

What if? Interaction with Recommendations

Martin Zürn

Malin Eiband

martin.zuern@campus.lmu.de

malin.eiband@ifi.lmu.de

LMU Munich

Munich, Germany

Daniel Buschek

daniel.buschek@uni-bayreuth.de

Research Group HCI + AI, Department of Computer

Science, University of Bayreuth

Bayreuth, Germany

ABSTRACT

Showing users recommended content has become a prominent way of integrating algorithmic decision-making in everyday intelligent applications (e.g. recommendations of films, music, news, routes). In this context, the research community has identified *What if?* questions as an approach for users to investigate and question such recommendations – yet many current applications seem limited in practically supporting this. We present a set of example GUIs and interaction techniques currently used in everyday recommendation systems in practice (e.g. Grammarly, Apple Music, Google Maps). Based on these example cases, we discuss possible UI extensions to explicitly support *What if?* interactions. From our analysis and reflection emerges the general approach of treating decision variables as a “first-class citizen” in UIs: We propose to 1) represent a recommended item’s decision variables in the user interface (and not just the item itself), and 2) to enable direct manipulation of these decision variables for *What if?* explorations.

CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability.

KEYWORDS

intelligent systems, recommender systems, explainability, intelligibility, interpretability, end-user debugging, interactive machine learning

ACM Reference Format:

Martin Zürn, Malin Eiband, and Daniel Buschek. 2020. *What if?* Interaction with Recommendations. In *Proceedings of the IUI workshop on Explainable Smart Systems and Algorithmic Transparency in Emerging Technologies (ExSS-ATEC’20)*. Cagliari, Italy, 4 pages.

1 INTRODUCTION

Artificial intelligence has been integrated into many everyday end-user products and services to improve the user experience and to help users navigate an ever-increasing amount of data. To this aim, many of these systems show users recommended content like films, music, routes and news based on complex algorithmic inferences. To date, the algorithmic decision-making process and the decision variables leading to a recommendation are often not accessible to users or hidden in the user interface, making it difficult for users to navigate this “inferred world” [4, 16].

To support users, the HCI research community has worked on several approaches. One of them are so-called *What if?* questions, which allow users to explore, investigate and question algorithmic decision-making. However, in comparison to work that focuses on deriving explicit *explanations* for system decisions (e.g. [14]), concrete design and interaction solutions for this exploratory approach in everyday use are still in their infancy. Notable exceptions include work by Lim and Dey [13] who built a *What if?* experimentation UI that allows users to set specific input sensor values and observe the respective system prediction for those values. Another example has been presented by Nguyen et al. [15] who introduced interactive sliders to adjust model parameters and observe the resulting change in the system output.

In this paper, we argue that *What if?* interaction with recommender systems offers rich – and till now underexplored – opportunities for user support: *What if?* exploration may 1) allow users to understand the system decision-making in an implicit way without the need for specifically crafted explanations, as argued in [5], and in this way exert their “right to an explanation” as part of the European Union *General Data Protection Regulation (GDPR)* [17]. Moreover, it may 2) let them become aware of system errors and, given options for feedback and correction, improve future recommendation, and thus 3) foster overall control of the system.

This is in line with calls for more expressive feedback and correction and fine-grained control options [6] in intelligent systems as well as more interactive system explanation [1, 6].

To provide inspiration for concrete design solutions, we present a set of example GUIs and interaction techniques currently used in everyday recommendation systems in practice, such as Grammarly, Apple Music, Google Maps, and the like. Based on these example cases, we discuss possible UI extensions to explicitly support *What if?* interactions. From our analysis and reflection emerges the general approach of treating decision variables as a “first-class citizen” in UIs: We propose to 1) represent a recommended item’s decision variables in the user interface (and not just the item itself), and 2) to enable direct manipulation of these decision variables for *What if?* explorations.

2 APPLICATION EXAMPLES

In the following section we first present for inspiration some selected, popular examples of interactions that can be interpreted as *What if?* interactions and that are well established and already familiar to users. For this purpose, we identified prominent use cases of recommender systems from different domains, such as shopping, entertainment, news, location-based services, social media, fitness and health, as well as typical decision variables used in these applications.

