AUD_I/O

SMART AND SOCIAL MUSIC CONTROLS

10/6/2021

MATS ERDKAMP
M.W.G.ERDKAMP@STUDENT.TUE.NL
FINAL BACHELOR PROJECT
FUTURE MOBILITY, INDUSTRIAL DESIGN

UNDER GUIDANCE OF

BASTIAN PFLEGING
B.PFLEGING@TUE.NL

DEBARGHA DEY
D.DEY@TUE.NL

Technische Universiteit Eindhoven
University of Technology

CONTENTS

EXECUTIVE SUMMARY

The following report discusses the design process of AUD_I/O, a car music controller redesigned from the ground up. The *music streaming* paradigm, combined with the ever-growing field of *big data,* has created an opportunity for new and improved interactions with music in the car. If correctly implemented, the resulting product could provide an easier, safer, and more social way to interact with music in the car. The AUD_I/O project aims to find out how the benefits brought about by music streaming and big data are best exploited to benefit the user the most. Implementations of these benefits include but are not limited to

- Music suggestions strengthened by circumstantial data (weather, speed, location, traffic)
- Music suggestions that take all passengers into account.
- Music sharing between cars

This new design fundament forces a kind of symbiotic relationship between the car and the music streaming service, where the vehicle becomes more than just another device that runs applications. Instead, the car becomes a crucial part of the system as a whole. The clearest example of this symbiotic relationship would be the influence of car sensor data on the music recommendation algorithm. The music streaming experience is augmented by a system that is hyper-specific to the car. Another area where the project tries to augment the music listening experience is on the social front. Again, social scenarios specific to the car are exploited to create the best experience. Driver and passengers can more easily bond over music. An interface has been created where driver and passenger can negotiate the prevalence of their music taste in a fun way. This social aspect has also been explored between different cars. This resulted in an inter-car song sharing system, and 'Traffic Jams', a system where a shared regional playlist is created for all the cars stuck in a traffic jam. These systems have been validated through user testing, with general response being positive. A real deployment of the systems described would be the logical next step in this design process.

PROLOGUE

Throughout my bachelor's program, artificial intelligence and related systems have played a big role. USE courses focused on robots and autonomous vehicles, and design projects focused on the impact of data (within Transformative Practices and DIGSIM) are all directly related to my fascination with A.I. and big data. It is therefore not strange that the Future Mobility squad is a good fit for me as a designer. While the AUD_I/O project is not directly related to the artificial intelligence involved with autonomous driving, it does touch upon another interest voiced in my vision: Improving interaction between users and big data-driven systems. Big data has been a product of the companies of the big 21st century, such as Facebook and Google. The fact that these are software companies first has resulted in big data residing in the back-end, away from the end-user. The end-user is limited in their interaction with the data, resulting in opaque systems that alienate users. New designs could make interactions more transparent, while also allowing end-users to harness the power of big data themselves.

It has been incredibly interesting to see all the different designs created in the Future Mobility squad. This fascination for not only my project but all the various projects around me leads me to believe that I made a great choice to do the final bachelor project as a part of this squad.

INTRODUCTION & RELATED WORKS

INTRODUCTION

Music streaming has exploded in popularity in the last decade, but many car interfaces are still very much radio-based with streaming shoehorned into the same interface. This project aims to rebuild car music interactions from the ground up, exploiting new and emerging technologies such as developments in the automotive industry, big data, machine learning, and of course music streaming. This radically different design fundament allows for novel interaction opportunities, as well as streamlining some of the already existing interactions. The project sets out to put on display what a well thought out integration between car and music streaming can achieve.

To get an understanding of why the aforementioned developments might lead to novel interactions, we need to have a look at how the individual industries have been transformed over the last few years. With music streaming taking prominence over other mediums such as radio and CD’s, or even digital purchasing of songs such as provided by iTunes, we can see a clear shift in how users listen to and discover new music. Music streaming services provide a more personalized experience compared to any other medium through the use of machine learning algorithms that recommend music to users based on their previously listened songs. This creates an highly personalized experience for all. In the words of Gustav Söderström, Spotify’s Chief R&D Officer:

“There isn’t just one Spotify experience. There are actually more like 345 million different Spotify experiences—one for each listener."

This focus on creating a personal music listening experience is further exemplified by services such as *Spotify Wrapped* (Spotify, n.d.) and *Spotify Only You* (Spotify, 2021b). These services aim to give the user a fun overview of their music listening habits, the first one being focused on listening trends over the year, and the second one honing in on some of the user’s unconventional music listening habits.

It has also never been easier share music with loved ones. Music has always been an inherently social activity. Think about festivals, concerts, sharing CD’s et cetera. But the streaming paradigm has made sharing music as easy as sharing a link to someone. This level of directness has not been achievable before. Furthermore, special tools are developed to make the music listening experience even more social. Systems such as *Spotify Group Sessions* (Spotify, 2020) allow for multiple people to take control over the music that is playing, while systems like *Spotify Blend* (Spotify, 2021) allow users to merge their musical tastes, making it even easier for them to bond over music. Not many projects have focused on bringing these social features on the car. This creates a design niche to be filled.

Looking at the developments in the automotive world, we can see that the space as a whole has been very disrupted over the last few years. The race towards autonomous vehicles has created a ripple effect on the industry at large. Cars have become driving super-sensors. Intel CEO Brian Krzanich claims that they will generate and consume up to 40 terabytes of data per eight hours of driving (Nelson, 2016). Most of this data is generated in service of the autonomous driving systems, but all of this data could potentially also be of use to aid different systems (such as music controls), Furthermore, software and human-computer interaction play an increasingly bigger role within the car. One clear example is the fact that companies like Tesla now actually roll out over-the-air software updates (Tesla, 2021), something that was unheard of just a few years back. The interiors of these cars have been adapting to accommodate this newfound ‘smartness’. Infotainment systems are moving towards being general purpose touchscreens, with capabilities very similar to what a smartphone can do. A vast amount of infotainment systems even run operating systems derived from a mobile operating system, such as Android Auto (Google, n.d.) or Apple Carplay (Apple, n.d.). These operating systems allow for easy conversion of mobile applications to car-specific software. This often results in interfaces that are stripped-down versions of their mobile counterparts. While this approach works fine for porting over basic functionalities, it fails to consider the emergent interaction opportunities that might arise from the software being specifically in a car. That is why this project takes a different approach. Instead of asking how music streaming can efficiently be ported to the car, we look at what the car can provide for the music streaming experience. This way of thinking forces a kind of symbiotic relationship, where the vehicle becomes more than just another device that runs applications. It instead becomes a crucial part of the system as a whole.

RELATED WORKS

In this chapter, the context in which the design exists is depicted. The design space is explored by looking at relevant research and designs. This mapping of the design space prevents repetition of designs already made, while also creating a clear image of the newest findings, creating a cutting-edge fundament on which to base the design.

Research

Spotify and 'Car Thing'

Spotify has been busy exploring how users interact with music in the car. They have been developing the *Car Thing* for a while now (Spotify, 2021a) Its original purpose was to help create a better understanding of users' listening habits while on the go, and has since been launched as a commercially available product. The *Car Thing* embodies the interest streaming services have in improving their services in the domain of the car. Spotify also labels the car as one of the most common places where people listen to music and a prime environment for exploring how people think about and engage with music (Martelaro et al., 2020). In this paper, they also explore a more conversation-based approach to music control. The extra focus available in a car for conversation allows for some interesting new interactions. For example, during a user test, the 'DJ-bot' told the driver about an upcoming Beyoncé concert, to which the driver excitedly responded "*I need tickets! Thank you for telling me*". These kinds of interactions highlight the unique context created by the car; drivers are relatively free to interact with a speech agent. As highlighted in the paper, it is important to note that the safety of these interactions is dependent on traffic circumstances. A speech-based system should only interact with the user during less engaging moments of driving.

Driver Distractions and Music

A large amount of research has been done on driver distractions and how it affects the driver's behavior. Young and Regan, (2007) reviewed a large amount of this literature, while outdated in some aspects, some insights still hold true to this day, such as decreased mean driver speed when the driver is interacting with a system in the car. Furthermore, distractions are categorized into four distinct types: visual, auditory, physical, and cognitive. Visual inputs that require manual interactions create the biggest reduction in driving performance since these create the greatest dual-task interference. That is not to say that auditory distractions do not reduce driving performance, if the auditory task is complex enough, driving performance will still be impaired. Other studies (Ranney et al., 2002) have compared the degradation in driving performance when comparing visual/manual with auditory versions of tasks. They conclude that the voice-based interface reduces peripheral impairment but does not appreciably reduce attentional impairment. Pfleging (2017) provides an extensive guide on how non-driving-related activities can be designed better. Through multimodal interactions, the necessity for time-consuming interactions can be reduced.

Contemporary Products

Currently, a lot of car manufacturers implement their own music interaction hardware and software. Most of these implementations are similar in concept but slightly different in execution. Through a comparison of systems, assumptions can be made about which interaction methods best fit the requirements of AUD_I/O. These assumptions helped shape the first design iteration of the project.

Touch Versus Physical

Currently, we see a clear divide taking shape in the automotive industry's dashboard designs. With this decade's touch-screen explosion, a lot of car manufacturers, such as Tesla, have decided to fully commit themselves to touch-based interactions. Tesla has even gone as far as putting the drive selector on-screen (Hawkins, 2021). Interestingly, this drive selector does not violate American federal motor vehicle rules, according to the National Highway Traffic Safety Agency. This development seemingly clashes with other recent legal developments, such as the German court prohibiting touchscreen operations while driving (Stegmaier, 2020)

Still seeing value in physical interactions, some companies such as Volvo have opted for a mix between touch and physical controls. Their approach seems to be that all interactions can be done on the touchscreen, but the physical inputs are there to simplify certain actions. Such an approach seems to the safer bet legally speaking.

Application-based Infotainment Systems

With the development of systems like Android Auto, or Apple Carplay, car interfaces are starting to take an approach similar to that of smartphones. The screen becomes a blank canvas on which applications can paint whatever they desire. This approach has been proven to work for phones, ever since the initial iPhone release wiped out most of the use of input buttons, such as physical keyboards, almost in an instant. This does mean that developing input devices for the car becomes a little more challenging. With no strict rules on what an app can and can't do, it becomes difficult to map input knobs across applications. We can limit our input device to provide functions for a singular non-driving related activity to alleviate some of these compatibility issues. In our case, we could create a control knob that would only control music streaming services. This does not mean that compatibility issues don't still arise. If an input explicitly performs a specific action, this action must be available across music streaming services. Alternatively, inputs should be non-specific, with button mappings being presented on the screen. In short, input hardware can be made with a certain NDRA in mind, but should not be hyper-specific in its use case.

Social Music Listening

Music streaming services collect user data to power their recommendation engine. This vast amount of stored user data has opened the doors to some new social music listening possibilities. Tools such as *Spotify Only You* (Spotify, 2021b) use this huge amount of data to paint a picture of what makes you unique as a music listener. Other services, like *Spotify Blend* (Spotify, 2021b) allow users to create a shared playlist that reflects both users' music tastes, allowing for a way to easily bond over music. This concept of bonding over music by combining and mixing tastes has been physicalized by Frolic Studio (2012). In their project *'Skube',* musical tastes become physicalized in the form of a cube. Bringing cubes together creates a shared playlists that all can enjoy. Furthermore, by changing orientation of the cubes, users can toggle between 'playlist' and 'discovery' mode. The former being music already known to the user, and the latter being tracks similar to what the users listens, but hasn't heard yet. This results in a system that simplifies some otherwise quite difficult to achieve tasks through use of a tangible interface.

The amount of projects that try something similar to 'Skube', and apply it to the interior of a car, seems to be very limited. This is one of the design niches that this project therefore aims to fill. Social music listening aided by big data, applied to the context of the car. The other design niche being the usage of car sensor data to aid music recommendations. Both areas provide the opportunity for a real valuable and novel design contribution.

