Utrecht University

Sharing science,
shaping tomorrow

INAUGURATION

Unravelling music through computation

Prof. dr. Anja Volk

April 16, 2024

Anja Volk

Unravelling music through computation

Inaugural lecture, Utrecht University, April 16, 2024

Watch the official livestream replay for the full experience: https://youtu.be/PiipeB34fbQ

Introduction

Mevrouw de rector magnificus, dear colleagues and friends here in the auditorium and online, liebe Familie und Freunde hier in der Aula und online: Thank you all for coming! I want to introduce in my lecture the field of *Music Information Computing* to you, which always begins with being fascinated and captivated by music.

So let me start with Billi Eilish's song 'What was I made for?' (2023)

She composed this song with her brother Finneas after months of having been 'incredibly uninspired', when she was wondering whether they simply couldn't do their craft anymore. A craft that had won her at that point several Grammys, two Guinness World Records, a Golden Globe award and an Oscar. At the end of a fruitless day, Billie was ready to quit and leave the studio. It was then that her brother sat down at the piano, started playing some chords and those first lines of the song came to her: 'I used to float, now I just fall down, I used to know but I am not sure now, what I was made for'. And so they turned their frustration of not having been able to produce anything useful for months into a marvellous song, which has since then touched so many people, who also wonder about the very human question of their purpose in life.

Let me jump back 200 years from Billie Eilish to Ada Lovelace in the 19th century, a mathematician and writer. She was the first visionary to express the potential of computers outside of mathematics, namely for music. You might wonder: computers in the 19th century? It was actually Charles Babbage's Analytical Engine that inspired her. This fully automatic calculating machine was envisioned by Babbage to evaluate any mathematical formula. It inspired Lovelace to speculate that in the future this engine 'might compose elaborate and scientific pieces of music of any degree of complexity or extent'. This was 1842.

Let me translate this into modern words: 'Assume, you do not have such a genius brother as Billie Eilish has, who knows how to unlock *your* songwriting creativity by playing some chords on the piano. Perhaps the musical output of such a machine could be an inspiration for you to compose a song at the end of a fruitless day, in collaboration with such a machine?' Some version of this idea is 200 years old.

Lovelace's vision about this analytical engine has been a great inspiration for Alan Turing's article on 'Computing Machinery and Intelligence' (Turing, 1950). Her thoughts prompted him to argue that computers can go beyond what humans ordered them to do and therefore can take us by surprise with what they produce.

In this picture, you see members of the Dutch team "Can AI Kick It" preparing their entry song for the first international *AI Song Contest* in 2020, organized here in the Netherlands. The goal was to use Artificial Intelligence as a partner in the creative process of composing a song. We are listening here together with rapper Willie Wartaal to some of the musical ideas our machine learning methods had

produced as starting points for composing a song together. As you might guess from our faces, computers can definitely take us by surprise, as Alan Turing had suggested. Music is a fantastic medium to show this. If we consider Ada Lovelace's speculation in 1842 on using the analytical engine for composing music, it becomes obvious that my rather young research area of *Music Information Computing* is connected to ideas of giants long before it started to actually come into existence.

My own way of engaging with music and computing did not come from the wish to compose music. It came from the desire of wanting to understand why music has such a central meaning in our lives. *How does music work the way it does for us humans?*

I grew up in a family of professional musicians, and I loved the sciences at school. In the end I studied both mathematics and musicology, which was at times rather exhausting because of the double workload. I came several times to the point of feeling I needed to quit one of these studies to survive my crazy life. But I could not decide which one without losing some part of my identity. This way I discovered what *I* really was made for – namely engaging both areas. Back then there existed no research area yet to combine music, mathematics, and computer science. I had only my conviction that musical and mathematical structures had a lot in common. I hoped this would lead me eventually to something.

So it is a tremendous joy and honour for me today that I have the privilege to introduce to you the area of *Music Information Computing*. It intersects with a number of interdisciplinary fields that I had the great opportunity to help establishing on an international level with many colleagues in the field, such as Music Information Retrieval, Computational Musicology, Mathematical Music Theory, and Music, Computing, and Health. In my talk I will take you along examples of questions about music that intersect with these areas – and why we would need to engage mathematics and computer science to investigate these questions.