We then describe other existing interactions outside the *What if?* context, which were not originally designed for this purpose nor are currently used in this context – but which we argue are promising candidates for use in a *What if?* setting for recommender systems.

The examples that we deem most relevant will now be presented briefly. Please note that this is not a comprehensive overview, but a collection meant to inspire reflections towards concrete UI elements that support *What if?* exploration.

2.1 *What if?* interactions in-use as of today

Here we present three example interactions, which already support *What if?* exploration. These examples cover “classic” item recommendation (e.g. restaurants), recommendation for productivity (e.g. routing) as well as recommendation for creativity (e.g. text).

2.1.1 Location: Foursquare. The service Foursquare provides recommendations for restaurants and other places based on users’ location and their search and check-in history.

While the app aims to find relevant venues at the current location, it also allows users to change their search location, for example in a map view. This gives them the option to discover recommendations they would receive if they were at a location different from their current one (i.e. *What if I would be at a different location?*).

2.1.2 Routing: Google Maps. Google Maps not only allows users to view maps and search for addresses and places, but also to generate a route from A to B using different means of transport. For car routes, Google Maps offers two different dimensions that can be explored with *What if?*, namely *time* and *path* of the route.

By default, Google Maps assumes that users start at the current time, and includes current traffic information about the route (such as traffic jams, closures, etc.) and estimated travel time. However, users can also specify an individual time for departure or arrival, and Google Maps will then provide an estimated travel time for that time (i.e. *What if I would start at a different time?*).

Another dimension is the path of the route. Google Maps provides a route suggestion and usually displays alternative routes greyed out on the map. In addition to the suggestions, the user can also customise the route by clicking on a point on the route and then freely moving it on the map (i.e. *What if I would take a different path?*).

The route length and the travel time are always displayed, so that users can easily compare different routing suggestions.

2.1.3 Text: Grammarly. Grammarly is an advanced spell checker that attempts to improve the quality of writing. In addition to providing traditional spell checking, it suggests alternative words and phrases that may better match a particular context and target audience.

For this, users can choose different options for the variables *audience*, *formality* and *domain*. Based on this choice, the system provides different suggestions for correction and improvement of the text. While a user might write for a specific audience and domain, the settings can be changed with one click. Thus, with this type of *What if?* exploration, the user can see the influence of these variables on the system recommendations (i.e. *What if I would change the audience?*).

2.2 Interactions applied to *What if?*

In this section, we consider interactions that are not currently used specifically for *What if?* exploration, but which we argue appear to be promising candidates in this respect.

Specifically, we present *feature weighting* because it is applicable to any kind of recommender system, a *temporal component* in form of a time axis since many recommender decisions change over time (both through new content and changing user profiles), as well as *virtual personas* and *emulation*, as these allow for interaction with and manipulation of physical objects.

We first look at how these interactions are currently used and then transfer them to the context of intelligent user interfaces by describing a scenario in which they could be used for *What if?* exploration.

2.2.1 Feature weighting: Apple Music. To alleviate the cold start problem, users are asked about their personal preferences when they first set up Apple Music (see [8]). Figure 1 illustrates this onboarding process: Users are given a number of items (first a set of genres, then artists) to interact with. Each item is displayed as a bubble, and users can click on it to increase its size (and thus its relevance for system decision-making) or remove it altogether. This allows for a playful interaction with the weighting of different items of a given set.

In the context of *What if?*, this interaction could be used to remove or weight certain features more or less to observe the effect on the model result. More concretely, in recommender systems that recommend further products based on characteristics of a product, certain characteristics could be weighted more strongly by means of this interaction, while others that users do not consider relevant can be removed (i.e. *What if I would change the importance of factor X?*).

2.2.2 Time axis: MacOS. Time Machine is a built-in backup solution in MacOS and allows the user to view a specific document in all available backup states. As soon as Time Machine is started for a document, the states are displayed piled up like a stack of cards, allowing the user to move through the versions along the time axis [2] (see Figure 2).