DESIGN PROCESS

PRODUCT DEVELOPMENT OVERVIEW

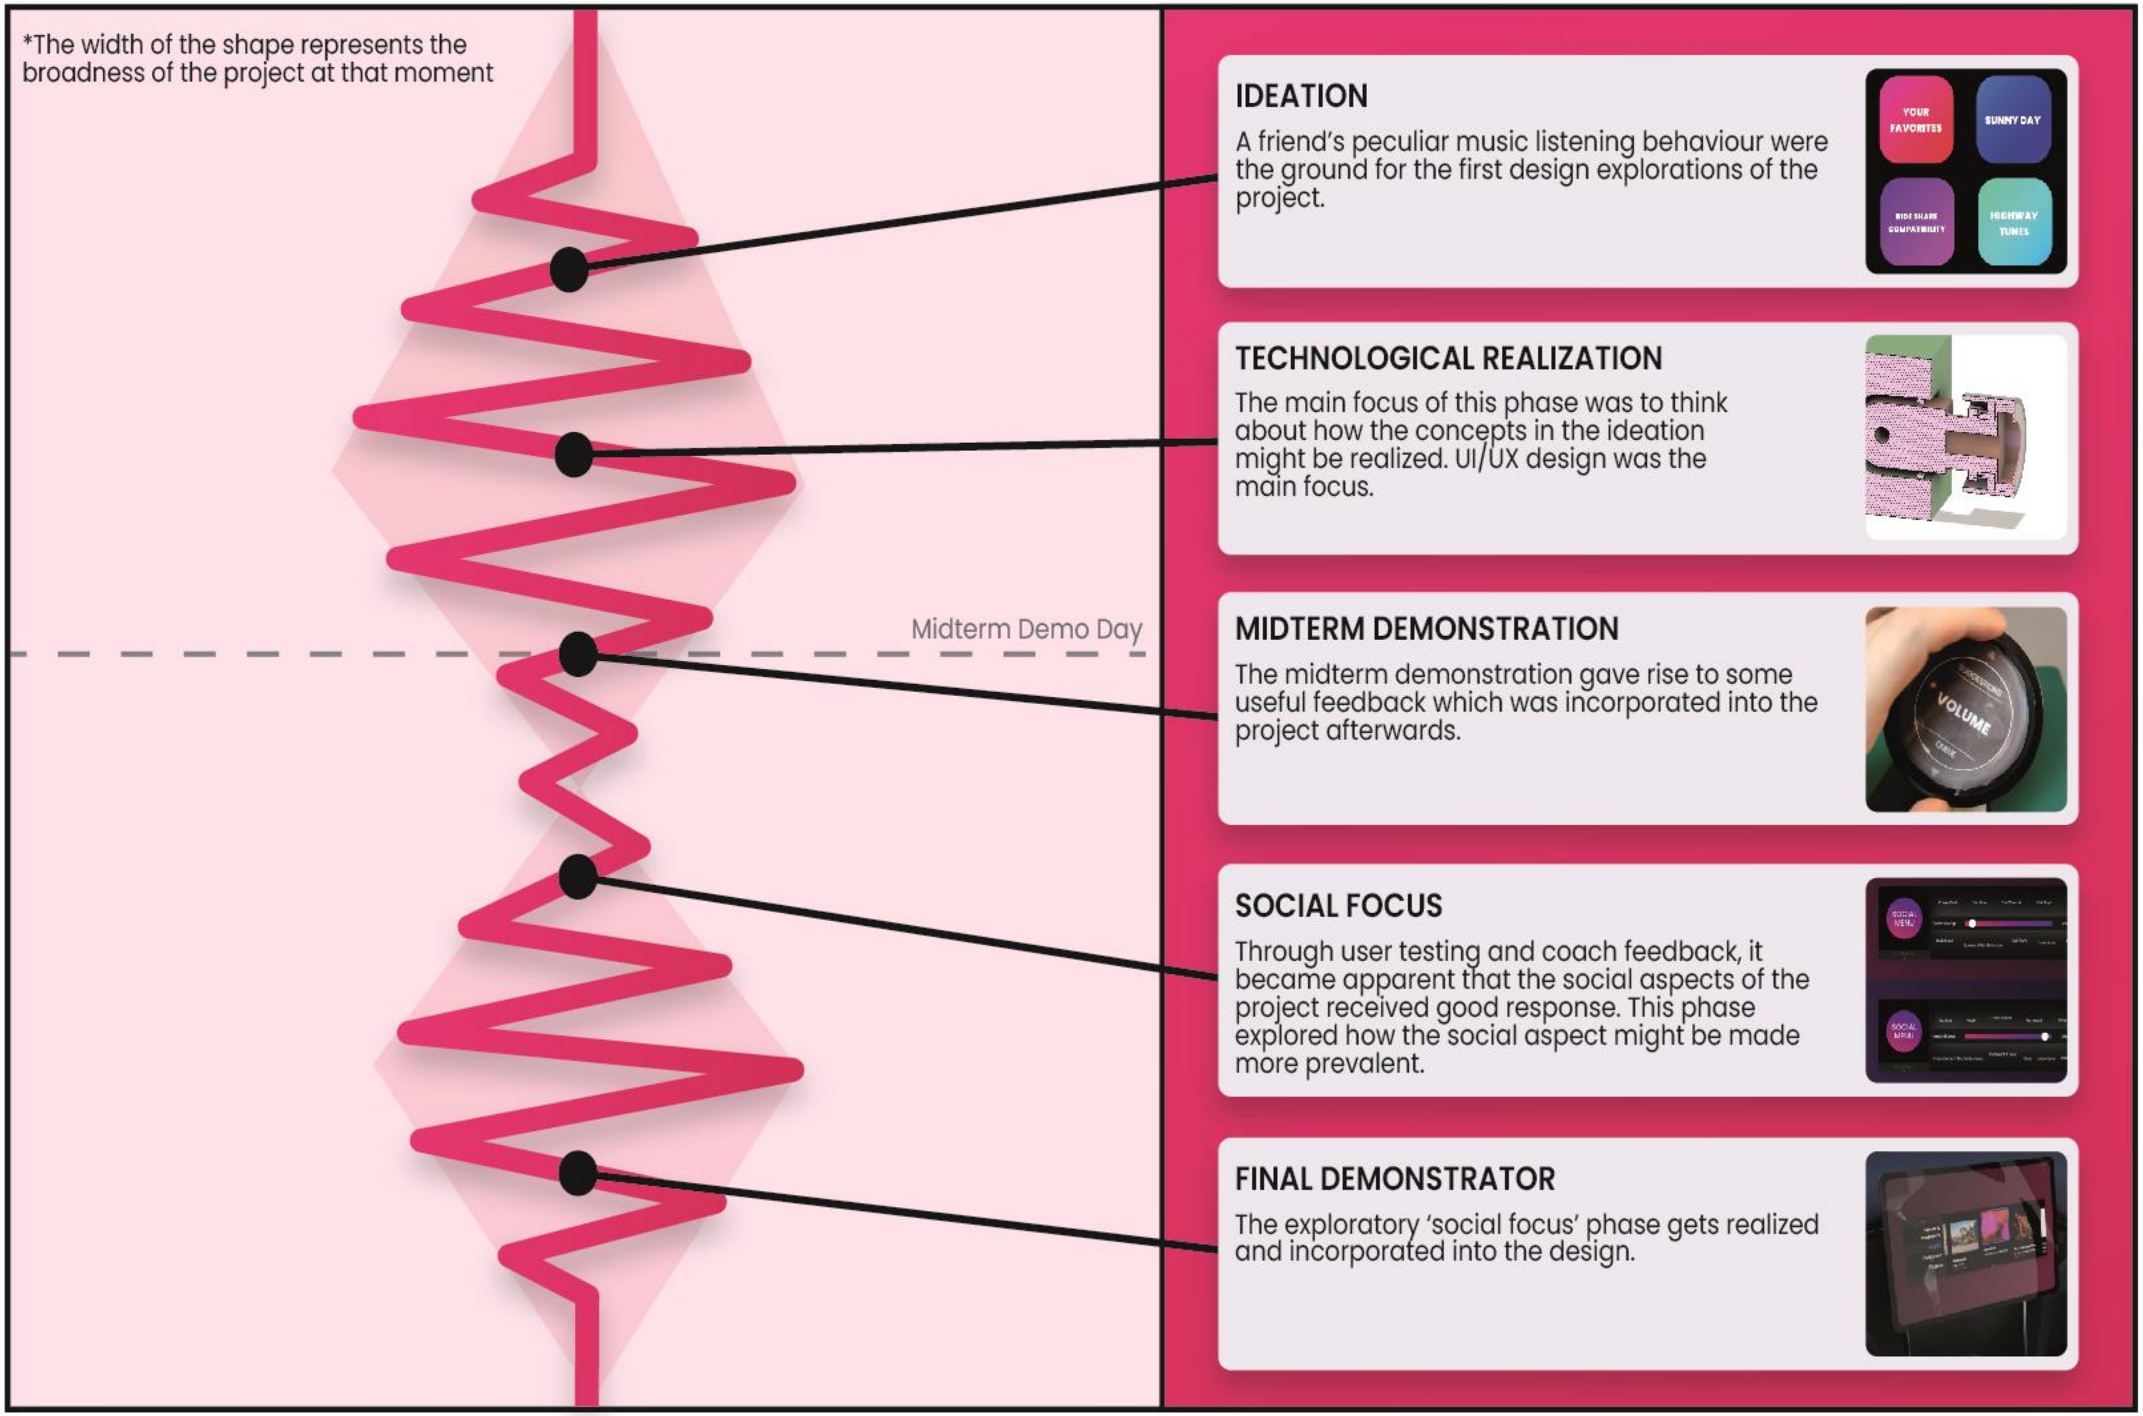

Figure 1: Design process overview

DESIGN PROCESS

The following section describes the pivotal moments in the AUD_I/O project. The different design phases are laid out, and the choices shaped the final design of the project are explained. Figure 1 gives an overview of the phases at large. The real timeline is not as black and white as depicted here, but the different phases discussed should give an insight into why certain major design choices were made.

Initial Ideation

The AUD_I/O project was born out of an interesting observation in the real world. A friend of the designer plays techno music exclusively on the highway. This quirky music listening habit sparked the initial explorations. Do more people listen to different music depending on the car's location? Are there other variables that influence the desired music? How can we utilize these usage quirks?

It quickly became obvious that these circumstance-based recommendations could largely rely on already existing sensors in the car. GPS, temperature, time, etc. are all already found in the car. Furthermore, an increasing amount of cars are connected to the internet. This would mean that additional data could be fetched from APIs for example. Thus, a system that, for example, plays techno when on the highway, was quickly deemed technologically feasible. After this realization, ideation moved towards how such a system might be implemented. This resulted in the first design iteration of AUD_I/O.

The first concept (figure 2) shows a barebones user interface. The goal of this interface was to provide an experience similar to the radio. A few options would be available and it's easy to switch between them. Similar to how a user might utilize radio stations (i.e. switch sometimes but generally set-and-forget). The playlists examples shown give a glimpse of what a sensor-powered recommendation engine might allow. Playlists are generated to be in line with the weather, location, and even one that takes the passenger into account. This passenger playlist served as the basis of what would later become the 'social suggestions' part of the project. General response to this 'journey-based music recommendation engine' was positive. This, combined with a lack of closely related works, resulted in this becoming the design nice that the project would explore.

This design niche was subsequently explored through mind mapping of the possibilities (figure 3). Through this exploration, it became apparent that the system would most likely consist of a 'smart' and a 'social' part. These aspects are further explored throughout the project.

Figure 2: First visualization of the concept

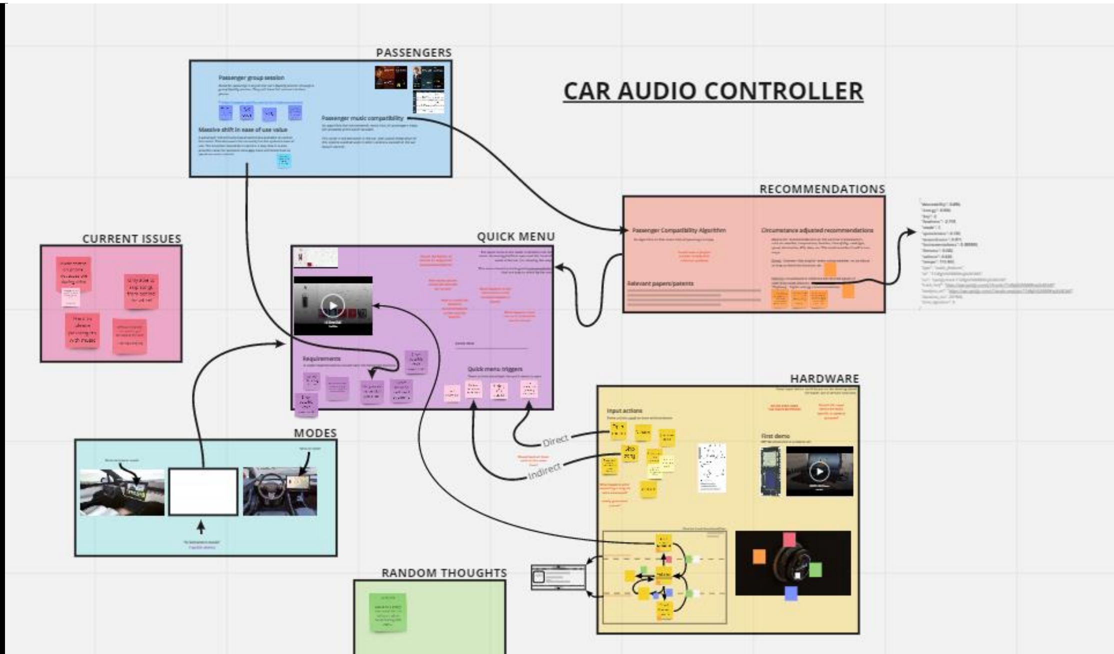

Figure 3: Mind map of the design space

Technological Realization

The aim of this phase was to conceptualize how the ideas generated during the ideation phase would be applied in the real world. The first phase presented more of a ‘set-and-forget’ approach to music control. This iteration tried to infuse the interface with more customizability. This choice was made after seeing the general usage patterns of streaming services in the car. A first iteration of an music streaming interface was created, that reflects the underlying novel design methods. Moreover, everything that users have come to expect from their music interaction devices, and more has been conceptualized. This includes skipping tracks, adjusting volume, stopping/resuming music. Lessons were learned about which parts of the design would be implementable within the project, and adjustments to the project’s focus were made accordingly.

Scoping the Design Domain

The automotive industry as a whole is moving towards a future of fully autonomous driving. Designing interactions for the car of today is very different from designing for a fully autonomous vehicle, or even a semi-autonomous vehicle. (Semi-)autonomous vehicles require different levels of attention from the driver, and this varying level of automation should ideally be reflected within the user interface (Pfleging, 2017). In this iteration of the design process, the choice was made to focus on current-day cars. This means that the systems are not designed with (semi-)autonomy in mind. By focusing on non-autonomous driving, the design needs to be usable in a system with a scarce amount of focus available. This results in a design that serves its purpose today, as well as during dangerous moments in a semi-autonomous vehicle when the level of automation is dropped. The user interface will most likely lose some of its purpose once the automotive industry adopts fully autonomous driving. The sensor augmented recommendations might still prove to be valuable in this future however. The convenience might outweigh the benefits from using the streaming service in full. To put it more concretely, users will most likely be able to use a full music streaming client in the future, comparable to a desktop experience. This would allow for more complex interactions, but it is unclear if that would outweigh the convenience benefits of the journey-based suggestions.

User Interface

Interface designs were made to see how the car sensor aided suggestions would possibly be represented in the infotainment system (figure 4). An iterative design process lead to the use of primarily horizontally scrolling menus. Other annoyances people had with their current car music solution were also addressed in the creation of this interface. One example would be the repeated skipping one has to do to get to a song they like.