Music and society

But first let me start with examples on how research about music is intertwined with important challenges we face in modern societies regarding our well-being, resilience, equity, coexistence, and inclusion. Because engaging and understanding music is 'not a luxury, but something as vital and urgent as the air we breathe'. This is how the Turkish writer Elif Shafak (2020) has coined the importance of arts in our lives. Or as Vera Lytovchenko (2022) has answered, when she was interviewed by the New York Times about giving impromptu concerts in the basement of her building in the Ukrainian city of Kharkiv since the beginning of the invasion: 'My music can show that we are still human. We need not just food or water. We need our culture. We are not like animals now. We still have our music.'

People have crossed borders and frontlines because of music. Loewana Weiss shared in last year's Remembrance Day on May 4th (2023) the story about the rescue of her SINTI grandfather during World War II. His family received a warning about the planned deportation to Auschwitz by an SS officer because he loved their music. This was one puzzle piece that allowed them to go underground and be saved.

Regimes fear the empowering role music can play, as the incarceration of the young Iranian student Shervin Hajipour has shown. He composed the song "Baraye" out of tweets of Iranian people fighting in the *Woman, Life, Freedom* movement. His video went viral – and the song was re-played throughout the world. You see here examples from Coldplay in Buenos Aires, and the Dutch Dance company 058, which

beautifully choreographed this song. The Time magazine named Shervin one of the 100 most influential people in the world in 2023.

Systematic investigations on the importance of having music in our lives have been reported in the Health Evidence Network Synthesis Report by the World Health Organization. It reports results from over 3000 studies that identify a major role of the arts in the promotion of health and wellbeing across our lifespan. For example, large-scale community-based music programmes for children who have been exposed to violence, have been found to reduce behavioural difficulties of these children. Other studies reported that music classes can improve social skills and reduce stress hormone levels, hyperactivity, and problem behaviours for children from lower-income backgrounds, or at risk of poorer socioemotional development.

Hence, the field of Music Information Computing is first of all about researching *music as a fundamental human trait*, and about developing technologies that enable novel forms of interaction with music that support and foster this central function of music in our lives.

Music as a fundamental human trait

Let's take a look now at a common musical experience to explain what research questions this field evokes.

Thoraya Maronesy is dedicated to interactive projects with strangers and experimented with the following question: 'What happens if you pin a circle on the ground, play music and invite strangers to *Dance here*?' Let's take a look.

Later in the video you see people dancing together, and you see them inviting strangers to join them. And as you can guess from these pictures, they start synchronizing their movements. This very joyful experience is based on what we call *entrainment*. In our roadmap paper on "Music, Computing, and Health" (Agres, Schaefer, Volk et al, 2021), we have described this process of perceptual entrainment as one of four main affordances that music offers in therapeutic contexts. Entrainment describes the coupling of perceptual processes to the rhythmic structure of music, it allows people to synchronize their movements with each other, it is a core mechanism which underlies a range of responses to music, and it can evoke positive feelings of togetherness and bonding. Entrainment is based on one of the most salient aspects of music, namely its often-regular temporal structure.

This leads to the following questions: What are the musical characteristics or features we need to hear when listening to music, to enable this process of entrainment? Could Thoraya have chosen *any* musical piece and people would have danced along?

The computational perspective would formulate it like this: what *information* from a digitized musical piece do we need to model to investigate how regular the temporal structure of that piece is? What musical information would help us to predict how humans would dance or clap along?

If we find a computational model for this, we might contribute to understanding how entrainment works. But we also might be able to recommend, for instance, dance music, or music suitable for running. We could create collaborative digital games based on music for bringing together children of very different backgrounds and make them play together by employing the principle of entrainment. I will get back to this application towards the end of my talk.

Modelling music information on rhythm: Inner Metric Analysis

My first example of computational modelling of music information addresses the temporal structure of music, which we also call *rhythmic-metric information*. The *beat* of the music, those regular spaced timepoints in the music that we would clap along, describe one part of this information. If we had to dance along, for instance a Waltz, we also need to be able to count along, such as 1, 2, 3, 1, 2, 3. And we need to perceive WHERE in the music the "1" is located, if we don't want to step on our partners' toes, should they start their moves on a different time point in the music.

So how does music tell us this information? Music theorists and cognitive scientists argue that music communicates so-called *accent patterns* to us, where we perceive the greatest accent on the "1", followed by somewhat weaker accents on the beats of 2 and 3, and even weaker accents on the time points in between these beats. The notation of the music usually communicates this with vertical lines we call *bar lines* – and the first note in a bar denotes the "1" with the strongest accent. However, we usually do not have a score in front of us when we listen – hence these accent structures need to be communicated to us in a different way. How does this work?

