manage and interact with daily (Morrogh, 2003; Wurman, 2001).

It is essential to rely on some kind of information filtering, and the information architect understands this requirement very well. The number of options to choose from has overwhelmingly multiplied in almost every field of our life: videos, music, books, travel itineraries and restaurants, health safety, courses or even dates with partners. As consumers, finding goods and services classified as "Long Tail" (Anderson, 2006) is an arduous task and does not exempt the use of special filters (Agner et al., 2020; Pandey, 2019). For this purpose, there are computer programs known as recommender systems.

Machine learning algorithmic models can support us in decision-making processes. These recommender engines are tools that belong to the category of information filtering and predict which classification a user will give to an item based on calculations and statistical inferences.

According to Rocca (2019), recommendation systems can be defined as algorithms designed to suggest items relevant to the user in a very general way. These items can be music to listen to, films to watch, texts to read, products to buy or anything else belonging to an extensive catalog. They aim to solve two types of problems: forecasting (data used to predict the evaluation a user will give to an item with which he has not yet interacted) and rating (when creating a finite list of items to be presented to the user), as explained by Pandey (2019). These ML-based systems play an increasingly central role in our lives, as we use them from e-commerce shopping to entertainment services such as music on Spotify or videos on YouTube. These include Facebook, LinkedIn or Twitter, as well as high click-through ads.

According to Pandey (Agner, Necyk and Renzi, 2020), the qualities of a reliable recommendation system are the following: recommended items should be relevant and interesting to users; items that users are not yet familiar with should preferably be presented and lists of items should have diversification.

Recommendation systems use the collaborative or content-based filtering approach, as explained below. There is currently a significant incidence of hybrid systems (those systems that combine more than one approach).

Collaborative filtering relies on users' past behavior (items purchased or ratings given) and other users' similar decisions. Collaborative filtering can recommend items of all kinds without assuming knowledge of their content. As Rocca (2019) remarked, it is based on the (questionable) assumption that people who have already accorded in the past will accord in the future and that people will like items comparable to what they preferred in the past. The interactions recorded over time generate new data and make the system increasingly effective. Therefore, the more the users interact with the items, the more accurate the recommendations become.

We can observe that the mode of data collection to feed the model can be explicit or implicit. Explicit types of data collection include: search terms, evaluating items with like or dislike icons or creating a list. Implicit types of data collection can be observing navigation times, considering songs listened to, viewing product lists or analyzing the social network. The implicit data reflect the records of the user's interactions with the system, which are interpreted as indicators of interest or disinterest.

With a broad set of tracked data, it is possible to describe clusters representing user communities with similar preferences to analyze their collective consumption behavior (Kathayat, 2019).

The second type of approach, the content-based one, requires information about both users and items, such as keywords representing each item and data for each user's profile. For example, age, gender, profession or any other personal data of the user, as well as genre, band, shows, country, music labels or other information for music files. Based on what the user liked in the past or is currently looking for, this type of algorithm finds and presents new similar items, but it has the inconvenience that it cannot become more efficient over time.

As Herlocker, Konstan and Rieldl (2000) explain, collaborative filtering technologies do not necessarily compete with content-based filtering: they can work together to provide a robust hybrid solution. Both approaches have strengths and weaknesses, so the hybrid approach is often adopted, fusing collaborative filtering with content-based methods and others (Pandey, 2019). The Netflix streaming platform is an excellent example of a hybrid system because its recommendations are based not only on consumption habits (collaborative approach) but also on similar video (content-based).

Despite their obvious benefits in mitigating the effects of information overload, and despite their contribution to more immersive and customized experiences, recommendation systems often receive severe criticism from various sources, including their scholars. For example, Irvine (2020) criticizes the use of algorithms because, by choosing practically everything for us and keeping us away from certain choices, they eliminate freedom of choice.

14.4.1 The Algorithms Implemented by Spotify

According to Iyengar (Irvine, 2020), the average individual makes over 70 conscious decisions a day. With the purpose to guide consumers through this labyrinth of choices, recommendation algorithms have become the most ubiquitous machine learning applications for products and services on the Web. A notable example of this is the Spotify platform.