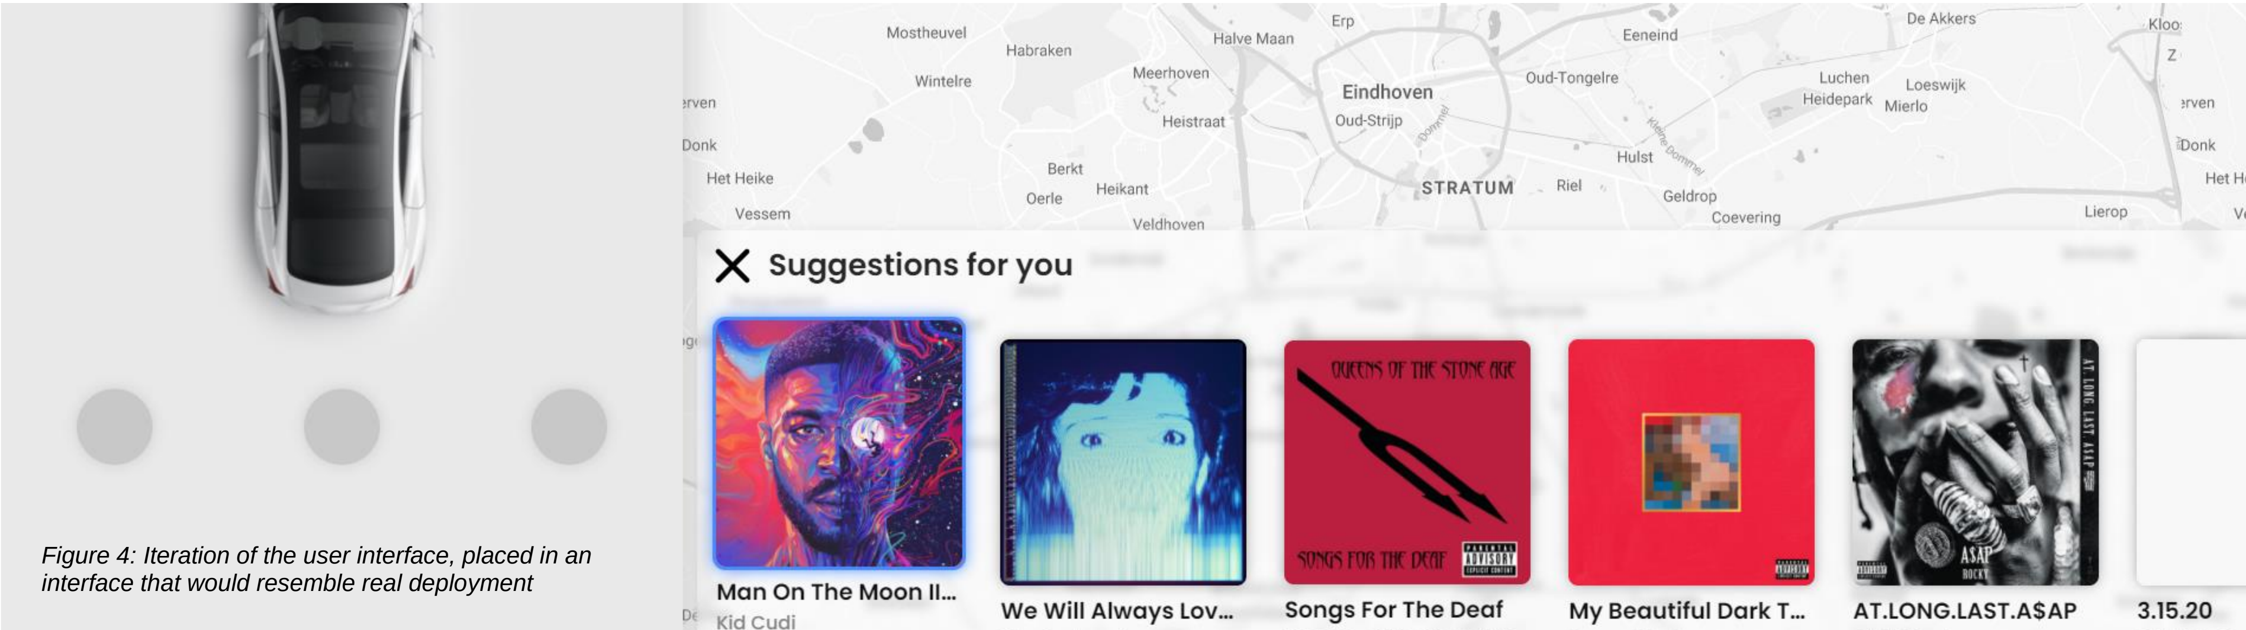

Figure 4: Iteration of the user interface, placed in an interface that would resemble real deployment

Input Device

User testing and research has lead to the creation of a physical input device (figure 5). Users seemed to respond well to car interfaces where they can navigate the user interface with a control knob. An additional benefit of designing with the physical knob in mind is that it forces a simplicity on the touchscreen interface. Interactions need to be achievable with limited inputs available. Confining the interactions to what can be done through physical input served as a good way to make the user interface more intelligible. A screen has been added to the control knob as well. Monitoring user behavior in other cars showed that the 'music menu' sometimes can get hidden when other interfaces take priority, such as navigation. By adding a small screen on the knob, most actions can still be performed, even in this minimized view. One could imagine the interaction to be similar to that of a smartwatch.

Realizing the Backend

This design phase quickly made it obvious that implementing the smart suggestions would not be feasible to a level comparable to the final product. One reason being time constraints, but more importantly, the smart suggestions concept is driven by big data. To gather enough data to power the machine learning part sufficiently was deemed unachievable within the scope of this project. Therefore, the focus shifted more towards fully realizing the social aspects of the project. To realize this social aspect, a compatibility algorithm was implemented. This algorithm allows users to combine their music tastes. This pivot towards realizing the social aspect to the fullest proved to be a good choice. Later testing shows that users more easily can get a sense of how the social system would be beneficial to them, as compared to the smart suggestions. This might be due to the promise of 'smarter suggestions' being hard to grasp. The enjoyment of the smart suggestions is directly linked to how well it is implemented. More concretely, users need to see how much better the smart suggestions actually would be. Conversely, with the social suggestions, the users could quickly fathom the concept and would even see its usefulness in this demonstrator phase. From a user-centered perspective, it also makes more sense to prioritize the social aspect in the testing, since the user would be more involved with this part of the design.

Voice Commands & Car Thing

While the smart suggestions can cover a big part of the user's interaction needs, they might desire more specific control at certain times. After a period of conceptualization, it was concluded that voice control would be the best option for these kinds of specific requests. Interestingly, just after conceptualizing these interactions, Spotify released Car Thing (Spotify, 2021a), a car audio player that implements these interactions in a very similar way. From this point onward, focus in the project shifted away human-computer interaction. This was done since the basic interactions were adequately mapped at this point, and to focus on providing novelty through the system.

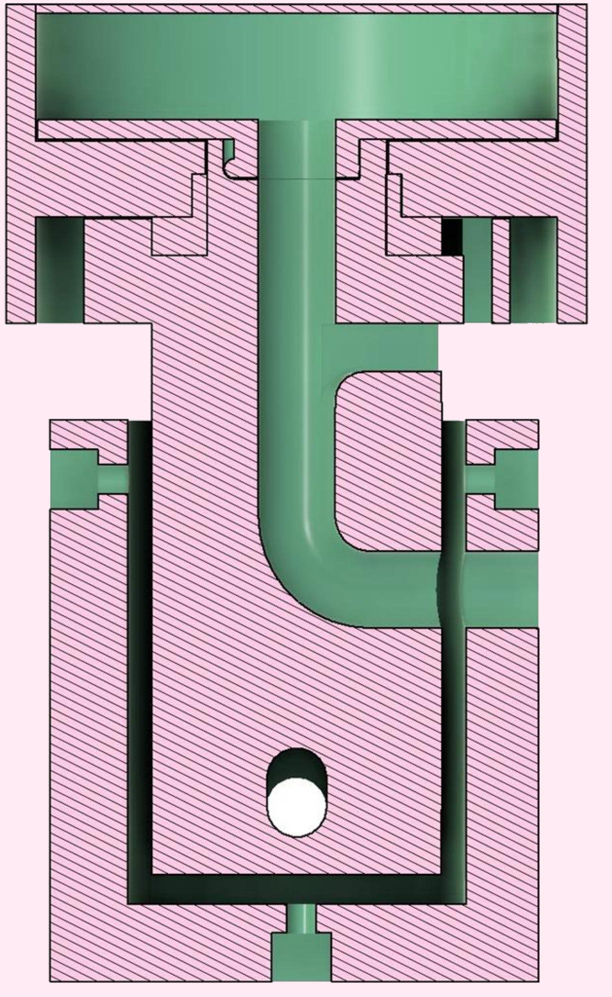

Figure 5: Cross-section of the 3D-print

Main iteration takeaways

Designing the user interface has brought to light some new and interesting design challenges and opportunities. The user interface can be aided by a physical input device. Further input improvements can be made through the inclusion of auditory inputs for some complex requests. Requests such as 'play my favorite hip-hop songs released in the 90s' are relatively easy to request through a voice command and are much harder to do through touch-based interactions (if not impossible). Exploration of what can be realized on the recommendation engine side resulted in a shift towards the social aspect. This part of the project was easier to fathom for users, even when implemented in a non-final form. The enjoyment of the smart suggestions is directly linked to how well it is implemented. More concretely, users need to see how much better the smart suggestions actually would be. It was for this reason that ultimately a more superficial implementation was chosen for the 'smart suggestions'.

The Midterm

The midterm demonstration proved to be a useful moment for feedback. Furthermore, during this time, the commercial side of the project was taken into consideration. The following sections describe these explorations in more detail.

Demo Feedback

One piece of feedback that was prevalent during the demo was the absence of the news. People seemed to like the fact that the radio periodically keeps them informed, and they miss this when using a streaming service.

An way of interacting with the news was conceptualized (figure 6). Since all systems are digital, we can record and playback news updates at any time. This realization served as the basis for this novel design. Users can go over to the news (& podcasts) section and listen to their favorite news update. These updates are periodically refreshed to be the latest version available. Some user testing was done on whether the system should still automatically play the news at certain times (appendix A), but most people seemed to agree that they would manually start it themselves when desired. One possible midway solution would be to implement popups when the system strongly believes that the user wants to listen to a news broadcast. An added benefit of this digital news listening is that users can now skip through sections that don't interest them. The 'queue' section gets populated with news topics, giving users the ability to easily skip between them.

The news section was conceptualized, but did not become a very prevalent part of the project, since it is not in line with the main design goal: creating a symbiotic relationship between the car and the music streaming service.

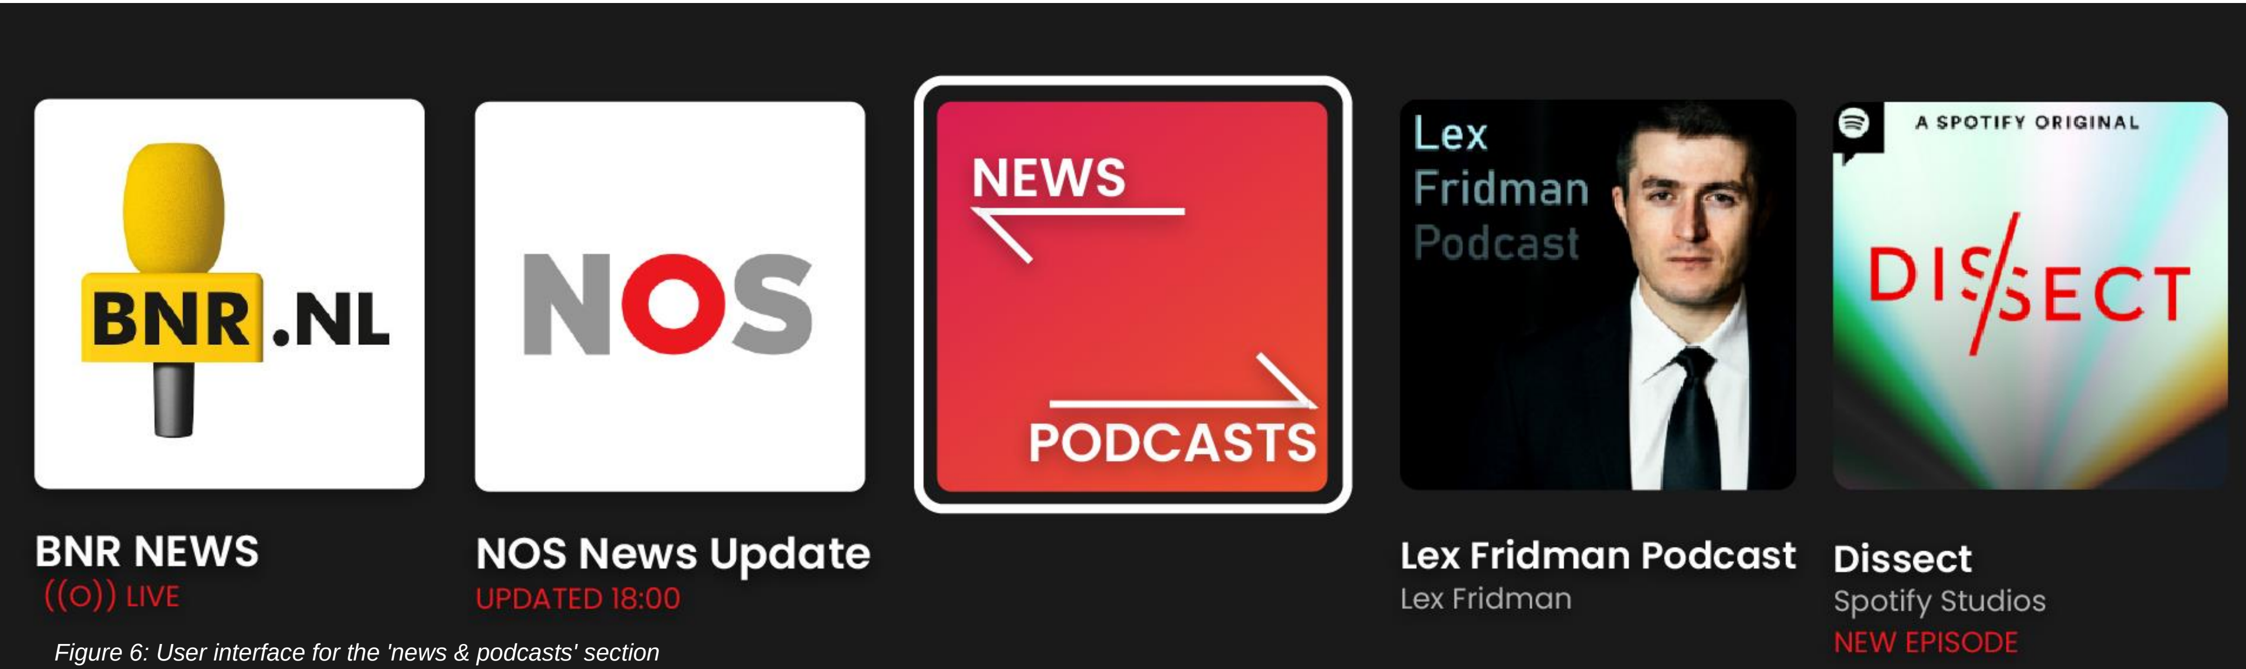

Figure 6: User interface for the 'news & podcasts' section

Business opportunities

Around the time of the midterm, the commercial value proposition of the project was also being developed. Multiple tools were used, such as value proposition canvasses, user journey maps, personas and more (figure 7). This commercial exploration showed quite promising results from the user as well as the service provider side.

Companies can provide quite significant value to the end user if they were to implement the AUD_I/O system. Car music control would become more convenient, since the system provides better suggestions than the status quo. This alone would provide enough incentives for some users to want and use this new system. The social aspects are entirely novel compared to systems that one can find in current day cars. This provides additional value for the end user. Importantly, both these upgrades come without significant downsides in other areas. There are barely any trade-offs identified when upgrading to this new system. One possible downside might be possible privacy concerns, but those could in theory be alleviated through certain privacy preserving computation methods. Companies are moving towards local computation of sensitive data. Voice assistants such as Apple's Siri are now using on-device speech recognition. This means that the sensitive data never leaves the device. Similar technology could in theory be used to for example, keep location and speed data in the car. It is however important to note that such measures might not have an effect on users that are not tech-savvy. They might not understand these systems, and therefore also not trust them. Manual control of what data is used could also be given to users, such that they can choose which data gets used or shared with third parties.