Let me introduce the computational model of *Inner Metrical Analysis* to you. It attempts to model how accent structures arise in a specific composition. The model takes a digitized musical piece as input, and generates metric weight profiles as output. These profiles describe the temporal structure of the piece (Volk, 2008a). Here you see as an example the analysis of the beginning of Mozart's Jupiter Symphony. The x-axis represents the time, and the blue lines on the y-axis depict the strength of the metric weight at that time point in the musical piece. The higher the blue line, the greater the metric weight. If I zoom into the structure, we can discern the following layers. The first beats of the bars have the strongest weight; the next layer is built by the weights of the third beats in all bars, followed by the layer of the second and fourth beats. The time points in between these beats receive the weakest weights. This is the typical accent hierarchy of a 4/4 – in which Mozart has composed the piece. However, this model does not know anything about bar lines.

The model only knows the time points at which notes begin, and it is agnostic to pitch information. Among these time points, that are here depicted with red bullets beneath the notes, the model searches for all possible pulses. By pulses we mean sets of onsets that are equally distanced in time. The model then assigns a metric weight to each note based on how many of these pulses coincide at that note. Very roughly speaking, if a note participates in many pulses, it receives a great metric weight. For the mathematicians among us: this model is rooted in Mathematical Music Theory. Musical structures, such as patterns and motives, have been modelled by Guerino Mazzola through simplicial complexes in Algebraic Topology. In this case, the metric weights are a quantification of the overlapping structures of the simplicial complex that is constructed from all the pulses detected in a musical piece (Fleischer-Volk, Mazzola, Noll, 2000). And this is how the simplicial complex looks like for this very short snippet of music.

Now what happens if we apply this model to Igor Stravinsky's *Danse Sacrale*? This piece is famous for its rhythmic complexity. You see here a short excerpt that shows how Stravinsky is switching back and forth between different time signatures, on this page between 2/8 and 3/8. When I ask my students in class to listen and to think about whether they perceive a regular temporal structure, usually half of the class

thinks its regular, the other thinks its irregular. Here is the analysis: according to this model, there really is no underlying regular metric structure present, as was in Mozart's piece. Hence, we can use this model to discern whether or not there is temporal regularity present. If Thoraya Maronesy had chosen this piece for her experiment – people might have had a more difficult time to dance along.

And so we used this model together with Bas de Haas as a starting point to automatically detect meter in digitized files (De Haas & Volk, 2016), which is an important task in Music Information Retrieval. We also used it to automatically classify dance music for music recommendation (Chew, Volk & Lee 2004). Moreover, we looked into describing genre characteristics, such as of Ragtime (Volk, 2008b). In Rags, the left-hand part usually communicates the metric structure of 2/4, while the right-hand part creates all kind of interesting rhythmic patterns that go against this metric structure, which we call *syncopation*. Let's listen to a short example.

When analysing many ragtime pieces, we systematically find metric weights in the left-hand part with strong accents on the two main beats in accordance with the time signature of 2/4. The different levels of syncopation in the right-hand parts lead to distortions in the profile. Here you see an example where the strongest accents are located between the main beats, which is very typical for ragtime. Hence the right-hand part has a complementary accent structure to that of the left-hand part, which I would argue is part of the fun of listening to this music. Interestingly, no pitch information is needed for this, which has also been shown in tapping experiments for ragtime music carried out by Snyder and Krumhansl (2001). These weight profiles hence capture aspects of our listening experience, and can be used for determining genre characteristics.

Modelling music information on melody

Another context in which this model was successfully employed, concerns the question of how we perceive melodies. In collaboration with the Meertens Institute in Amsterdam, which is a Dutch cultural heritage institution, we worked in a joined NWO-CATCH project on the following task. The institute's folk song collection of *Onder de groene linde* should be made accessible to the general public via a search engine. It contains more than 7000 recordings of folksongs made in the 1950s to 1980s. Until this project, anyone who was searching for a specific song usually had to call in to the institute and then an expert would go and search for it in the catalogue. In times of digitization, the wish arose to enable users to search themselves in the collection and access the richness of this Dutch musical heritage.

This collection had the following specifics: folk songs are usually transmitted orally, without any notation involved, hence in a word-of-mouth kind of process. One person sings a song to another person, who then sings this again to the next person. As you can imagine, we hardly ever exactly re-create the song we have heard. This process of transmission of songs over centuries leads to the creation of several variations of what was originally one song. While the historic process itself is lost, and all we have now are the recordings, the musicologists of the Meertens institute have classified these variations. All songs that are sufficiently similar to each other, are classified into a so-called *tune family*. These melodies are hypothesized to have a common origin. Hence, a search engine for this collection needs to be able to capture the similarity relations between songs. So we interviewed the musicologists on the nature of the similarity they perceived, such that we would be able to develop a computational model for it.