Spotify is a digital, cloud-based music enterprise—founded in Stockholm in 2008—offering cross-device access to over 50 million songs, as well as podcasts and videos (Boyd, 2020). There are over 200 million users worldwide, 100 millions of whom are premium subscribers (Kelley, 2020). According to Johnson (2020), with such a volume of music files, the platform's challenge is: How do you recommend songs with an appeal to its listeners?

As the author stated, there are significant differences between the consumption of films on entertainment platforms and the consumption of music: (i) an extreme asymmetry in the size of the catalogs (Spotify's volume of options is enormous compared to Netflix's); (ii) unlike films, the consumption of songs can happen repeatedly; (iii) the music market is much more typified as being niche than that of movies or series.

Algorithms have changed the way people find, listen and interact with music (Irvine, 2020). This success' components are well-known playlists generated, personalized and cured, such as Discover Weekly, Daily Mixes, Release Radar or recommended suggestions of artists. The service has become one of the most successful platforms in streaming worldwide. Not only because it connects users through music, but also because it employs complementary algorithmic approaches to its recommendations. The methods that Spotify's recommendation system uses are three (Cornell, 2020; Ciocca, 2020):

(i) Collaborative filtering;
(ii) Natural language processing (NLP); and
(iii) Audio models.

One of the first companies to use so-called collaborative filtering (CF) was Netflix, considering users' film ratings to generate the learning about films to recommend to other similar users. After Netflix's success, employing collaborative filtering became popular among companies and a starting point for any recommendation model.

Collaborative filtering and natural language processing are beneficial for connecting listeners to the music that other listeners hear and comment on. Collaborative filtering is a type of algorithm based on metrics, streaming data and user visits to artists' pages. It starts by analyzing a user's behavior, comparing it to other users' behavior. Collaborative filtering approaches revolve around the strategy of determining user preferences from historical behavioral data patterns. Consumer behavior leaves a trail of data, generated through implicit and explicit feedback, which is collected (Johnson, 2020; Ciocca, 2020). If two users listen to the same sets of music or artists, their tastes are likely to be aligned, explains Irvine (2020). The connections between listeners create a list of recommendations that similar listeners appreciate.

It should be noted that, unlike Netflix, initially based on a star system assigned by users, Spotify opted for implicit feedback to train its model. According to Pasick (Irvine, 2020), examples of this feedback could include: songs played in repetition and songs ignored after a 10-second timeout. User data can also come from explicit feedback, such as the heart button in "Discover Weekly" or songs saved in the library or the "Liked from Radio" playlist.

For greater accuracy, Spotify also uses natural language processing to analyze the playlist as a "document itself" (Johnson, 2015). The NLP algorithm is based on blog posts and articles about music to connect music to artists. This model's source data are common words, tracking metadata, news and other Internet texts. At stake is a computer's ability to understand human speech as it is spoken (Ciocca, 2020).

Those algorithms tirelessly scour the Web to determine what users are saying about specific artists and songs. The goal is to identify clusters to determine what adjectives and terminology refer to artists and songs and which other artists and songs connect to them in the comments. The system classifies artists linked by comments on the network into vectors to recommend similar content to other Spotify users.

The third recommendation method is audio models—useful for analyzing raw audio data and recommending music that has not yet become popular. Adding a third model is vital to give more accuracy. Unlike the first two, this algorithm recognizes songs that are still unknown to the public because the audio models capture and consider tunes with few listeners.

This algorithm applies convolutional neural networks that use clustering techniques to identify similarities in tempo, key, mode, time and sonority of audio tracks (Cornell, 2020; Ciocca, 2020). Convolutional neural networks technology is already employed for facial recognition, and it was modified for audio data instead of pixels. After processing, the neural network builds an understanding of each song according to its technical characteristics. It allows us to understand the fundamental similarities between the songs and, therefore, which users may appreciate them based on their consumption history.

The three models described above are used with clustering techniques to produce personalized recommendations of songs, podcasts and playlists, such as the "Discover Weekly." It should be emphasized that the recommendation models only work well because they are connected to a broader ecosystem. This ecosystem includes gigantic amounts of data captured through user interactions (a fact people are not always aware of) and uses numerous clusters to make it work in huge matrices (Cornell, 2020; Ciocca, 2020).