As for how the business side would look, there are multiple opportunities on how this could take shape. The system could be integrated on the car manufacturer level (e.g. Tesla), or on the streaming service side (e.g. Spotify). All options have certain advantages, such as the car manufacturer being able to provide tighter integration with the car's interfaces, while the streaming service can provide a better integration with the recommendation engine. The streaming service provider would hypothetically also see more adoption. More people might switch music streaming services compared to buying a car to make use of the service. One could argue that companies like Tesla aim to provide a superior software experience in every way. In that sense this system would add to that experience. Even if this were to be the case, an argument could be made that the streaming service can still provide the service, and license it through an exclusivity deal. Thus, there seem to be more options when implementing from the streaming service side. They have the required expertise, and their potential market size is bigger. They can onboard car manufacturers, which could create a ripple effect where no manufacturer wants to stay behind.

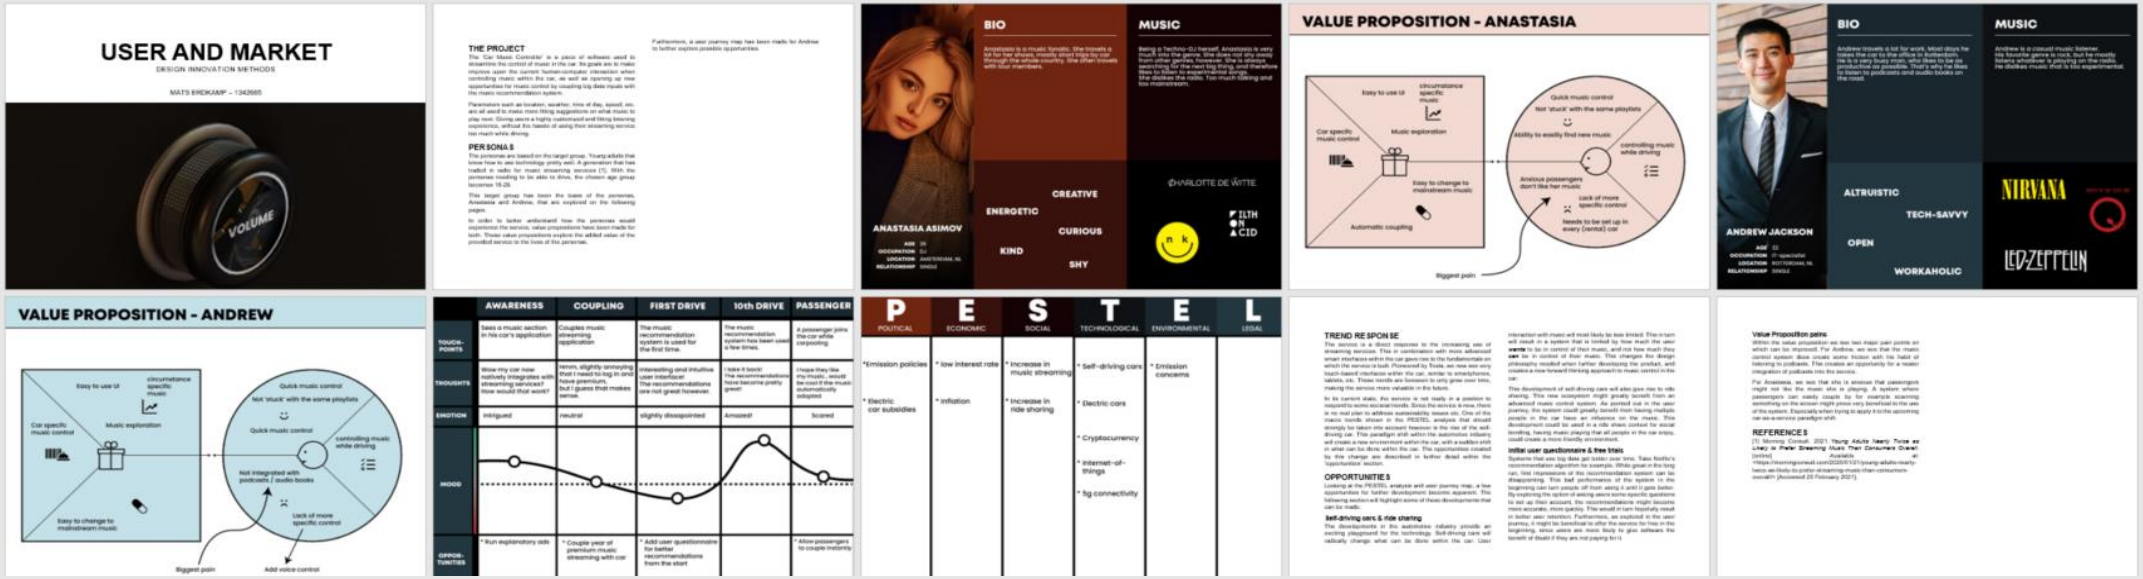

Figure 7: Partial overview of the business analysis

Social Focus

While the technical realization phase already discusses a shift towards the social aspects of the design, this phase takes that even further. Through some coach feedback and user testing, it became apparent that it might be beneficial to the user experience if users are allowed to have more direct interaction with social systems. This lead to a new ideation phase which broadened the design scope for the first time since the project's inception. Some of the new systems that were developed in this phase are highlighted in the following sections.

Social Slider

Originally, the idea for the social listening was to take the taste of both users and combine them in an even manner. After some feedback, it became clear that there might be merit to a more involved social system, one where users can for example negotiate whose music taste should be more prevalent. This resulted in the conceptualization of the social slider. This menu appears once a passenger has coupled to the car. Here, users can control how much their music taste should be represented in the social suggestions. Furthermore, users are allowed to blacklist certain artists that they absolutely don't want to hear. The 'design' section of the report gives an explanation on how this technology works in the product.

Inter-car connectivity

The development of the social systems in the car raised an interesting new topic; why not allow for social interactions between cars? This question served as the start of the inter-car communications systems. A few concepts were created that explore the possible opportunities that arise when assuming that all cars' music sessions are coupled.

The first idea that came to mind was inter-car song sharing. This system would allow users to share a song, when the system deems it safe to do so. A specific instance of this song sharing would during the downtime when waiting at the traffic light. Passengers can share their currently playing song to a neighboring car. It is limited to being the currently playing song to keep cognitive load low and to reduce spamming (i.e. sending generally unpleasant songs).

Second, the concept of a regional session was created. This would mean that all cars in a certain region could join this session, and listen to the songs playing. One area where this might be especially useful is during a traffic jam. The system would detect the traffic jam and would subsequently create a playlist where cars can freely join. This could in theory create a communal feeling in the traffic jam, since it would (according to user test participants) be fun to see everyone sing/jam to the same songs. This specific application of regional playlists got given the moniker 'Traffic *Jams*'. With 'Jams' also being slang for music. (i.e. 'this is my jam!').

User Testing

All of these social additions were user tested afterwards. Users were shown a picture and informed of how the system would function in practice (appendix E). Qualitative questions were asked about to get a sense of the usefulness of the conceptualized systems. The general response to all proposed systems was positive. Some systems were generally deemed more of a gimmick than others. But, the results of the user testing definitely made it seem like there is a lot room for exploration of the social aspects. This has therefore become a topic to definitely further explore in future works. The slider had the most positive response overall, and was therefore chosen as the flagship social feature of the project. As a way to communicate the inter-car communication capabilities, the stoplight song sharing was chosen.

THE DESIGN

THE DESIGN

The AUD_I/O project has been created with the developments of the automotive and music industries in mind. The following section gives an overview of the novel interactions and technologies that have arisen by focusing on the latest and greatest technologies in the respective industries.

The Overview

The design aims to make the music listening experience in the car smarter and more social. The 'smarter' aspect of the design mostly comes from the integration of car sensor data into the music recommendation engine. By using data such as the location, speed, ETA, weather, and more, the music recommendation engine can be augmented to appropriately respond to different situations within the journey. These journey-based recommendations allow for a better user experience, since the music recommendation engine now has extra parameters to strengthen its presented suggestions. The fact that the recommendation engine gets augmented by these extra parameters is not obfuscated from the user. For example, playlists might give indications as to why they are currently being recommended. An example would be the playlists clearly showing users that it has been made with the 'highway' and 'sunny weather' in mind. This ultimately results in an experience that is, but maybe more importantly, feels more personalized.

The second foundational pillar of the AUD_I/O project is the social controls. The AUD_I/O project allows users to easily couple to the car's music listening session. Driver and passenger can subsequently negotiate whose music taste should be more prevalent during the drive through interaction with a slider. This slider impacts the influence each passenger has on the recommendation engine. More information on the inner workings of this algorithm can be found in the 'social controls' section. The social controls are beneficial to the user experience in two major ways. First off, the fact that all passengers can couple to the streaming session results in music that all can enjoy. This results in greater overall enjoyment of the music that is playing. Secondly, the fact that the recommended music is a product of the musical tastes of all passengers, allows for discussion and social bonding over the recommendations. "You like this band? I love them too!" or similar conversations might be sparked by the AUD_I/O interface.

Smart Suggestions

The main emergent opportunity of tightly integrating the car with music streaming are the smart suggestions. By connecting car sensor data to the music recommendation engine, more fitting suggestions can be generated for the end user. Currently, streaming services add various metadata to songs. This metadata provides an abstracted representation of the song, which in turn can be used to inform the recommendation engine. In the case of Spotify, this metadata are called the audio features (Spotify, n.d.-b). 'Energy', 'valence' and 'danceability' are examples of the available audio features. All of these features are given a value between 0 and 1. Car sensor data allows us to cleverly choose in what range we want the audio features of the suggested songs to be. We might imagine that the user's desired song energy is linked to the speed at which they are driving. Over the medium-to-long-term, patterns might start to form in user streaming habits. Jazz might always get skipped on the highway, or sad music might be preferred when it rains. Through machine learning, the system can start to detect these user interaction patterns. This creates a self-learning system, where every newly played or skipped song informs the system of the user's desires at given circumstances.

Such a system could decipher usage patterns at the individual scale, but also at a macro scale. There might be general usage patterns shared by a great share of users. These general usage patterns can be used to create global 'rules' (i.e. sad music when it rains). The benefit of these global rules is that they can alleviate some of the cold-start problems a new user might have. To put it more concretely, they can help power the smart suggestions of a user that has not had enough interactions with the system for it to give reliable suggestions. Thus, a new user might start with smart suggestion rules that are deduced from the macro trends. As time passes, the influence of personal usage patterns can increase.

Technical Implementation Synopsis

The AUD_I/O interactive demonstrator uses a superficial implementation of the system previously described. This system creates two personalized playlists for the user. 'Sunday Drive' and 'Highway Tunes'. The generation process of these playlists is as follows. First, the users fifty most listened songs of all time are requested from the Spotify API (Spotify, n.d.-c), after which the audio features of these songs are requested (Spotify, n.d.-b). Spotify defines energy as follows: *'Energy is a measure from 0.0 to 1.0 and represents a perceptual measure of intensity and activity. Typically, energetic tracks feel fast, loud, and noisy. For example, death metal has high energy, while a Bach prelude scores low on the scale. Perceptual features contributing to this attribute include dynamic range, perceived loudness, timbre, onset rate, and general entropy.'*. The top fifteen highest energy songs are allocated to the 'Highway Tunes' playlist, while the lowest energy songs are put in the 'Sunday Drive' playlist. The result is two playlists where the user can get a general impression of how the two situations might result in different recommendations. It is important to note that the link between highway and higher energy music has been chosen by the designer. Real implementations would not make assumptions like these, but rather recognize usage patterns in real-world deployment, as described in the previous section.

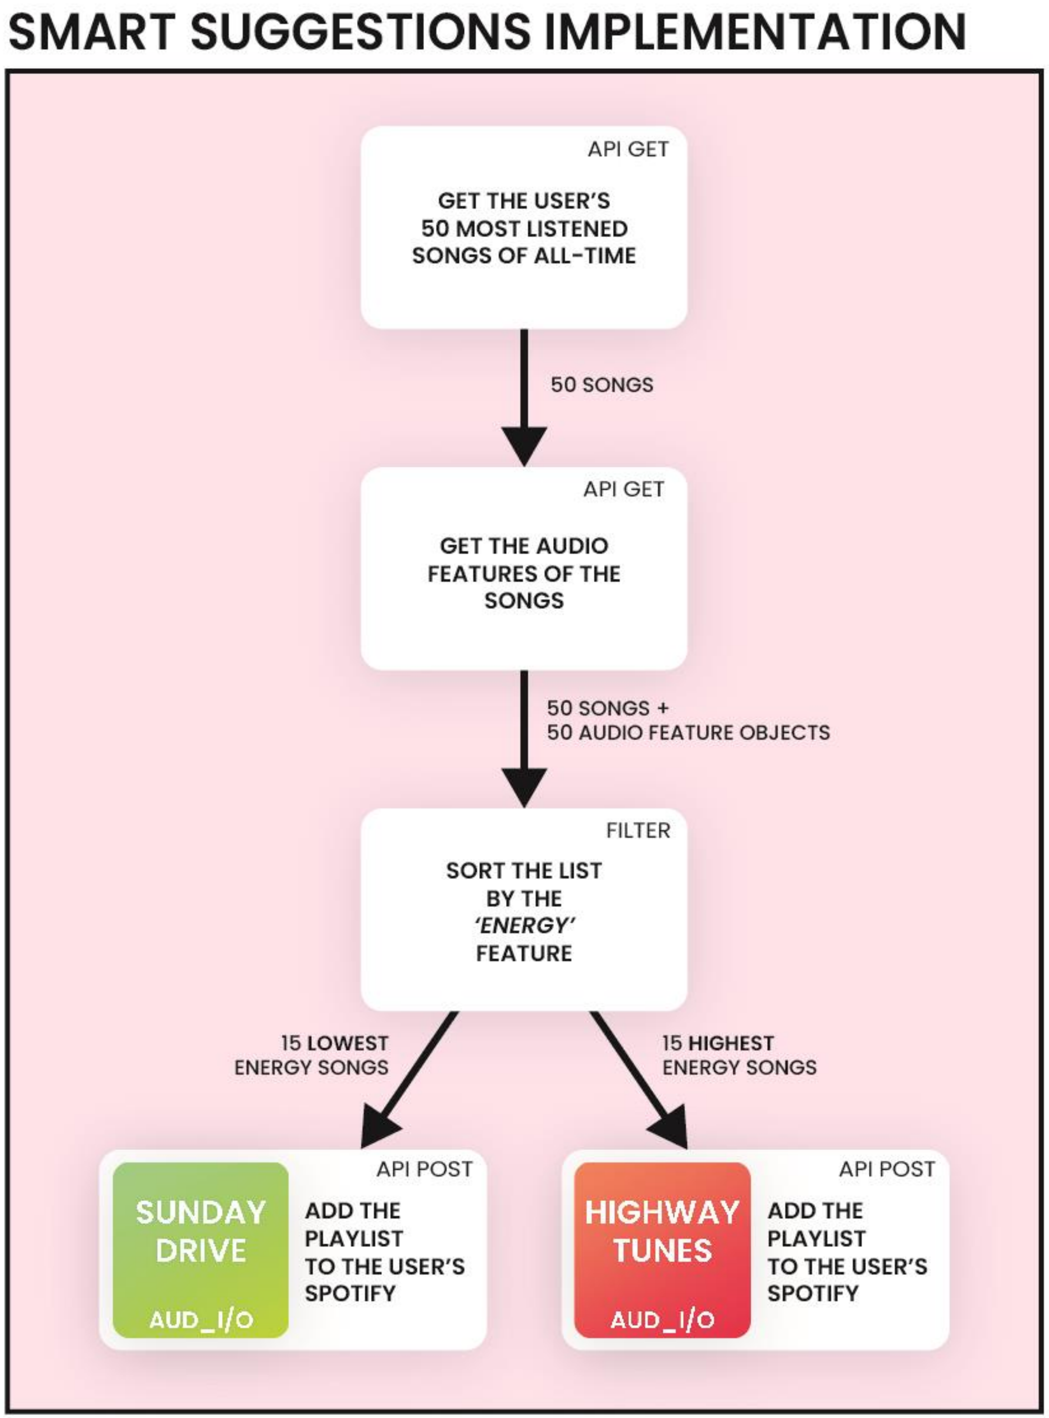

Figure 8: Schematic overview of how the smart playlists are generated in the demonstrator

Social Suggestions

The other key part of the AUD_I/O project is the social controls. This part of the project focuses on the inherently social nature of music listening, and allows for some novel interaction methods to aid this inherently social experience. First off, passengers can easily join a car's music session, which in turn lets them influence the smart suggestions. Their influence on the system can be negotiated through the use of the social slider (figure 9). This slider gives users an intuitive sense of how much of their music will be represented, and creates the opportunity for some social negotiation. Ten artists are shown to give an impression of the recommendations that will be generated. Other social interactions that become available through a focus on music streaming, are a karaoke mode ('Caraoke') and Inter-car music sharing (figure 10). This latter concept has been further explored in design context of a traffic jam, and goes under the moniker of 'Traffic Jams' (figure 11). Traffic Jams creates a regional playlist which all cars in the traffic jam can join freely. This in turn would create a more communal and fun feeling during the traffic jam. User testing revealed that participants saw promise in the idea of seeing neighboring cars sing along to the same songs.