As computer scientists we expected an elaborate list of musical characteristics the musicologists would mention. However, their answer was something like 'Well, this is intuitive – you just listen to the music, and then you know! It is really not that difficult!' For a computer scientist this was somewhat of an

unexpected answer. But it is a very typical answer on how we as humans tend to function. We become experts over years in a certain area, but we are often rather poor at describing *how* we are doing what we are doing.

We then sat down with the musicologists and listened together to the songs and designed an annotation study to unravel the concept of melodic similarity. The experts annotated the musical characteristics that had come up in our discussions for 360 melodies. They annotated how important they were for each melody to classify it into a certain tune family.

Our quantitative analysis of these annotations revealed that the most important musical characteristics the musicologists considered, are repeated patterns or *motifs* between songs (Volk & Van Kranenburg, 2012). By this we mean short snippets of the melody that reoccur, as you can see here with the blue and red circled group of notes in two songs.

In parallel to these investigations with the musicologists, we were working on several computational models which might capture these relations between the melodies. The most successful one, introduced in Peter van Kranenburg's PhD thesis (2010), became the backbone of the online search engine. You can click on the button *find similar melodies*, and then you receive a list of melodies that are similar in their melodic characteristics to your query. The underlying algorithm has been inspired by the concept of *sequence alignment* from bioinformatics. This is a method for identifying regions of similarity in the sequences of DNA or RNA. In our case, we do not align DNA sequences, but melodies.

The alignment algorithm takes two melodies, aligns them note by note, and measures at each instance how well the notes match each other, or how many gaps need to be inserted. Based on this, it outputs a distance value for the two melodies. With this distance value, we can successfully classify 99% of the melodies into their correct tune family. This means not only that we could build a successful search engine for the collection with this model, but also that the musicologists had worked with great consistency when assigning tune families based on how they perceived the similarity between songs. Otherwise, we would never have been able to re-create this with such a computational model. Hence, experts are excellent at what they are doing, even if they cannot explicate very well *how* they are doing what they are doing.

This successful alignment model led to another question. The insight from our annotation study had been that repeated patterns are the most important carriers of information for the similarity. But the alignment model takes every note into account and does not consider specific patterns. Was something wrong with the musicologist's intuition, and how could we test the role of these patterns for the similarity?

We evaluated the role of these patterns in the following way with master student Peter Boot (Boot, Volk, de Haas, 2016). We threw away those notes of the melody that were not part of an annotated pattern. Hence, we created what we called a *sparse melody*. This sparse melody would then contain the most important information regarding similarity as the musicologists had argued. Then we classified these sparse melodies into tune families and tested whether the answer would still be correct.

Indeed, with our computational experiments we figured out that we can throw away 60 percent of the notes in a song and the accuracy rate only drops to 0.88. This is still a pretty good classification result when you throw away more than half of the song material. And it was a much better classification result

than when we would randomly throw away 60 percent of the notes. This way we learned that indeed the annotated patterns from the musicologists contribute significantly to the similarity between songs.

We had started this investigation into how we perceive melodies with musicologists who rather implicitly used their expertise on assigning tune families. Throughout the computational modelling process and the collaboration, we were able to make this implicit knowledge more explicit. This demonstrates an important way on how we gain insights about music by modelling music information computationally, next to delivering a very useful search engine.

Modelling music information on harmony

My next example of modelling musical structures concerns harmony in music, which Hendrik Vincent Koops investigated in his PhD thesis (Koops, 2019). An important inspiration for this topic came from our collaboration with *Chordify*, a music e-learning platform.

Chordify allows users to play along whatever Youtube song they like by automatically recognising chords, and aligning them to the music. And by the way, Chordify's beginning started at our department, when Bas de Haas collaborated with the software technology group on employing the functional programming language of Haskell in his *HarmTrace* model. The underlying computational task is called *Automatic Chord Estimation*, and it requires an algorithm to detect at what time point in the audio signal which chord is present. Since Chordify allows users to edit these chord sequences, they found out that users disagree to quite some extend on which chords they think should be played. The sometimes rather excessive amount of edits sparked the question of whether listeners have any common ground on what they consider a suitable chord sequence.