14.5 RESEARCH METHOD: USERS' INTERVIEWS

As Kuniavsky (2003) explained, observation can be crucial for UX research. However, to deeply understand the user experience, it is necessary to ask him or her questions—and this is what an interview is about.

According to Courage and Baxter (2005), an interview is often used to study the user experience. An interview is a guided conversation in which one person seeks information from another. There are a variety of types of interviews, depending on restrictions and needs. A semistructured interview is a combination of structured and unstructured types. The interviewer begins with a set of questions to answer but can deviate from them from time to time. This kind of interview is not so structured, which is easier to analyze data, but it has the advantage of unexpected data.

A screening questionnaire was applied among students in order to select interviewees. Being a Spotify subscriber and an advanced user of the platform were selection requirements. As the data collection process took place during pandemic times, the 2020 second semester, remote communication was mediated by apps like Zoom and similar apparatuses. The questionnaire was useful to select the interviewees and to get to know them better.

As we developed qualitative research, there was no need to select a large sample of the population. We applied one-on-one semistructured interviews to a small sample of design and marketing students from four Brazilian universities with distinct characteristics and backgrounds: Universidade Federal de Pernambuco—Campus Caruaru; Pontifícia Universidade Católica do Paraná—Curitiba; Universidade Estadual Paulista—São Paulo; Universidade Federal Fluminense—Niterói.

The participants selected for the interview are undergraduate university students aged 20–39, mainly 20–25 years old. They are Spotify subscribers for over two years or more, managing personal profiles. The questionnaire and the interviews revealed that there are many media practices in common among them: daily access to the Internet for many hours; use of applications such as e-mail, news apps, social networks, instant messages and entertainment apps, among others; use of streaming content in smart TVs, notebooks or smartphones; file-sharing; and e-learning sites.

Interviewees may also be users of other streaming platforms, like Netflix, YouTube or GloboPlay. They consume music, preferably via Wi-Fi or downloaded songs. They use the app at home, car or on the bus; and while working or playing sports. During the coronavirus pandemic, some reported that they like to listen to Spotify while cleaning their households.

14.5.1 Interviewing Results

Here we present a summary of the responses to the semistructured interviews with Spotify users.

When answering how to find new music on the Spotify platform, some users surprised us by saying that they use Google and YouTube (a competing service). Some users were relying on social networks or recommendation from friends:

> "Sometimes I go on YouTube because I like to watch videos of the artists. I end up researching the artist on YouTube."
>
> "I find bands on the Internet, on Google, or a friend who knows a lot tells me about it. And mainly I find on YouTube, only then I migrate to Spotify."
>
> "I miss something that exists on YouTube Music: the related songs."
>
> "Recommendations are very similar to the Netflix catalogue: I keep passing by, just passing by, and I don't listen to anything."

However, Spotify's recommendations may play an essential role in the user experience. Users told stories that represented their emotional involvement with the platform's sharing resources and reaffirmed some social or cultural benefits provided by the recommendations.

> "I'm quite satisfied because before I had Spotify, my friends knew a lot more bands than me."
>
> "A friend of mine sent me a print of a sad song that I was listening to on Spotify. His web version has shown him. I hadn't talked to him for a long time, and this made us reconnect."
>
> "I ended up knowing a lot of music because of Spotify's recommendations."
>
> "When I was a teenager, I didn't have money to buy CDs, and I only listened to music through piracy. The positive thing about streaming is that the horizon gets much wider, right?"

Continuing the interview, users discussed their customized recommendations and how they could identify them instead of generic contents. Some interviewees indicated navigational, comprehensibility and information overload problems (classic issues related to information architecture). For some others, the personalized category and grouping labeling system has come to attention.

> "I don't think the lists of recommendations are highlighted enough because there is a lot of information, and I always get a little lost."
> "I just don't explore anymore because they recommend a lot all the time."
> "I noticed the group 'Songs that miss you,' and also the title 'Maybe you like it.' … These titles are funny."
> "Why do you have Daily Mix 1, 2, 3 and 4? What's the difference among them? I never really understood."

Some interviewees have shown that they are uncomfortable with generalist recommendations and always prefer personalized ones:

> "I don't think they directly promote songs for which they were paid to promote. On Netflix, I can feel it. … [I]t's terrible, so indiscreet."