Technical Implementation Synopsis

To allow demo users to experience the social suggestions, a music compatibility algorithm has been implemented. Adding up all the streams in the users' streaming history was deemed insufficient since it comes with two major flaws. Firstly, the user with the most streams would shift the recommendations in their favor. Secondly, just summing streams does not inform us to what degree the songs are actually compatible between users. A novel algorithm has been implemented that aims to minimize both these problems. This algorithm uses the median rank of a given number of streams. An explanation with pseudocode is depicted in the figure 12. The resulting total 'compatibility' is the sum of all the individual scores. This summation process can be altered through the addition of an influence multiplier. This multiplier can let one person's scores be more influential. This is how the social slider functions. The underlying formula for figure 8 would approximately be:

$$TotalScore = Score(famerdkamp) * 0.9 + Score(pieter2509) * 0.1$$

From this formula we can deduce that it is more important in this scenario that 'famerdkamp' likes a song. The same formulas can be implemented to work for albums and artists as well. The album version is implemented in the demonstrator. As a result, users can quickly see and verify the influence of the social slider on the recommended albums (figure 13).

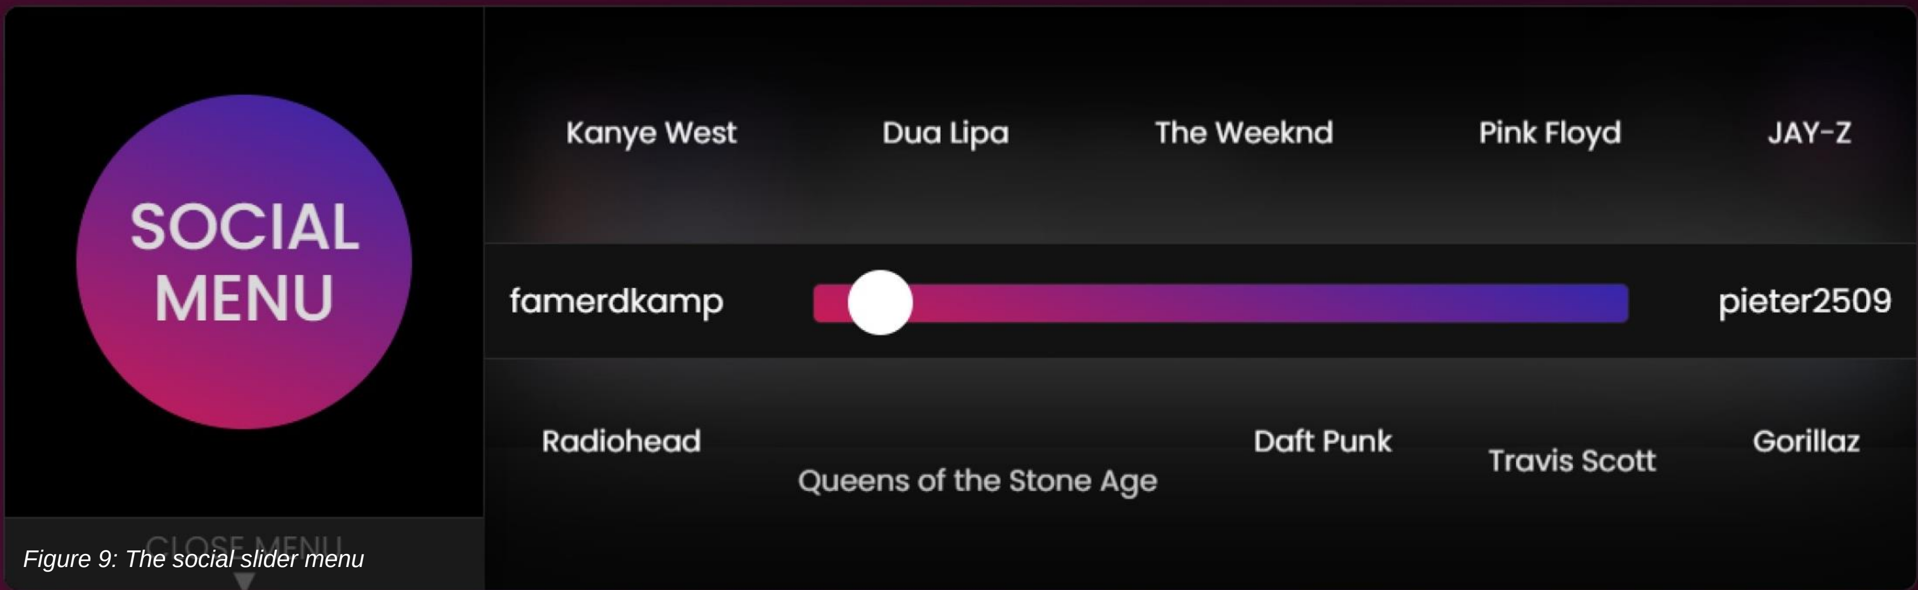

Figure 9: The social slider menu

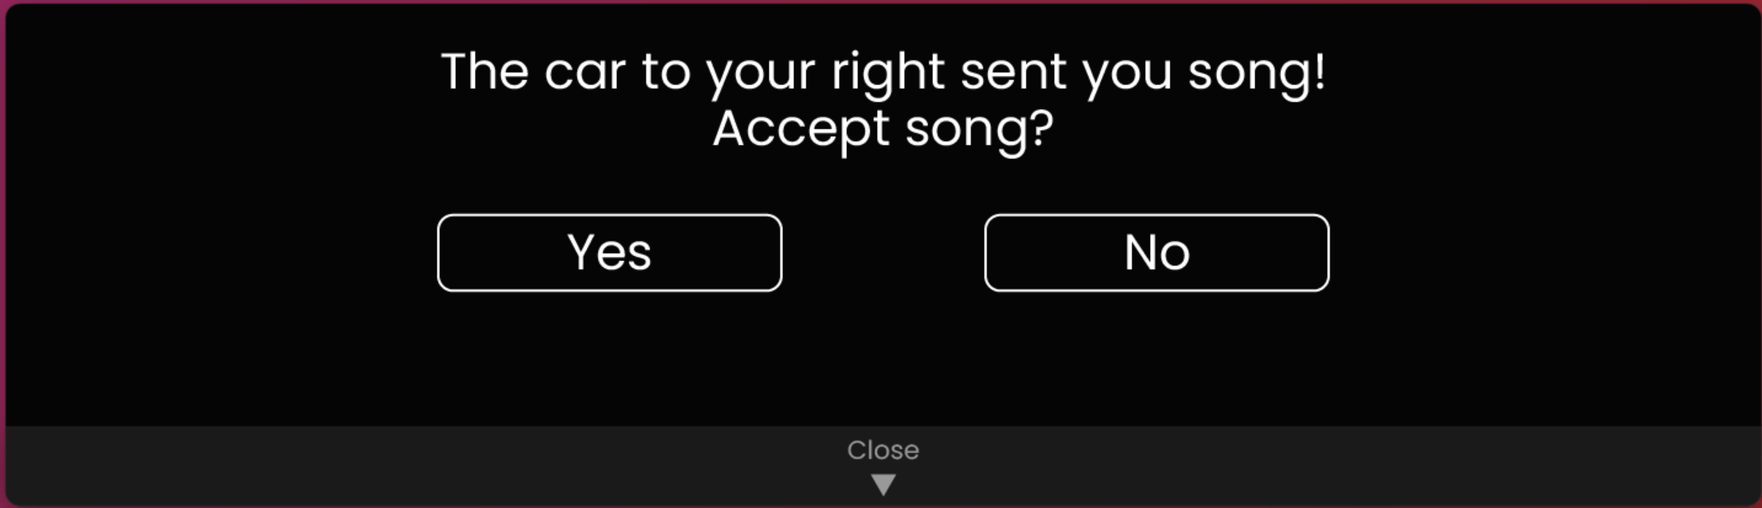

Figure 10: Example of how inter-car song sharing would look on the receiving end

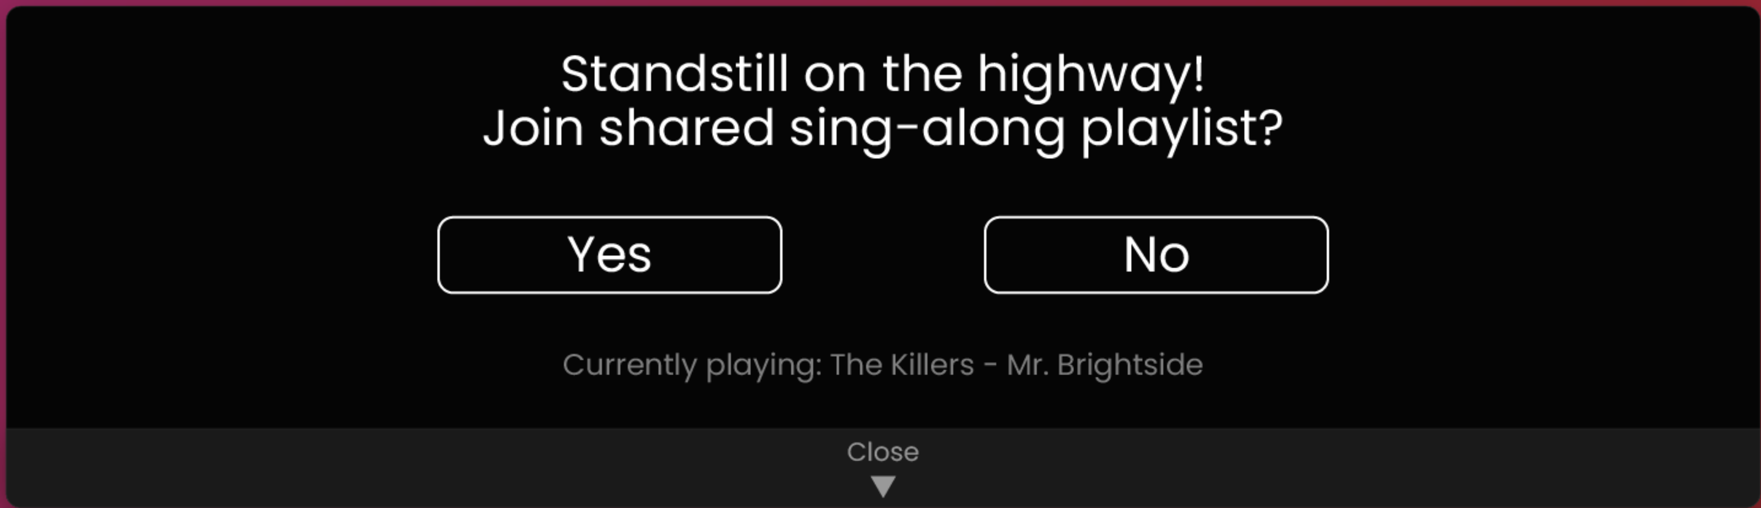

Figure 11: Example of how inter-car regional playlist system could be applied