Since user preferences may change over time, a time axis-like interaction could enable users to apply the status of a recommendation model from the past to today’s inventory. In this way, a “time travel” would reveal system learning over time (i.e. *What if today’s inventory is recommended using my profile from 10 years ago?*).

Moreover, this interaction might be applicable for content which inherently changes over time, such as news or social media posts. The time dimension of the model could also be fixed here, and instead the time dimension of the inventory could be altered (i.e. *What if music from 10+ years ago would be recommended today?*).

2.2.3 Virtual persona: The Sims 4 & Memojis. In the game *The Sims 4* Electronic Arts Inc. [7] players create virtual characters which lives they direct. During this process, players have various possibilities to edit character traits and appearance in great detail – from the width of the nostrils to the distance between the eyes. A similar concept can be found in the *Memojis* feature on iOS [3], where users can create their lookalike as an emoji.

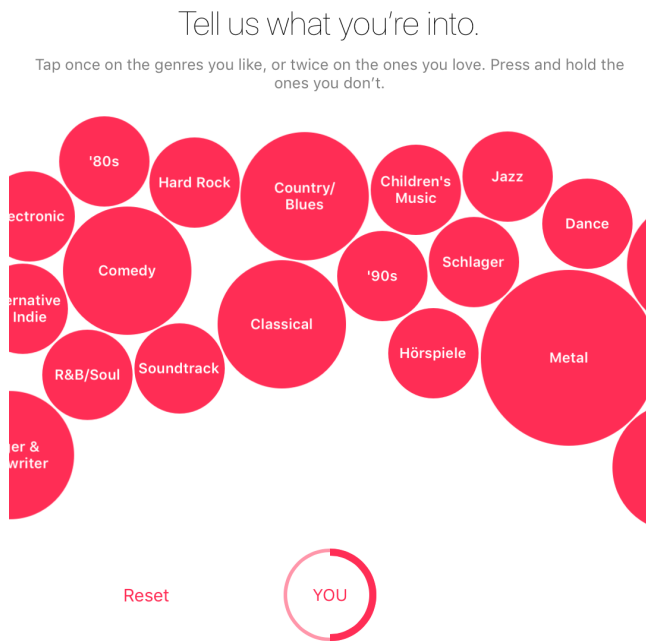


Figure 1: As part of Apple Music’s onboarding, users weight different artists and genres so that the system gets to know their personal preferences.




Figure 2: Time Machine allows users to navigate through the history of a data item.

Transferred to the context of *What if?* interaction, the appearance of such a virtual persona could be used to visualise a specific user model. Allowing users to change their persona (e.g. adapt gender, age or interests) could give them the option to explore system recommendations for other user models. This could be particularly insightful in recommendation systems where the physical appearance is an important factor, e.g. in dating applications such as Tinder or personalised fitness apps such as Freeletics (i.e. *What if I would change my physical shape?*).

2.2.4 Emulation: Car Emulator. *Car Emulator* [9, 10], shown in Figure 3, is an application which lets developers reproduce almost

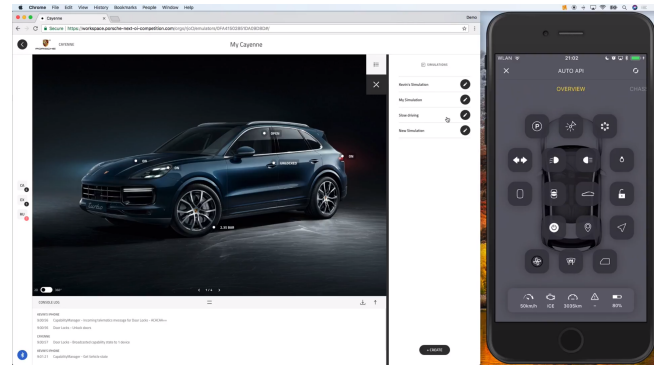


Figure 3: Car Emulator allows to create and manipulate an object’s state [10].

any state of a vehicle for software testing (e.g. open/close doors, lights on/off, etc.).