To investigate this question systematically, we asked four highly trained musical annotators to determine the chords in 50 popular songs, and we constructed the *Chordify Annotator Subjectivity Dataset* (Koops et al, 2019). We investigated how much the annotators agreed on the chord labels, and we compared different granularities of information about these chords. For instance, we found only 73 percent overlap for the traditional major–minor chords and 54 percent overlap for the most complex chord labels. These numbers show that not only Chordify users with all kinds of different musical expertise disagree with each other, but also highly trained annotators disagree a lot. Hence, we have to take subjectivity into account in computational models of harmony. We will never find the one-fits-all solution of chord labels for a given musical piece.

Vincent Koops then realized the first approach to personalizing chord labels, namely by modeling harmonic interval-based chord representations. He integrated these representations from multiple annotators and deep-learned them from audio recordings. He showed that with this model we can then personalize chord labels for individual listeners and hence offer better chord labels to them (Koops et al, 2018).

With these examples on extracting musical information for rhythm, melody, and harmony I have named already a number of different computational approaches for modelling musical structures. There are many more we have employed, since musical structures are very complex and diverse. Hence, music delivers very interesting data for computer and information scientists, and provides a playground for model development and model comparison.

Application of musical structures in novel interaction technologies

In the following I want to highlight the application of musical structures in enabling new forms of interactions with music through technology. Some of my examples were already intertwined with interactions, like playing chords along a Youtube song or interacting with a collection of folk songs. In the following, I want to highlight one specific application area, namely the field of *Music, Computing, and Health*. In our roadmap paper on this topic (Agres, Schaefer, Volk et al, 2021), which resulted from a week-long workshop at the Lorentz Center in Leiden, we identified a wide range of scenarios in which computation can play an importing role.

Here I want to highlight two scenarios: First, the analysis of musical interactions in clinical improvisations. Computational models for analysing these improvisations in music therapy allow to study patient-client interactions. In an overview paper with Tinka Veldhuis, and music therapists Jos de Backer and Katrien Foubert (Volk et al., 2023a) we argue, for instance, how they can support the initial clinical assessment of the client, or the monitoring of the process the client takes during therapy.

Second, the development of *serious* or *applied games* for music therapy. These games should facilitate additional training moments for clients at home, in between weekly sessions with therapists that are usually constrained in time and frequency. While therapists have mentioned the need for such applications already in 2019 in our workshop, the COVID pandemic that followed made this need for technological applications in administering music therapy even more urgent.

With Sander Bakkes, Laurien Hakvoort, and Rebecca Schaefer, we worked on the first applied game for *Musical Attention Control Training* (Volk et al, 2023b). This is an intervention method in Neurologic Music Therapy which optimizes attention processes in the context of cognitive rehabilitation. Structured music-making that activates different parts of the brain and that stimulates neural connectivity, allows to train attention processes. This is crucial for clients with attention deficits, for instance related to ADHD or Parkinson Disease (PD). Master student Ermis Chalkiadakis developed the game *The Last Minute Gig* targeted at clients with Parkinson's (Chalkiadakis, 2022). The game emulates the role of the therapist, by playing music to the client, to which the client plays along by pressing a button which automatically produces chords and steers a guitarist avatar. The client is asked to adjust their playing along whenever they recognize a change in the music, which follows established protocols in neurologic music therapy.

We evaluated this game in a pilot study with around 20 healthy participants playing the game on 10 different days within 2 weeks, and a control group which did not play the game. For both groups, we tested three cognitive functions before and after these two weeks. While we see a slight improvement for the group who played the game, these improvements are not yet significant. More importantly, we learned from this pilot evaluation that those participants who had greater musical experience enjoyed the game less, felt less challenged and felt the least skill increase over the two weeks. From these findings we conclude that such a game is a very promising context for improving automatic music generation algorithms. We need to generate a lot of music for these repeated sessions, and the music needs to be complex and interesting such that players enjoy the game and keep playing for improving their attention.

Developing and evaluating this game required the collaboration with experts in music therapy, neuroscience and game design. My next example adds even more disciplines. Funded by grants from the

Dynamics of Youth theme of Utrecht University, we have developed the game *KenDo* in the *AlgoRhythm* project. This project seeks to support children with sensory processing issues, as is often the case for people with *Autism Spectrum Disorder*. Research in music therapy has indicated that musical rhythms can be employed to detect and ameliorate these sensory processing difficulties.