Although some respondents were unaware of "Discover Weekly" playlist, others have confirmed to use it. Furthermore, lists or layouts generated to handle specific access devices may generate some confusion.

> "This week, 'Discover Weekly' has just one song that I already knew. These are rock, new wave and classic rock."
> "I always use it. In the mobile application, it's one thing; in the desktop and on the website, the recommendations are different. It's a bug."

We can say that, in general, the respondents considered "Discover Weekly" and other Spotify recommendations as of high relevance. However, the excess of recommended items was again criticized.

> "'Discover Weekly' is excellent. Spotify was killer because I liked 90% of it."
> "The only [negative] point is the excessive amount of songs, of artists that I have never heard of. … I don't have time to stop, to listen, and deepen!"

When asked how their recommendations were created, users could not confidently say what implicit actions were collected. This confirmed suspicions that users cannot build a complete or reliable mental model about how the system works.

> "I've never tried to understand this."
> "Only the searches …"
> "I thought that by adding in a playlist of favorites, I would be giving an indication that I liked. … But now I have doubts if this really instructs the algorithm."
> "Click and follow specific genres, click on releases of the week, I don't know."

"Spotify believes that because I'm listening a lot to a song, I must really like it. It's something indirect, never direct."

"I imagine it takes the styles and artists that I mostly listen to."

"If I listened to a song with the subcategories pop and rock, and another song that is rock and indie, Spotify will show me other bands with the category common to them, which is rock … although associated with other subgenres."

"Maybe they see how much time you spend listening to an artist."

We asked users how they could improve their recommendations. Although there are users who already do this or have a notion of how to do it, the answers indicated that not everyone employs methods to give feedback to the model. Some were thinking about this for the first time.

"I never thought about this."

"I don't know how to interact with the recommendations to improve them."

"I just don't need to show what I'm enjoying."

"I don't feel that giving like or dislike influences the recommendations."

"There were rare cases of songs that I haven't liked. Then I clicked on that little sign that looked like 'Forbidden.' I asked it to hide it so as not to pollute."

"On YouTube Music, the choices are explicit; … in Spotify, it seems to be something organic."

"I hardly evaluate … I don't feel comfortable showing it to anyone."

We asked users who gave feedback if they are satisfied with the response time and whether the system can keep up with their state of mind when receiving feedback. Many of them reported that the recommendations are not up to speed (but this was not unanimous).

"I think it takes a little bit, about a week to change everything."

"I would like the updates to be more frequent than once per week."

"I change my state of mind very quickly. On Spotify, it will take me a week to get a result."

"Recommendations do not go with the ideal speed to my moment of life, my state of mind."

"I noticed a change in my 'Discover Weekly' because of the quarantine: my playlists were very depressive. I think you can give some signs, but the change doesn't happen until a while later."

"The recommendations do not follow my life moment. They stay more or less in the same 'vibe.' If it's a depressive romantic song, Spotify will reinforce this psychological model."

"They don't follow my moments because they are commercial recommendations."

"I don't think it accompanies my state of mind."

"Yes, when I'm happy, starting to date and in love, it usually follows my mood."

In one of the reports, it was made clear that the human-algorithm interaction deserves a more restrained attention because it can have unpredictable consequences.

> "I try not to listen to sad music because of medical advice. I suffer from depression. But it's not every day that Spotify recommends me to listen to happy music."

Our research tried to find out whether users are comfortable having their consumption data collected and whether they saw any sign of a threat to privacy. Most interviewees demonstrated that they are aware of the risks posed by data collection and other privacy issues:

> "It's invasive because Spotify doesn't give me the awareness that he's measuring [and sharing] all this. I get a little scared, you know?"
>
> "I feel I don't have many choices if I prefer not to share my data. But I honestly don't feel uncomfortable."
>
> "I'm a little afraid about sharing data with platforms. There is a website that you connect with [Spotify's] login, and it gives you an invasive result. They found out that young Brazilians listened to depressive music during the quarantine."