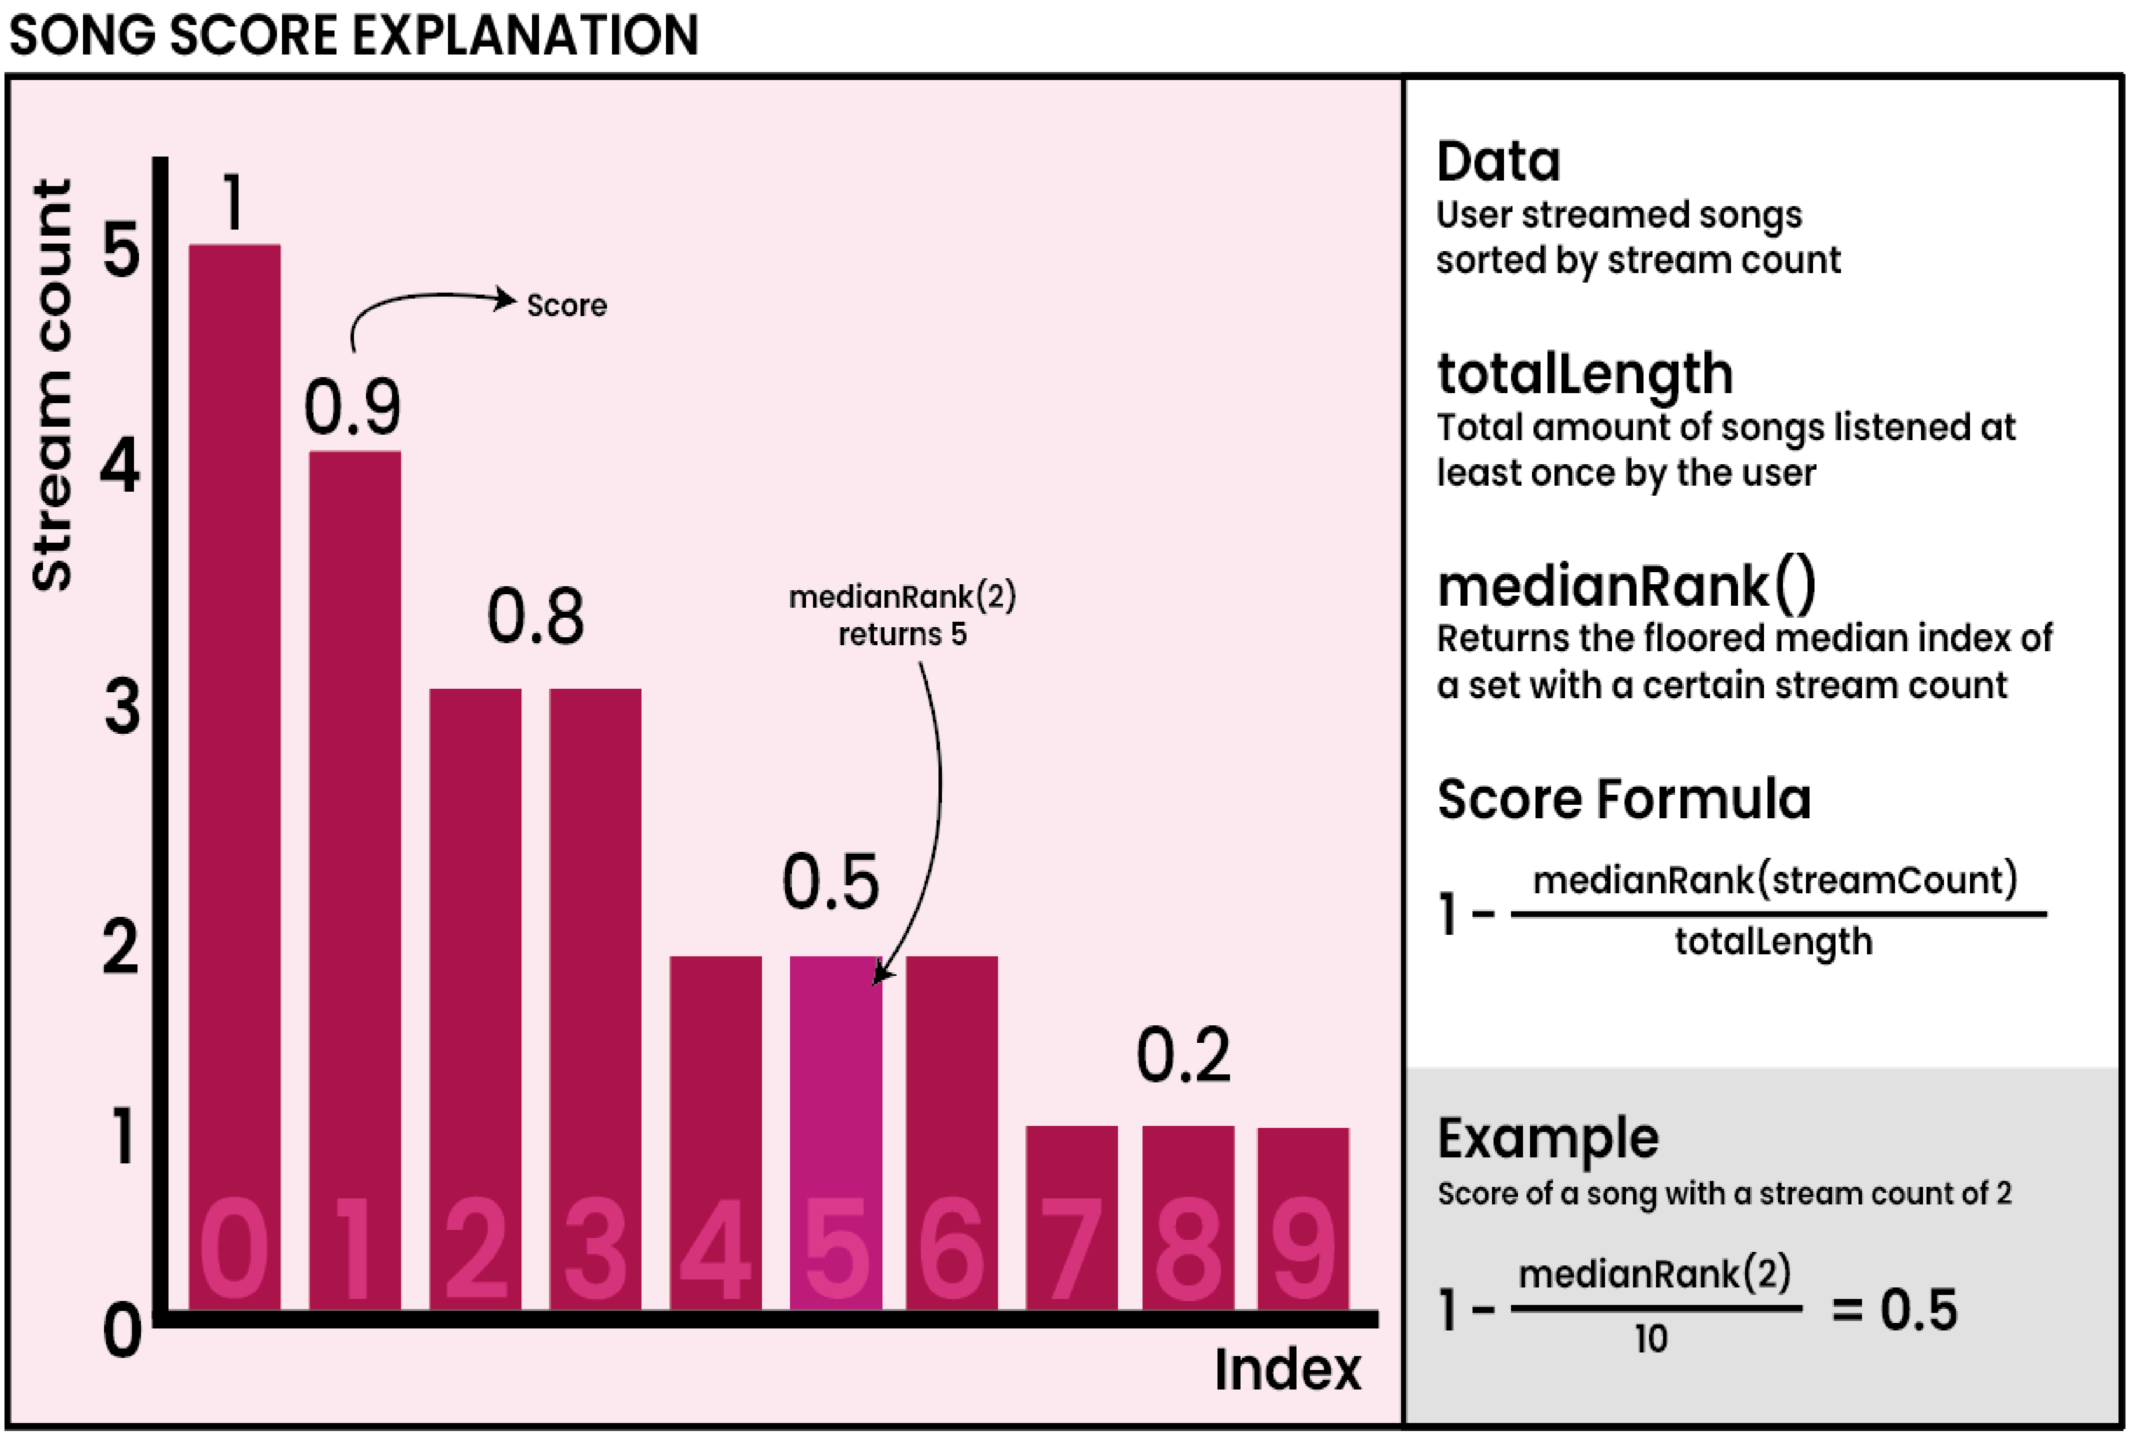

Figure 12: Explanation of how song scores are calculated

The Interface

In an effort to bring the concepts presented closer to reality, an interface was created to show off how the smart/social features might be implemented in a real setting (figure 13). The interface encapsulates all the novel interactions made possible by the AUD_I/O project, and aims to present them to the user in a way that would be similar to an actual implementation in the car.

The interaction flow of the menu has been optimized for easy and safe navigation. A general overview of how the menus are connected can be found in figure 14. Suggestions can be seen at a glance, and previously tedious tasks, such as skipping multiple songs, are made easier. Interestingly, after development of the interactions, Spotify themselves brought 'Car Thing' out into the world. Car Thing is a standalone car music that is clipped on top of the dashboard. The interactions in the Car Thing project are very similar to what was conceptualized for the AUD_I/O project.

It was therefore chosen to not further validate the interface design choices made, but instead focusing on the novelty that the AUD_I/O project brings: smart and social music suggestions.

The Demo

The demonstrator interface is available for the public to explore (appendix C). This demonstrator is the same as the one presented during the demonstration day at the end of the project. Users can log in with their Spotify account to try the demo. Furthermore, a ninety second demo video was created to communicate how the interface would function to the public (appendix A). Additionally, to see the system in uninterrupted use, a short usage video was recorded. (appendix B).