To support the rhythmic training, we have transformed the music training concepts at the University Medical Center Utrecht into a digital game that the children can play in the safety of their homes. The game idea was developed by this group of bachelor students (and please forgive me when I stop reading out all names in the coming slides, there are simply too many). Ken, the main character of the game, wants to learn sword fighting in the Dojo, and they can do this only by ear because they are blind. Hence, a drummer presents rhythmic patterns, which Ken needs to listen to. The task of the child is to listen together with Ken to these rhythms, and then they reproduce these rhythms by tapping on the interface. By doing so, they steer the avatar of Ken who starts sword fighting when the child taps. This way the child helps Ken to learn the art of sword fighting.

Apart from the student group I already mentioned, all of these experts and students from different fields and institutions are involved in developing, testing, and adapting the game in a truly interdisciplinary manner. In a long-term perspective, we want to learn whether and how children improve their sensory processing if they play the game over longer periods.

My last example concerns a collaborative music-based game in the context of social inclusion. Video games have become a very popular form of entertainment, but most of them rely heavily on visual cues. This excludes a rather large group of gamers, namely the visually impaired. It is estimated that around 2 billion people world-wide live with some form of visual impairment. In the Netherlands, three-quarters of visually impaired children follow () programs at regular schools, but they are facing serious barriers when it comes to participation and social engagement. Playing with each other is a very important form of social interaction. Not being able to play the video games their sighted peers play, poses a serious hurdle of participation for these children. Hence games that can be played and enjoyed both by sighted and visually impaired children are urgently needed.

This is what *Bongo beats: tap with me* attempts – a music-based collaborative game which requires playing on a bongo along a song that is played. One player either sees or feels via tactile vibrations, when the bongo drums need to be hit, and then informs the other player to indeed then hit one of the drums or both. According to our pilot studies, the players collaborated well with their teammates, enjoyed the game and showed a willingness to keep playing it. We also learned that if they did not like the specific music we had selected, they enjoyed the game to a much lower extend. This effect makes it necessary that players are able to choose a musical piece they would like to play along. Hence, we need to offer a whole range of musical pieces in the game.

Determining for many pieces the time points to which the players should play the bongo part, requires automatic beat detection algorithms. With this task we are back to the topic with which I began the lecture about extracting information from digitized music to determine metric information. We need to be able to automatically extract beats to then assign time points in the game to which the drums need to be hit. Apart from the master students Evangelos Potamianos (Potamanios, 2022) and Marc Ferrigi (Ferrigi, 2023) who worked on the development, evaluation and adaptation of game in their master theses, this collaboration invited again a new set of experts. This time experts on visual impairment and human computer interaction – and the most important experts, the visually impaired players themselves.

It is a great joy collaborating with the Bartimeus foundation on ideas on how to improve the game in the future and to learn from their feedback.

All of the game examples I presented here, are based on employing musical structures. Hence, understanding musical structures is a prerequisite for building these interaction technologies. Here is a bit of a bird's eye view of the different topics I have touched upon with my examples. I hope it is obvious by now that realizing these topics is only possible through collaborations with different groups in our department, and outside of the department, with experts from different faculties and institutions.

The importance of interdisciplinarity in research

Hence, let me address interdisciplinarity in general, and I come back to this circle we started with.

Circles of academic disciplines have established their own methods and values. Anyone who is involved in more than one discipline knows how almost bewildering different our disciplines can be. So we have to work with these differences when we step outside our established circles. There are simply too many urgent questions we cannot tackle within one discipline alone. I have shown you a number of them involving the complexity of music.

But working with differences between disciplines also applies to educating the next generation of students. When they come to us to study a certain field at an age where they ask themselves 'What was I made for?'- they need a circle describing what the field is about and what expertise they will acquire. However, in teaching the master course *Sound and Music Technology*, I also experience the useful exercise of stepping outside of one's disciplinary circle for the students. It requires them to embrace fields that are sometimes lacking the preciseness of definitions a mathematician would demand, and to deal with the sometimes messy experiences of human beings.

Interdisciplinarity in computer science education

This interdisciplinary education sometimes leads to novel projects I did not have in mind when preparing the course. For instance, the game for visually impaired children was initiated by students in this course. I learned about the music examples from the Iranian *Woman, Life, Freedom* movement from students when we compiled our collaborative December playlists together. We do this each year to widen the options for listening to music beyond the typical songs that are played that month. This provides a fascinating way of learning for me, what societal and personal topics our students care about – when they contribute songs about climate change, social injustice, gender identities, love, fear and hope of the young generation. In 2021 we were even invited to NPO Radio 1 to present this playlist.