Users reported issues related to other platforms that supposedly use private information for hyper-segmented advertising:

> "If you say something, then Instagram is offering an advertisement about it!"
>
> "It's your data and your emotions, music reflects it so much. It's another tool for the company to 'play' with you, so to speak. Selling products to you."
>
> "I said: I am in the mood to buy a skate. Then I went for a drink of water, and when I came back, a roller-skating ad appeared. It's crazy and frightening at the same time!"

Sometimes, respondents have shown a variety of "conspiracy theories" associated with their data's possible irregular use:

> "As in the WestWorld series, there is an exploration of the human mind and desires. I hope that Spotify does not use our information to [share with] others."
>
> "I don't interact, I don't give like or dislike, I don't give answers … I don't feel comfortable."
>
> "I used to share everything about my life and my family, but over time I gained the awareness that it is not exactly a danger of being attacked by a bandit but is more related to manipulation."
>
> "We think we are in control, but we are not. We would be very different if it weren't for this artificial intelligence that keeps throwing things in our faces."

"When you go on the Internet, you're already willing to throw your privacy away."

"I worry if they take my bank account or use my money."

14.6 CONCLUSIONS AND NOTES FOR DISCUSSION

The emergence of the current ML-based recommendation systems has strongly impacted user experiences, with the justification of contributing to smoothing the "information overload" problem. Algorithms demand that we update our look at information architecture—a discipline that traditionally deals with information landscapes' structural design by synthesizing organization, navigation, labeling and search systems.

Our study encouraged us to examine how information architecture has been reshaped by the introduction of collaborative filtering and machine learning approaches. By moving beyond its third phase—that of pervasive information, as identified by Resmini and Rosati—currently undergoing a rapid transformation toward algorithms' structured approach, information architecture is being compelled to coexist with new technological and automated ways to ensure findability and discoverability. That is because we live in times where gigantic catalogs full of options such as Spotify's may make users' navigation and decision-making unviable.

This research tried to understand how young Brazilian users consume streaming content on the Spotify platform. Our goal was to continue previous research, where we reported conclusions on the use of the Netflix platform (Agner, Necyk and Renzi, 2020).

Some user responses showed that recommender systems might involve dealing with typical information architecture issues such as the nomination and category grouping system. They revealed that information overload, findability and comprehensibility problems are not yet completely solved. As one respondent stated: "There is a lot of information, and I always get a little lost."

Experts and institutions involved in user experience and machine learning research generally agree that transparency and explainability are requirements for systems development and design, aiming to overcome the "black box" feeling. ML-based interfaces should inform users how they work, i.e., what implicit or explicit data is collected to create recommendation lists. However, we realized that users generally could not understand the fundamentals of how the system works and do not build a suitable mental model of it.

Explainability is essential because, according to Herlocker et al. (2000), it guarantees users' confidence in the recommendation system. Moreover, it contributes to greater understanding, involvement, education and acceptance of the system as a valuable decision-making aid. So, UX designers should strive to clarify the system's logic, which collects and handles user interaction data to generate personalized recommendations. It will help users to build a more concrete mental model that is closer to reality.

Most respondents have been sensitive to the suspicion that Spotify (and other services or networks) captures—in a non-transparent way—manages and shares their interaction data. When asked about their data's privacy, respondents pointed out fears of obscure use, robbery or marketing manipulation. As in our previous survey, these suspicions shaped what Lovejoy and Holbrook once called "conspiracy theories." A statement made by one interviewee highlighted this sort of concern: "It's your data and your emotions; music reflects it so much. It's another tool for the company to 'play' with you, so to speak."

The problem is that these theories influence how users will interact: some may avoid explicitly reporting feedbacks. It was the case for the user who stated: "I don't interact, I don't give like or dislike, I don't give answers. … I don't feel comfortable."

Some of the users seemed unaware of how they can interact with the algorithm to make the recommendations more in line with their goals or mood. Frequent users do not know precisely how to interact with recommended content to instruct the algorithm to improve lists and better suit their profile and personal preferences. They have not shown to be confident in their knowledge about the universe of possible inputs taken into account by the algorithms.

Human-algorithm interaction emerges as the new frontier of studies involving user experience design and information architecture. It seems that UX designers are not always aware of the recommendation system issues raised by this research. We have addressed some emerging problems that may have been so far neglected by UX professionals. The information collected with young Brazilian Spotify users showed several aspects of recommendation systems that deserve further investigation.