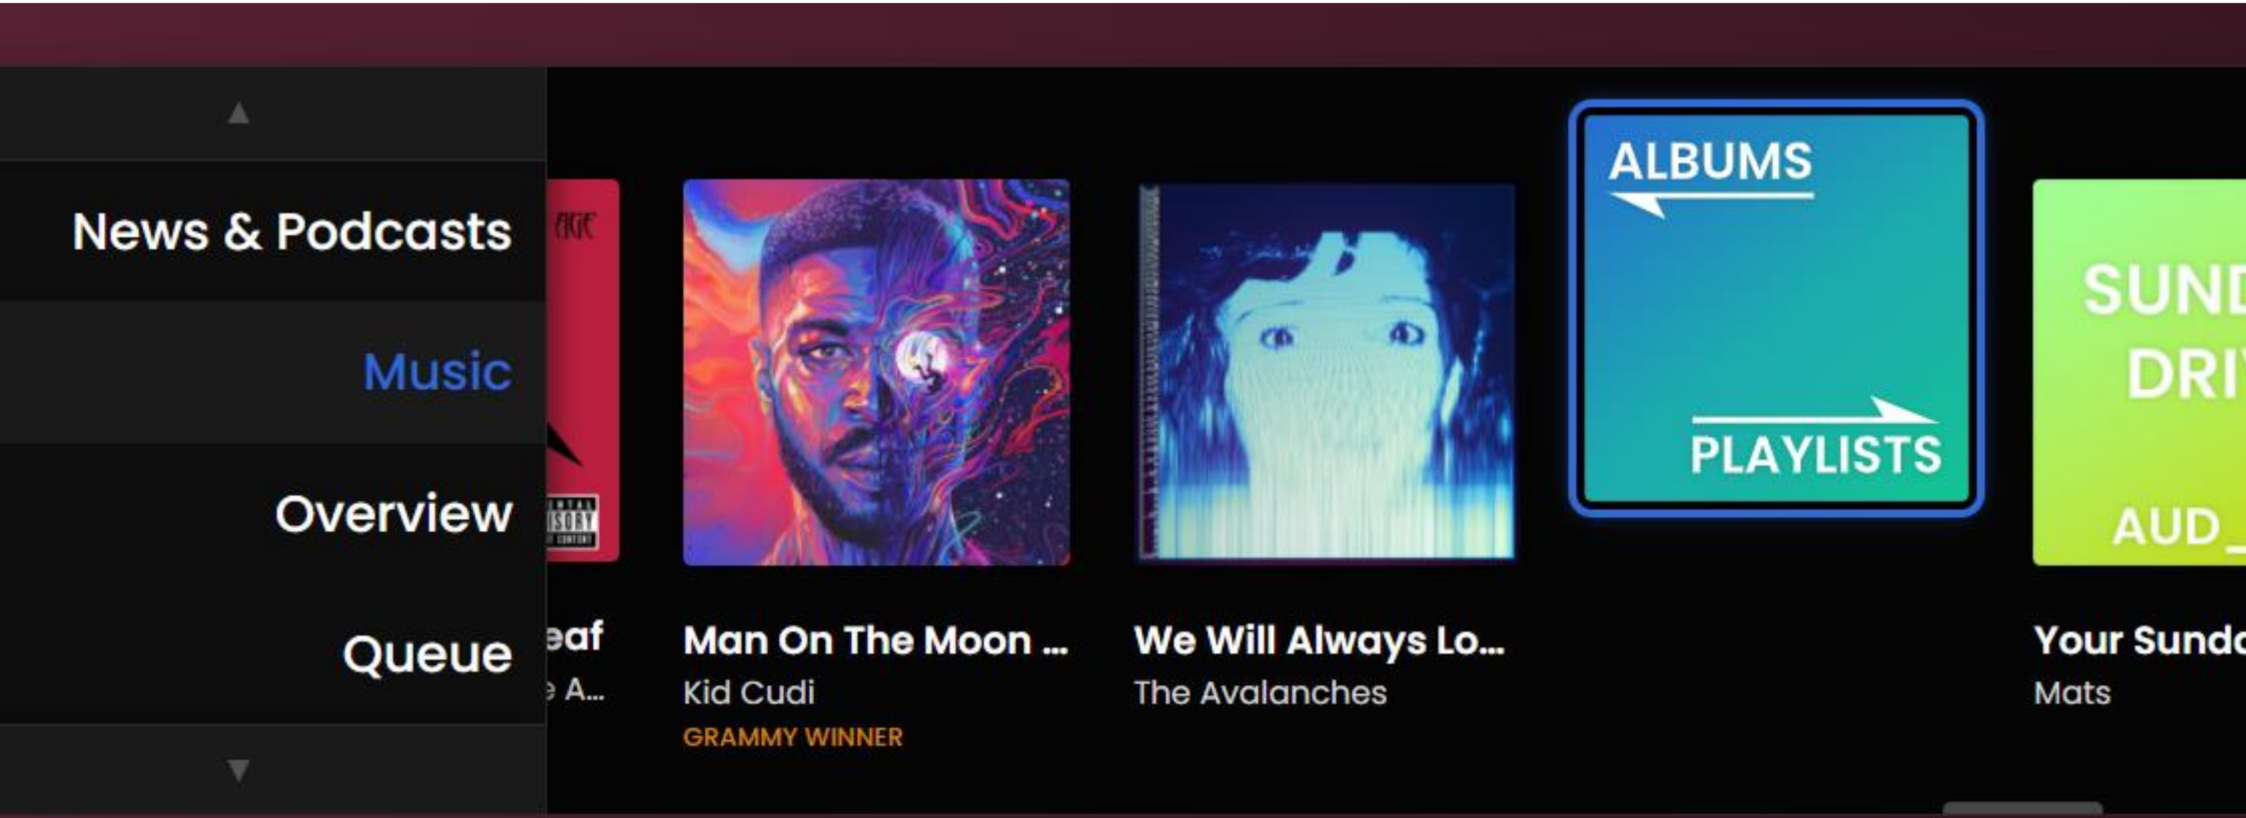

Figure 13: photo of the digital interface on the ''Music' tab

INTERFACE LAYOUT

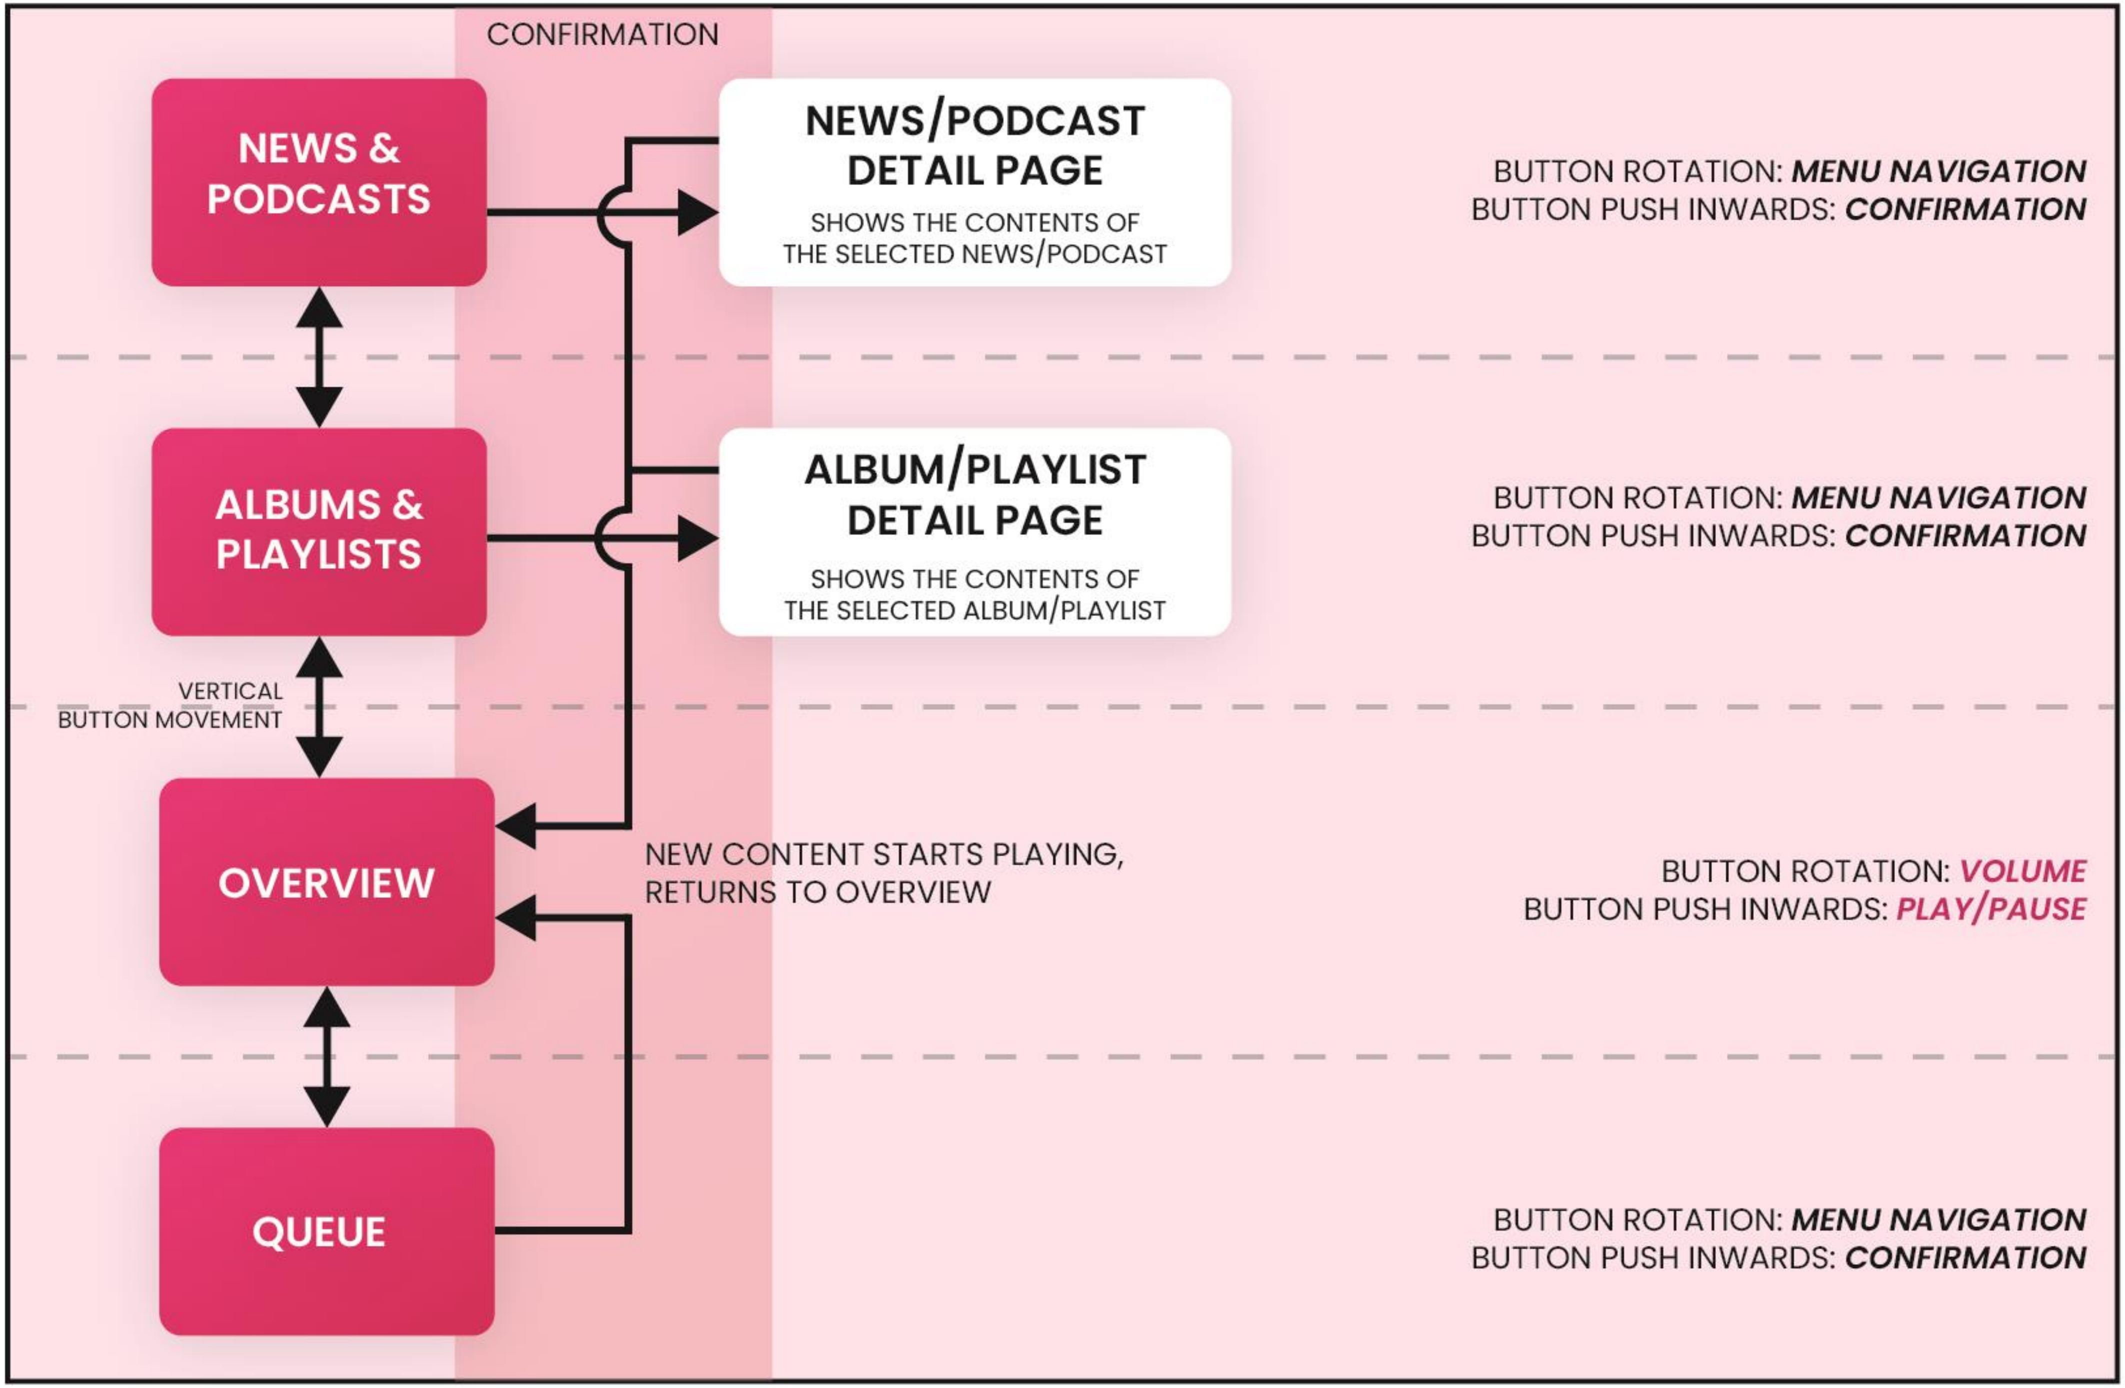

Figure 14: General overview of the menus and how they are connected to each other

The Input Device

A control knob was created to aid the navigation of the on-screen interface (figure 15). While not necessarily novel in its execution, its creation did result in certain benefits for the project as a whole.

First off, initial user testing showed that users like this kind of physical interaction with their infotainment system within the car. This belief is further cemented by the plethora of car manufacturers that make use of similar physical input devices.

Because of this input device, the screen interface was designed with physical navigation in mind. This forced a kind of simplicity on the interaction flow. The control knob can rotate, push inwards, and move in the vertical directions. These constraints were thus also put on the interface, which results in a clean and seamless user experience.

Furthermore, the display on the knob serves as a minimized way to control the music. All interactions can in theory still be done with the big menu closed, this might be desirable when, for example, the navigation is taking up the entirety of the infotainment screen. The knob would in theory serve a role similar to that of a smart watch. Convenient for quick interactions, but bigger tasks still benefit from having the entire menu open. This application would need to be further explored, since this was not within the scope of this project.

While not entirely new in its concept, the control knob served a crucial role in the exploration of how a commercialized product implementation might look. It helped bridge the gap from conceptual design to actually deployable product.

Figure 15: Photograph of the created physical input device

DISCUSSION & CONCLUSION

DISCUSSION

Limitations of the Research

User Testing

The nature of the relationships between the researcher and users can greatly influence how the users interact and respond to demonstrations of the design. Partially due to COVID-19 restrictions, most tests have been done with people familiar with the researcher. This might have steered responses to be more positive than what the general population's consensus would be. The negative effects of this subconscious steering are seemingly limited however, since the project is a design project where user opinions were only used to validate the designer's choices. There are therefore no research results that have been compromised by the relationship between designer and users. At most, some design features might get a worse response from the general public. This hypothetical discrepancy would be easy to examine through future user tests on a broader audience in the future. Anecdotal evidence from publicly showcasing at a demonstration day seems to suggest that the general public does respond quite well to the ideas presented. More scientific proof would be crucial if the project continues to be worked on in some shape or form.

Dumb 'Smart Suggestions'

One of foundational pillars of this project is the smart, journey-based, suggestions. These smart suggestions have been thought about intensively, but where not implemented to a degree that would resemble the final product. Users could therefore not test how it would be to have the smart suggestions actually available during a drive. The general response to the promise of such a system was good, but, this response could easily be biased since the hypothetical scenario created in the user tests promised better suggestions. It would ultimately be up to the user to judge if the suggestions are actually better. The future works section explores how the usefulness of the smart suggestions can be tested in a more scientific manner.

Future Works

Looking at the project as a whole, there are certain design opportunities that were identified, but were left unexplored. The following section describes some of the further design/research that can be done to forward the developments of advanced car music control systems.

Real deployment

One of the most straightforward future steps to take would be the actual deployment of the system. Especially the 'smart suggestions' part of the project would greatly benefit from being realized in a more technical manner. Studies need to be performed on to what degree the car sensor data actually strengthens the recommendations. Due to the complex, data-driven nature of the recommendation engine, user testing such a system in its infancy can be quite difficult. Intermediate feasibility tests can be performed where the journey based suggestions are secretly curated by a team of researchers. This should in theory give a general idea on if the idea shows promise. Subsequently, steps can be taken to realize the machine learning powered suggestion engine. To fix the cold-start problem that a new recommendation engine might have, user tests could be performed where a set of users are asked about what they music they would define to be fitting for the situation (i.e. sunny highway or rainy village). The results of these tests could then be used to tune the initial parameters of the recommendation engine. To give an example, certain genres could be set as more prevalent on certain road types (i.e. more hard rock on the high way or sadder music when it rains).

If all these tests prove to be promising, the next step would be to take the system actually to production. In this case it would be possible to create the music streaming – car sensor data implementation well before deploying the new interface to the user. User behavior on the old interface can be of service to bootstrap the deployment of the new recommendation engine. In more concrete terms, users' song listening behaviors get monitored similarly to how a streaming service might already do that. The circumstantial car sensor data gets added as an added layer of information. Once the new interface then launches, it will have a significant amount of data to base the smart suggestions on, leading to a better initial experience for the end-user.

As the ultimate validation test, a recommendation engine with and without car sensor data available could be compared through user testing. The car sensor data should be able to make a discernable difference in the user experience. If this is not the case, then either the recommendation engine needs to be redesigned, or the user would not see any benefits from the additional sensor data, rendering the sensor data useless.

Physiological Sensing

As highlighted by multiple users and designer contemporaries, there is an opportunity to add physiological sensing into the mix to provide even more fitting suggestions. In this project the choice was made to focus specifically on creating a symbiotic relationship between the car and music streaming services. Physiology is an area of research of its own and was therefore deemed to not be within the scope of this project.

There are however definitely applications where physiological sensing could prove to be a great addition. One of such use cases would be an implementation with SentiRide (Joosse, 2021), a project that was created at the same time as AUD_I/O, at the same university. The SentiRide project aims to keep bus drivers focused through a number of interventions, such as chair vibrations and visual cues. The underlying technology of the AUD_I/O project could in theory be used to add music as a means of controlling the focus level.

These kinds of additions would be beneficial from a road safety perspective, which ties into the next section of future works, the regulatory control.