I am delighted to experience the eagerness of students to engage in projects that take our responsibility serious to develop technology that embraces all groups in society. We still have a long way to go with this responsibility. For being able to develop 'algorithms that work for everyone', as the Dutch AI coalition has coined it, we need the expertise of scientists and students from all parts of society. When kids ask themselves 'What was I made for?', we need to work towards a future in which as many girls, boys and everyone in between and beyond will think that computer and information sciences might be one of their exciting options. And they might think of themselves as having multiple talents beyond narrow disciplinary circles that are going to be very useful in their future careers. The field of Music Information Computing is certainly an exemplary field showcasing the usefulness of interdisciplinarity.

Diversity and inclusion in computer science

To ensure that everyone with a talent and eagerness for pursuing a career in the sciences feels encouraged to do so, we also need to re-think how we connect our past with our future. I still remember when I walked for the first time into the Senate Hall here, with all of its portraits of dignitaries on the walls to honour those who came before us. It is a very festive room - you will all get to see it during the reception. But it also made me feel uneasy – as if I did not belong there. I applaud the university's initiative to add portraits of all female professors, but obviously, we are still lacking other dimensions of diversity here – a very visible one would be someone with a different skin colour than my own.

For future generations, we need to do better. We cannot afford to waste any precious talent because someone might think they are not made for a specific path since they do not see anybody like them on such a path already. It was an amazing experience for me to initiate together with three colleagues the WIMIR mentoring program in the field of MIR in 2016. This program serves meanwhile all underrepresented groups and runs yearly on all 5 continents. I hope that in the future researchers will be as diverse as the participants you see here who invested their time into this program, for which I am deeply grateful.

I also want to thank our department for supporting the founding of the Women in Information and Computing Sciences network and the ICS Diversity committee. I am very grateful for the many collaborators within and outside of these committees and their ongoing efforts to help us becoming an inclusive department to researchers and students of very different backgrounds.

Let me come back to my experience of entering the Senate Hall with all these dignitaries at the wall. When I received a photograph from one of our PhD graduations in this room a couple of years ago, my then 6-year-old son got very excited 'Mom, I know these guys, they are all famous pirates'. Now his beginner's mind really helped me entering next time this room. When heaviness descended upon me, I whispered to myself 'They are all famous pirates'. For the future, I hope that I will have the opportunity to whisper encouraging words into the ears of others who are doubting their path of what they were made for.

With my expertise in interdisciplinary research, I want to continue building bridges between disciplines and people and making new connections. Hence, I am very grateful that I can contribute to the Dynamics of Youth community at Utrecht University. This strategic theme is all about establishing interdisciplinary connections across our different faculties to answer crucial questions for future generations.

And for the field of Music Information Computing, I want to engage in it as a gate opener for exciting interdisciplinary research involving computer and information sciences. I want to develop music technologies which help unlock music's fundamental meaning in our lives for people with very different backgrounds and abilities, and specifically for the next generation.

Thanks

There are so many people to thank for the fact that I am standing here today talking about Music Information Computing. You have seen many of their names on the slides, already beyond what I was

able to read out loud. Obviously, doing science is not a one-woman-show, and establishing new research areas specifically not. Hence my deep gratitude goes out to the many researchers who collaborated on projects with me, and who joined forces in establishing the strong connections we see today between music and computer science, by founding and running societies, conferences, journals, or mentoring programs.

Utrecht University has been an active part of these endeavours for many years, since Frans Wiering and Remco Veltkamp started computational research into music here more than 20 years ago. So we have become an important player within an international network in this research. I am very proud of that and grateful to all contributors at our department over these many years. Also, I thank our students who are curious and committed to contributing their ideas to the field and inspire me to think along new avenues. Today, I also want to thank Frans Wiering and Mirjam Visscher for organizing our shared musical experience that we will engage in after my lecture and when we leave the auditorium. For the artwork on my title slide I thank Tamara Mtsentlitze and Dall-E.

My biggest gratitude goes out to my family for their steadfast support. My husband changed twice countries with me for my job hunt. He always encouraged me to pursue the path that he believed I was made for. When we moved to the Netherlands with a 5-months-old baby, one job, no permanent housing and no childcare, my mom moved with us and stayed for many months to be there for our daughter. My son is probably too grown up now to believe that I have graduated as a pirate, but perhaps *professor* will also do. It was my daughter who introduced me to Billie Eilish's song we heard at the beginning. I was cooking and it was her second attempt at drawing my attention to this song. This time I did pay attention and stopped stirring the pot while listening together with her.

I strive to continue – actually I hope we all strive to continue – to listen to the next generation when they seek to draw our attention to what matters to them, and to be inspired by them as much as we are inspired by those who came before us and on whose shoulders we stand.