To conclude, we reaffirm the need for UX designers and information architects to update themselves to better understand the technical nuances of ML-based recommendation systems. Such dynamics should contribute to building experiences in which users feel confident and satisfied interacting and collaborating with these systems.

REFERENCES

Agner, L., Necyk, B., Renzi, A.: Recommendation Systems and Machine Learning: Mapping the User Experience. In: Marcus A., Rosenzweig E. (eds.), *Design, User Experience, and Usability. Design for Contemporary Interactive Environments. HCII 2020. Lecture Notes in Computer Science*, vol. 12201. Springer, Cham (2020). doi:10.1007/978-3-030-49760-6_1

Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P.N., Inkpen, K., Teevan, J., Kikin-Gil, R., Horvitz, E.: Guidelines for Human-AI Interaction. CHI '19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. May 2019. Paper No.: 3. Pages 1–13. https://doi.org/10.1145/3290605.3300233

Anderson, C.: *A Cauda Longa: Do mercado de Massa para o Mercado de Nicho*. Elsevier, Rio de Janeiro (2006).

Bawden, D., Robinson, L.: Information Overload: An Introduction. *Oxford Research Encyclopedias*, (2020). https://oxfordre.com/politics/view/10.1093/acrefore/9780190228637.001.0001/acrefore-9780190228637-e-1360. doi:10.1093/acrefore/9780190228637.013.1360

Bengio, Y.: Yoshua Bengio on Intelligent Machines. (2016). https://youtu.be/ePUSElR0o9o, accessed 27 Dec 2019.

Benyon, D., Resmini, A.: User Experience in Cross-Channel Ecosystems. In: *31st British Human-Computer Interaction Conference*, Sunderland, UK (2017).

Boyd, C.: How Spotify Recommends Your New Favorite Artist: A Story of Data, Taste, and a Very Effective Recommender System. https://towardsdatascience.com/how-spotify-recommends-your-new-favorite-artist-8c1850512af0

Budiu, R.: Can Users Control and Understand a UI Driven by Machine Learning? Nielsen Norman Group, (2018). https://www.nngroup.com/articles/machine-learning-ux/, accessed 02 Jan 2020.

Ciocca, Sophia. How Does Spotify Know You So Well? (2017). https://medium.com/s/story/spotifys-discover-weekly-how-machine-learning-finds-your-new-music-19a41ab76efe. accessed 15 Oct 2020.

Cramer, Henriette; Kim, Juho, Confronting the Tensions Where UX Meets AI; Introducing a New ACM Interactions Forum. *ACM Interactions* 26(6), 69–71, Nov-Dec 2019.

Cornell University. Networks. How does Spotify Know You So Well? (2018). https://blogs.cornell.edu/info2040/2018/09/17/how-does-spotify-know-you-so-well/. accessed 13 Oct 2020.

Courage, C., Baxter, K.: *Understanding Your Users: A Practical Guide to User Requirements Methods, Tools, and Techniques*. Morgan Kaufmann, San Francisco (2005).

European Commission's High-Level Expert Group on Artificial Intelligence. Ethics Guidelines for Trustworthy AI. (2019) https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai, accessed 09 Oct 2020.

Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., Srikumar, M.: *Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI* (January 15, 2020). Berkman Klein Center Research Publication No. 2020–1, doi: 10.2139/ssrn.3518482, Available at SSRN: https://ssrn.com/abstract=3518482, accessed 09 Oct 2020.

Google. People+AI Research (PAIR). (2018). https://pair.withgoogle.com/

Harley, A.: Individualized Recommendations: Users' Expectations & Assumptions. (2018). https://www.nngroup.com/articles/recommendation-expectations/, accessed 28 Nov 2019.

Harley, A.: UX Guidelines for Recommended Content. https://www.nngroup.com/articles/recommendation-guidelines/, accessed 28 Nov 2019.

Hechler, Eberhard, Oberhofer, Martin, Schaeck, Thomas: *Deploying AI in the Enterprise: IT Approaches for Design, DevOps, Governance, Change Management, Blockchain, and Quantum Computing*. 1st ed., Apress, Berkeley, CA (2020). doi: 10.1007/978-1-4842-6206-1

Herlocker, J.L., Konstan, J.A., Riedl, J.: Explaining Collaborative Filtering Recommendations. In *Proc. of CSCW 2000*. ACM, New York, 2000, 241–250.