Such an emulator-based UI approach might be used also to let users manipulate system decision variables in recommendation systems. For example, users could change product characteristics such as colour or brand – and see to what extent system recommendations change. This could not only be applied to the currently recommended product, but also to past recommendations (i.e. *What if this product would have different characteristics?*).

3 DISCUSSION

The examples presented in the last section are meant to provide inspiration for our community to work towards a richer UI design vocabulary for *What if?* exploration in recommender systems – in particular one based on direct manipulation of decision variables. In this section, we summarise general desiderata for such a design vocabulary. Moreover, we discuss overarching implications of a *What if?* approach to supporting users in interaction with recommendations.

3.1 Conceptual view: “Recommendation item” includes its key decision variables

Designing for *What if?* raises the basic question of how we define a “recommendation”. In this paper, we suggest a new comprehensive notion: A recommendation comprises not only the recommended content itself (e.g. a film, product, etc.) but also the key *decision variables*. This view highlights that design solutions for *What if?* could – and should – not just engage users in exploration with the recommended content itself, but also with the decision variables based on which this recommendation was given. Casting these variables as an inherent part of the recommendation highlights their need to be presented in the UI. More concretely, this would enable new direct manipulation interactions for users to influence the decision-making process in different dimensions or at different stages, not only by providing feedback via star ratings or giving a “thumbs up” or “thumbs down”.

3.2 Design solution principle: Rich direct manipulation of recommendations

Most notably, the presented examples foster *direct manipulation* – and through that a sense of *immediateness*: For example, when assigning weights to artists and genres in Apple Music, users do not have to look through a list of filter criteria but instead can immediately grasp and experiment with the system decision-making in a playful way. Similarly, in Google Maps, customising a route is done via dragging the suggested route on the map. Moreover, *Time Machine* and *Car Emulator* support direct manipulation through *affordances* (i.e. through a card stack and a real-life object).

We argue that direct manipulation of recommendations and decision variables in combination with affordances form promising research avenues towards finding ways to integrate *What if?* exploration into everyday applications.

3.3 Challenges

While many expert systems for interactive machine learning such as Tensorboard¹ already support *What if?* interactions, designing them for laypersons in everyday use is challenging. These challenges include designing exploration in a way which does not distract from the actual task users want to do [5], and which does not overwhelm users [11, 12]. Moreover, one fundamental premise for designing for *What if?* in this context should be that *users are not always interested in exploration* – the HCI research community should therefore think about how *What if?* can be seamlessly integrated into the interface, for example via on-demand approaches. Finally, while we argue that *What if?* might foster overall (feeling of) control of the system, future work should explore if and how often people actually make use of this possibility.

4 CONCLUSION

We presented a set of example UIs from various current applications to inspire working towards a richer UI design vocabulary for interactions with recommendations, in particular supporting *What if?* exploration. Our suggestion for a promising design solution principle here has two key takeaways:

First, we propose to conceptually view key decision variables as an integral part of any “recommendation item” shown to the user. In other words, such decision variables need to have a representation in the UI. Second, we propose to then design direct manipulation interactions for these representations of the decision variables. Here, our set of examples provides ideas for possible starting points to work towards concrete UI designs.

REFERENCES

- [1] Ashraf Abdul, Jo Vermeulen, Danding Wang, Brian Y. Lim, and Mohan Kankanhalli. 2018. Trends and Trajectories for Explainable, Accountable and Intelligible Systems: An HCI Research Agenda. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (*CHI '18*). ACM, New York, NY, USA, Article 582, 18 pages. <https://doi.org/10.1145/3173574.3174156>
- [2] Apple, Inc. 2018. *Use Time Machine to restore deleted files or older files - Apple Support*. Retrieved 2019-12-18 from <https://support.apple.com/en-us/HT209152>
- [3] Apple, Inc. 2019. *How to use Memoji on your iPhone or iPad Pro - Apple Support*. Retrieved 2019-12-18 from <https://support.apple.com/en-us/HT208986>
- [4] Alan F. Blackwell. 2015. Interacting with an Inferred World: The Challenge of Machine Learning for Humane Computer Interaction. In *Proceedings of The Fifth*