Ik heb gezegd.

We will now perform with our departmental choir Bobby McFerrin's Song Seven from his album *Circlesongs* and we will walk out of the auditorium with this song. I sing and clap this song each year with the students of the *Sound and Music Technology* course, as a way to practice some of the musical concepts learned in the course – like rhythm, melody and harmony. Bobby McFerrin is dedicated to letting everybody join in making music, so let's do this together now. The easiest way to join in is this body clapping part. We in the choir will start with clapping only, and we then add a very short basic melody that we will repeat, and then add more voices to it. Please join in the clapping and or the singing of the basic melody. We will employ Utrecht University's teaching models which says *It is okay to make mistakes*.

References

Agres, K., Schaefer, R.S., Volk, A., Van Hooren, S., Holzapfel, A., Dalla Bella, S., Müller, M., De Witte, M., Herremans, D., Ramirez Mendelez, R., Neerincx, M., Ruiz, S., Meredith, D., Dimitriades, T. & Magee, & W.L.(2021). Music, Computing, and Health: A Roadmap for the Current and Future Roles of Music Technology for Health Care and Well-Being, Music & Science. Volume 4, pp 1-32.

Boot, P., Volk, A., & De Haas, W.B. (2016). Evaluating the Role of Repeated Patterns in Folk Song Classification and Compression, Journal of New Music Research, 45(3), 223-238.

Chalkiadakis, E. (2022). Developing and evaluating a Musical Attention Control Training computer game application, MSc Thesis Utrecht University.

Chew, E., Volk, A., Lee C.Y. (2004). Dance Music Classification Using Inner Metric Analysis. In: The Next Wave in Computing, Optimization and Decision Technologies. Springer OR/CS Interfaces Series No.29, 355-370.

De Haas, W.B. & Volk, A. (2016). Meter detection in symbolic music using Inner Metric Analysis, In Proceedings of 17th ISMIR Conference, New York, USA, 441-44.

Ferrigi, M. (2023). Improving the Player Experience of Collaborative Multiplayer Games for Visually Impaired Children. Master thesis Utrecht University.

Fleischer-Volk, A., Mazzola, G., & Noll, T. (2000). Computergestützte Musiktheorie, Musiktheorie, 15(4), pp. 314-325.

Koops, H.V. (2019). Computational modelling of variance in musical harmony, PhD thesis Utrecht University, 2019.

Koops, H.V., De Haas, W. B., Burgoyne, J.A., Branson, J., Kent-Muller, A. & Volk, A. (2019). Annotator subjectivity in harmony annotations of popular music, Journal of New Music Research, 48 (3), pp. 232-252.

Koops, H.V., de Haas, W.B,, Bransen, J. and Volk, A. (2018). Automatic chord label personalization through deep learning of shared harmonic interval profiles, Neural Computing and Applications, pp. 1 – 11.

Potamianos, E. (2022). Developing an inclusive multiplayer rhythm game for visually impaired children, MSc thesis Utrecht University.

Snyder, J., & Krumhansl, C.L. (2001). Tapping to Ragtime: Cues to Pulse Finding, Music Perception, Vol. 18, No. 4, 455–489.

Turing, A.M. (1950). Computing Machinery and Intelligence. Mind 49: 433-460.

Van Kranenburg, P. (2010). A Computational Approach to Content-Based Retrieval of Folk Song Melodies. Ph.D. thesis, Utrecht University, 2010.

Volk, A. (2008a). Persistence and Change: Local and Global Components of Metre Induction using Inner Metric Analysis, Journal of Mathematics and Music, 2(2), 99-115.

Volk, A. (2008b),.The Study of Syncopation using Inner Metric Analysis: Linking Theoretical and Experimental Analysis of Metre in Music, Journal of New Music Research, 37(4), 259-273.

Volk, A. & van Kranenburg, P. (2012). Melodic similarity among folk songs: An annotation study on similarity-based categorization in music, Musicae Scientiae, 16(3), 317-339.

Volk, A., Veldhuis, T., Foubert, K., & de Backer, J. (2023a). Towards Computational Music Analysis for Music Therapy. In *Proceedings of the 24th International Society for Music Information Retrieval Conference, Milan, Italy, 2023.* (pp. 247-256).

Volk, A., Chalkiadakis, E., Bakkes, S., Hakvoort, L. & Schaefer, R.S. (2023b). Developing and Evaluating a Musical Attention Control Training Game Application, Proceedings of the Sound and Music Computing Conference, Stockholm, Sweden, pp. 261-268.