Regulatory control

The AUD_I/O project has been designed to completely be in service to the end-user. The recommendation engine's main goal is to please the user as much as possible. Interesting discussions arise however when looking at the same technology through the lens of a regulatory body. One could imagine that regulatory bodies would see benefit in using this technology in the name of road safety. Calmer music could be played when a driver is speeding, for example. Such implementations would need to be carefully thought out, since a system that too strongly intervenes might feel like a violation of free choice. This, of course, is the case with many safety rules, such as safety belts, helmets on bikes, and more. The created safety benefits should generally outweigh the burden of the restrictions. Not informing people that the system is trying to steer their behavior would also be a possibility, but that in turn opens up a whole new plethora of ethical concerns.

All of the aforementioned points should be seriously taken into account when exploring when exploring what the AUD_I/O projects underlying technologies could offer through the lens of public safety.

CONCLUSION

The AUD_I/O project has highlighted some novel interaction opportunities that have arisen from combining the cutting edge of the automotive industry with the new developments in music streaming. By involving the car as a crucial part of the system, a more symbiotic relationship is created. The currently realized demonstrator was able to give users a clear image of what could be possible, to great response. Newly emerged interaction opportunities include the implementation of car sensor data into the music recommendation engine, and novel ways to socially listen to music in and between cars. Both of these opportunities show great promise and deserve to be further explored individually. Some of the lessons of the social aspect could even be applied to different contexts. Examples include ride-sharing, but also completely different contexts like an office setting.

To conclude, AUD_I/O provides a vision for the future where the car and music are more tightly integrated. It makes music interactions smarter, and more social, filling the design niches identified in the related works.

ACKNOWLEDGEMENTS

I would like to thank coaches Debargha Dey and Bastian Pfleging for their continued support throughout the design process, and Niels Winkelman for his contributions towards realizing the 3D-print. Furthermore, all user test participants are thanked for their time and effort.

BIBLIOGRAPHY

Apple. (n.d.). *iOS - CarPlay*. Retrieved June 9, 2021, from https://www.apple.com/ios/carplay/

Frolic Studio. (2012). *Skube, rethinking the home audio listening experience*. FROLIC Studio | Strategic Product Design & Development. https://www.frolicstudio.com/portfolio/skube

Google. (n.d.). *Auto*. Android. Retrieved June 9, 2021, from https://www.android.com/auto/

Hawkins, A. J. (2021, March 25). *Tesla's new touchscreen drive selector doesn't violate the rules, NHTSA says*. The Verge. https://www.theverge.com/2021/3/25/22351044/tesla-nhtsa-touchscreen-gear-selector-fmvss

Joosse, N. (2021, June 4). *SentiRide*. Eindhoven University of Technology. https://demoday.id.tue.nl/projects/bB1yelkjp9

Martelaro, N., Mennicken, S., Thom, J., Cramer, H., & Ju, W. (2020). Using Remote Controlled Speech Agents to Explore Music Experience in Context. *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. Published. https://doi.org/10.1145/3357236.3395440

Nelson, P. (2016, December 7). *Just one autonomous car will use 4,000 GB of data/day*. Network World. https://www.networkworld.com/article/3147892/one-autonomous-car-will-use-4000-gb-of-dataday.html

Ranney, T. A., Harbluk, J. L., & Noy, Y. I. (2002). The Effects of Voice Technology on Test Track Driving Performance: Implications for Driver Distraction. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *46*(22), 1814–1818. https://doi.org/10.1177/154193120204602208

Spotify. (n.d.-a). *2020 Wrapped*. Retrieved June 9, 2021, from https://www.spotify.com/nl/wrapped/

Spotify. (n.d.-b). *Get Audio Features for a Track | Spotify for Developers*. Retrieved June 9, 2021, from https://developer.spotify.com/console/get-audio-features-track/

Spotify. (n.d.-c). *Get User's Top Artists and Tracks | Spotify for Developers*. Retrieved June 9, 2021, from https://developer.spotify.com/console/get-current-user-top-artists-and-tracks/

Spotify. (2020, July 28). *Your Squad Can Now Stream Simultaneously Using Spotify's Group Session Beta*. https://newsroom.spotify.com/2020-07-28/your-squad-can-now-stream-simultaneously-using-spotifys-group-session-beta/

Spotify. (2021a). *Introducing Car Thing from Spotify*. https://carthing.spotify.com/

Spotify. (2021b, June 3). *Celebrate Your Unique Listening Style With Spotify's Only You In-App Experience*. https://newsroom.spotify.com/2021-06-02/celebrate-your-unique-listening-style-with-spotifys-only-you-in-app-experience/

Stegmaier, G. (2020, August 5). *Urteil nach Tesla-Unfall*. auto motor und sport. https://www.auto-motor-und-sport.de/tech-zukunft/tesla-unfall-gerichtsurteil-gegen-touchscreenbedienung-fuer-scheibenwischer/

Tesla. (2021, May 11). *Software Updates*. https://www.tesla.com/support/software-updates

Young, K., & Regan, M. (2007). Driver distraction: A review of the literature. *Driver Distraction: A Review of the Literature.*, *2007*, 379–405. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.469.4178&rep=rep1&type=pdf

APPENDICES

APPENDIX A: Demo day web page:

https://demoday.id.tue.nl/projects/B87LqmZL1E

APPENDIX B: Extended demonstrator video:

https://vimeo.com/561347824

APPENDIX C: Demonstrator link:

http://matserdkamp.com/car/mc

APPENDIX D: User Test 1

QUESTIONNAIRE AUD_I/O ITERATIVE PROCESS

GENERAL QUESTIONS

[introduction explaining the goals of the research]

1. age
2. profession
3. how tech-savvy would you consider yourself?
4. how often do you drive a car? (non-COVID situation)
    a. what kind of car mostly? How does the infotainment system look?
    b. how often does the destination differ. (i.e. not just to work)

CURRENT MUSIC LISTENING HABITS

[participants are questioned about their current music listening habits in the car]

1. Could you describe how you currently listen to audio in the car?
2. You describe mostly using [radio/streaming], why? Could you highlight some of the pros and cons of both systems.

3. How often do you change radio/streaming, what prompts a change?
4. How explorative is your music listening?
    a. Do you listen to albums on release week?
    b. How often do you discover new artists
        i. Would you *like* to discover them more frequently?
        ii. Is this 'discovery' active (searching) or accidental (hearing it somewhere)?

THE NEWS

1. Do you actively listen to the news in the car? (why not?)
2. How would you combine streaming with news?
    [questions subject to change throughout the conversation and only serve as prompts]
    a. Would you prefer the news at set times (18:00), or rather have manual control?
        i. Would you prefer auto-play, or rather get a prompt ('new news broadcast available, press to listen now')
    b. If made possible, would you skip the news per section? Or would you rather sit through the whole thing? (show how the news would be broken up per section) Tradeoff would be that the news is NOT live.
    c. Would you want GPS-style (breaking) news interruptions? (why not?)
3. How do sports fit into this? (same sub questions as above)
4. Would you change these settings on a per drive basis? Or just once?

RADIO SHOWS / PODCASTS

1. When listening to the radio, do you actively seek out or avoid radio shows? What do you like dislike?
2. Do you ever listen to podcasts? How do the two options compare in your opinion.
3. Would you implement talking sections into a music playlist?

PAPER PROTOTYPE TEST

[subjects will be asked to map certain interactions. A paper prototype with unlabeled interactions will be given to the participant. Together with the initiator certain interactions will be mapped to certain actions. Participants will argue the choices they make.]

CLOSING QUESTIONS

1. How would you imagine the system we just created affecting your listening habits?
2. Wherein lies the most value for you personally?

[Thank you for participation]

APPENDIX E: User Test 2

QUESTIONNAIRE AUD_I/O SOCIAL CONTROLS

We are going to go through a few examples of the social part of the project. I explain the general idea, and then we have a small discussion on if and how you would use it.

Here we see the albums suggestions section, and as you can see one of the highlights is ‘loved by steven’. Would it aid your choice making if you knew a friend liked that album?

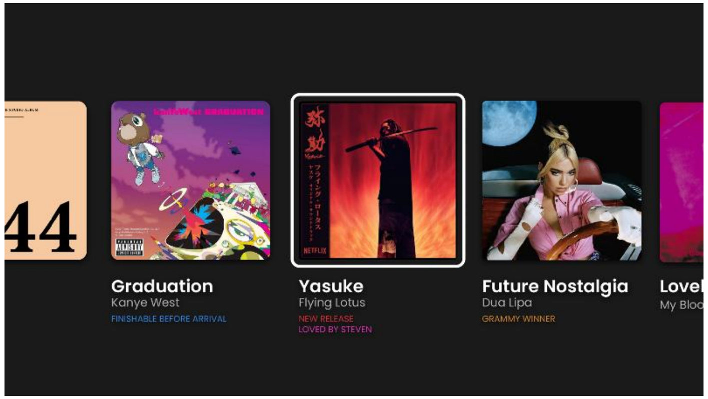

The next mode would be karaoke mode. People like to sing in the car, and this mode makes that even more convenient. We see a bit of gamification also happen, since users are assigned a score, similar to singstar. Would you use this? Why (not)?

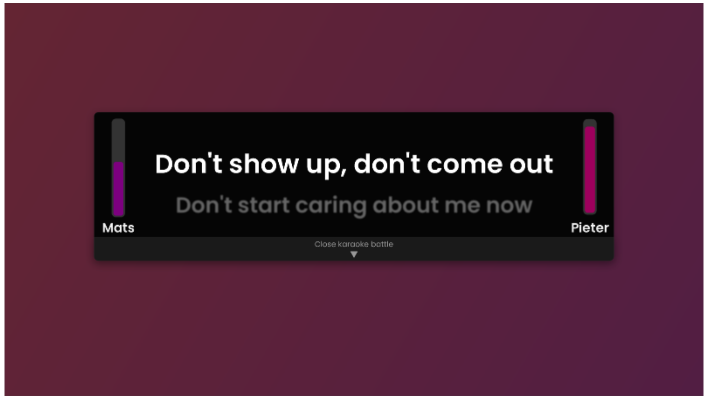

Third, we have the social slider, once a passenger couples to the car’s music session, the social slider can be used to set each person’s influence on the recommendations at large. Would you use this system? With whom would(nt) you use this?

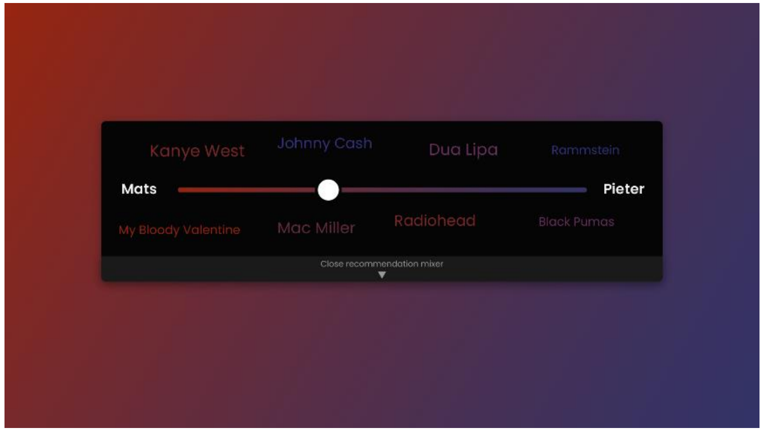

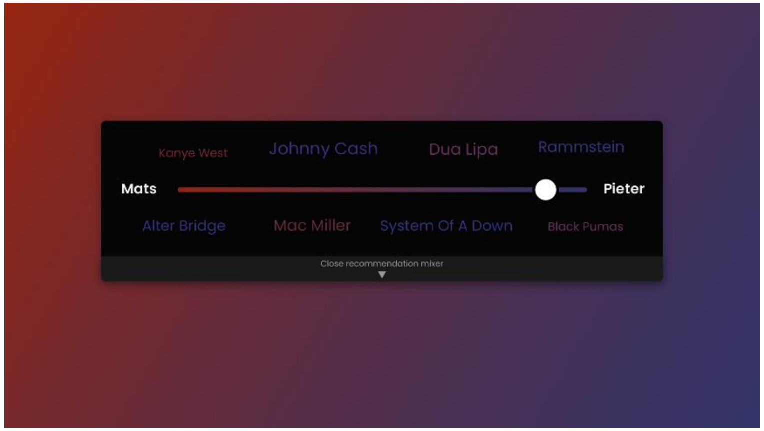

We spoke a bit about the social control within the car, but we can also consider inter-car communication. In this example, we are standing still at a traffic light and the person to your right sent you a song. Would you see yourself accepting this request? Would you send your songs to other road users? Could you describe what would be fun about this system working?

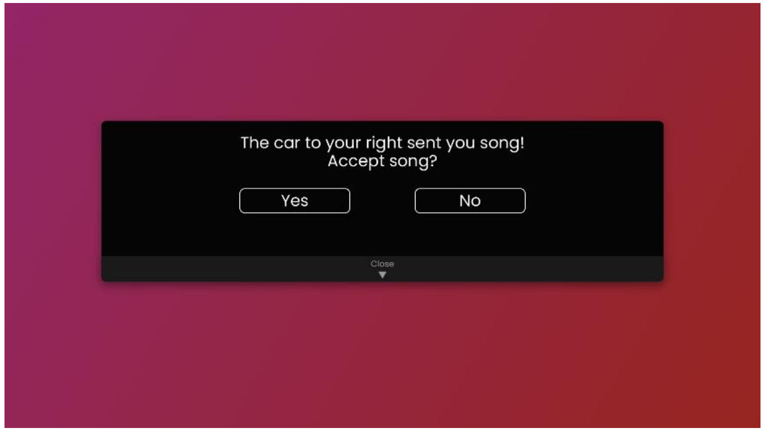

The final implementation would be the regional playlist. The system has a global overview of all the cars, and can make playlists accordingly, in this example, we see that a special playlist has been made for all the people stuck in a traffic jam. Would you join this session? Why (not)? Could you describe what would be fun about this system working?

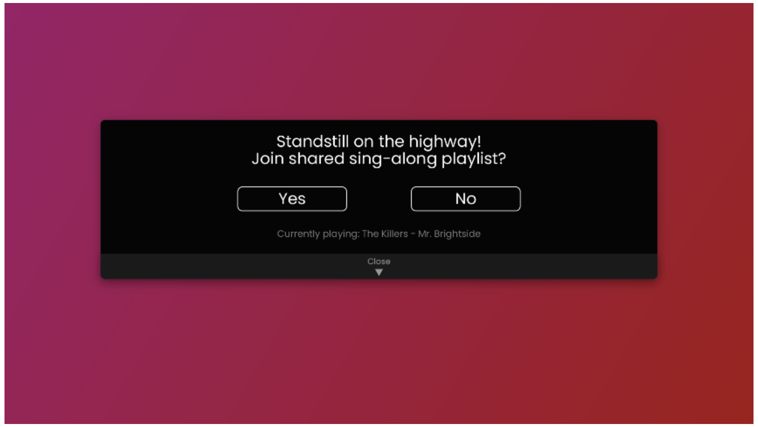

Thank you for your time, this concludes the examples. Do you by chance have any ideas of your own on how this social aspect could be implemented?