- Decennial Aarhus Conference on Critical Alternatives* (Aarhus, Denmark) (*CA '15*). Aarhus University Press, 169–180. <https://doi.org/10.7146/aahcc.v1i1.21197>
- [5] Malin Eiband, Charlotte Anlauff, Tim Ordenewitz, Martin Zürn, and Heinrich Hussmann. 2019. Understanding Algorithms Through Exploration: Supporting Knowledge Acquisition in Primary Tasks. In *Proceedings of Mensch Und Computer 2019* (Hamburg, Germany) (*MuC'19*). ACM, New York, NY, USA, 127–136. <https://doi.org/10.1145/3340764.3340772>
- [6] Malin Eiband, Sarah Theres Völkel, Daniel Buschek, Sophia Cook, and Heinrich Hussmann. 2019. When People and Algorithms Meet: User-reported Problems in Intelligent Everyday Applications. In *Proceedings of the 24th International Conference on Intelligent User Interfaces* (Marina del Rey, California) (*IUI '19*). ACM, New York, NY, USA, 96–106. <https://doi.org/10.1145/3301275.3302262>
- [7] Electronic Arts Inc. [n.d.]. *Create Your Sim Demo - The Sims 4 Official Site*. Retrieved 2019-12-18 from <https://www.ea.com/games/the-sims/the-sims-4/pc/create-a-sim-demo>
- [8] GottaBeMobile. 2015. *Apple Music Setup - YouTube*. Retrieved 2019-12-18 from https://youtu.be/ydYah_BXt9A?t=30
- [9] High-Mobility GmbH. 2017. How To Use HIGH MOBILITY's Car Emulator. *HIGH MOBILITY Magazine* (2017). Retrieved 2019-12-18 from <https://medium.com/high-mobility/how-to-use-high-mobilitys-car-emulator-244200558a8a>
- [10] High-Mobility GmbH. 2018. *Porsche NEXT OI Competition: Car Emulator - YouTube*. Retrieved 2019-12-18 from https://youtu.be/y_-Wp7uETBA
- [11] Todd Kulesza, Margaret Burnett, Weng-Keen Wong, and Simone Stumpf. 2015. Principles of Explanatory Debugging to Personalize Interactive Machine Learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces* (Atlanta, Georgia, USA) (*IUI '15*). ACM, New York, NY, USA, 126–137. <https://doi.org/10.1145/2678025.2701399>
- [12] Brian Y. Lim and Anind K. Dey. 2011. Design of an Intelligible Mobile Context-aware Application. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services* (Stockholm, Sweden) (*MobileHCI '11*). ACM, New York, NY, USA, 157–166. <https://doi.org/10.1145/2037373.2037399>
- [13] Brian Y. Lim, Anind K. Dey, and Daniel Avrahami. 2009. Why and Why Not Explanations Improve the Intelligibility of Context-aware Intelligent Systems. *Proceedings of the 27th International Conference on Human factors in Computing Systems*, 2119–2128. <https://doi.org/10.1145/1518701.1519023>
- [14] Tim Miller. 2019. Explanation in Artificial Intelligence: Insights from the Social Sciences. *Artificial Intelligence* 267 (2019), 1 – 38. <https://doi.org/10.1016/j.artint.2018.07.007>
- [15] An T. Nguyen, Matthew Lease, and Byron C. Wallace. 2019. Mash: Software Tools for Developing Interactive and Transparent Machine Learning Systems. In *Explainable Smart Systems Workshop at the 24th International Conference on Intelligent User Interfaces* (*IUI '19*).
- [16] Advait Sarkar. 2015. Confidence, Command, Complexity: Metamodels for Structured Interaction with Machine Intelligence. In *Proceedings of the 26th Annual Conference of the Psychology of Programming Interest Group* (*PPIG 2015*). 23–36.
- [17] The European Parliament and Council. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union* (2016).

¹<https://www.tensorflow.org/tensorboard>, accessed 17 December 2019.