Honda, Hugo, Facure, Matheus, Yaohao, Peng: Os Três Tipos de Aprendizado de Máquina. (2017). https://lamfo-unb.github.io/2017/07/27/tres-tipos-am/

Irvine, M.: Music to My Ears: De-Blackboxing Spotify's Recommendation Engine. (2019). https://blogs.commons.georgetown.edu/cctp-607-spring2019/2019/05/06/music-to-my-ears-de-blackboxing-spotifys-recommendation-algorithm/

Johnson, C.: Algorithmic Music Recommendations at Spotify. (2014). https://www.slideshare.net/MrChrisJohnson/algorithmic-music-recommendations-at-spotify

Kathayat, V.: How Netflix Uses AI for Content Creation and Recommendation. (2019). https://medium.com/swlh/how-netflix-uses-ai-for-content-creation-and-recommendation-c1919efc0 af4. Accessed 15 Jan 2020.

Kelley, C.: Spotify Reaches 100 Million Paid Subscribers Ahead of Apple Music. *Forbes*, (2019). Accessed November 11 2020. https://www.forbes.com/sites/caitlinkelley/2019/04/29/spotify-100-million-paid-subscribers-apple-music/?sh=52fe3161117f

Kuniavsky, M.: *Observing the User Experience: A Practitioner's Guide to User Research*. Morgan Kaufmann, Burlington, MA (2003).

Lévy, Pierre: Tecnologias da inteligência: o futuro do pensamento na era da informática. In: *Editora* 34, 203p. São Paulo, SP (1993).

Lovejoy, J., Holbrook, J.: Human-Centered Machine Learning. https://medium.com/google-design/human-centered-machine-learning-a770d10562cd, accessed 10 Dec 2019.

Morrogh, Earl: *Information Architecture: An Emerging 21st Century Profession*. Prentice Hall, New Jersey (2003), 194p.

Morville, P.: *Ambient Findability*. O'Reilly Media, Sebastopol, CA (2005).

Nielsen, J.: Mental Models. (2010). https://www.nngroup.com/articles/mental-models/, accessed 13 Oct 2020.

Oliveira, H.P.C., Vidotti, S., Bentes, V.: *Arquitetura de informação pervasiva [recurso eletrônico]*. 1a. ed. Cultura Acadêmica, São Paulo (2015).

Pandey, P.: The Remarkable World of Recommender Systems. https://towardsdatascience.com/the-remarkable-world-of-recommender-systems-bff4b9cbe6a7, accessed 28 Jan 2020.

Renzi, A.B.: UX heuristics for cross-channel interactive scenarios. In: Marcus, Aaron., Wang, Wentao (eds.), *Design, User Experience, and Usability: Theory, Methodology, and Management*. Springer, Vancouver, BC, pp. 481–491 (2017).

Resmini, Andreas, Rosatti, Luca: *Pervasive Information Architecture – Designing Cross Channel User Experiences*. Morgan Kaufmann – inprint of Elsevier, Burlington, MA (2011).

Rocca, B.: Introduction to Recommender Systems. (2019). https://towardsdatascience.com/introduct ion-to-recommender-systems-6c66cf15ada, accessed 19 Dec 2019.

Rosenfeld, Louis, Morville, Peter, Arango, Jorge: *Information Architecture for the Web and Beyond*. 4th ed. O'Reilly, Sebastopol (2015).

Wallach, Flohr, Kaltenhauser: Beyond the Buzzwords: On the Perspective of AI in UX and Vice Versa (HCII 2020). Dieter P. Wallach (B), Lukas A. Flohr, and Annika Kaltenhauser. Springer Nature, Switzerland (2020). H. Degen and L. Reinerman-Jones (eds.), *HCII 2020, LNCS 12217*, pp. 146–166 (2020). doi:10.1007/978-3-030-50334-5_10

Willis, D.: The UX Umbrella. (2014). http://www.dswillis.com/talks/2014/4/the-ux-umbrella

Wurman, Richard S.: *Information Anxiety 2*. Que, Indianapolis (2001), 308p.