Augmenting the experience of playing the piano: controlling audio processing through ancillary gestures

Niccolò Granieri

A thesis submitted in partial fulfilment of the requirements of Birmingham City University for the degree of Doctor of Philosophy

Arts, Design and Media
Birmingham City University
Royal Birmingham Conservatoire
United Kingdom
July 2020

To my parents, Elide and Michele

Acknowledgements

I would like to thank my supervisors and mentors Prof. Lamberto Coccioli, Dr. James Dooley, Dr. Jason Hockman and Dr. Tychonas Michailidis for their invaluable support throughout these years of research. Even though being a team of five was a challenge in itself, they provided me with motivation and advice whilst also challenging my ideas.

I would also like to extend my thanks to Dr. Jamie Bullock who infused me with his love for research and without whom I would not have started my doctoral journey in the first place. My first year in the Integra Lab as a research assistant has been one of the most significant years in my life for my career.

I extend my thanks to the Royal Birmingham Conservatoire for funding this research and supporting conference attendance. Its resources, facilities and cohort of musicians were invaluable to my research project.

I would also like to thank my PhD research colleagues and Integra Lab fellow researchers for the silent working sessions, the uncontrollable laughter and collaborative efforts that have been undertaken throughout these years. I would specifically like to thank Dr. Joe Wright for his support and friendship without which I would have been lost, and *software wizard* and friend Jefferson Bledsoe, that co-developed Reach with me saving me from many sleepless nights.

This dissertation and the various tests would not have been possible without the patience, time and effort of the pianists that took part. In this regard, I would like to acknowledge the contribution of Thomas Harris, Elisei Stefanescu, Larence Shore, Ruifan Xuan, Anthony McLean, Kelly Ho, Stella Roberts, Elliott Sanson, Miriam Fornari, Richard Stenton, Juliet Jeffries, Cobalt Piano Duo (Kate Halsall & Fumiko Miyachi), Rosie Tunley, Chloe Knibbs, Maria Chiara Argirò, Jorge Rodrigo Quintás, Michael Gee, Piera Onacko and Andrew Woodhead.

I have been lucky enough to have everlasting friendships in Italy that have remained just as strong throughout my five years in the UK. I would like to thank Francesca and Tommaso for the surprise visits, the laughter and support in the moments of need, and Beatrice for being the best friend anyone could ask for.

Even though this thesis is already dedicated to my parents, I would like to thank them again for their endless support, and for pushing me to pursue my goals and dreams. Together with them, I would like to thank the entirety of my extended family including Diego and Barbara for the scheduled weekend Skype sessions, the laughter, the discussions and all the love. Last but certainly not least, I would like to thank Selina for being by my side from day one, first as a friend and then as my partner, sharing with me everything this PhD had to give.

Related Publication

The thesis may include text from the following publications.

1. Granieri, N. (2017) from piano to piano. *Research Conference 2017 (RESCON17)*, Birmingham City University, Birmingham, United Kingdom.

2. Granieri, N, (2018) Reach - Designing keyboard instruments with pianists in mind. *Sound, Image and Interaction Doctoral Symposium (SIIDS18).* Madeira Interactive Technologies Institute, Madeira, Portugal.

3. Granieri, N., Dooley, J., Michailidis, T., (2019) Harnessing ancillary micro-gestures in piano technique: implementing microgestural control into an expressive keyboard-based hyper-instrument. In: R. Hepworth-Sawyer, J. Hogson, J. Paterson and R. Toulson, eds. *Innovation in Music - Performance, Production, Technology and Business*. London: Routledge, pp.269-282.

4. Granieri, N., Dooley, J. (2019) Reach, a keyboard-based gesture recognition system for live piano sound modulation, In: *Proceedings of New Interfaces for Musical Expression 2019 (NIME19).* Porto Alegre, Brasil, 2-6 June 2019.

5. Granieri, N., Dooley, J., Michailidis, T., (2020) Retaining Pianistic Virtuosity in #MIs: Exploring Pre-Existing Gestural Nuances for Live Sound Modulation

through a Comparative Study. In: *Perspectives On Music Production - Innovation in Music*. London: Routledge.

Abstract

Pianists spend many years practicing on their instrument. As a result they develop alongside their pianistic technique a set of gestural nuances that enable them to perform expressively and establish their own acoustic signature on the piano. This *mute layer* of nuanced gestures is rarely taken into consideration when developing new keyboard-based gestural interfaces. These often usually require new gestural vocabularies to be learned resulting in a disruptive experience for the pianist. The main objective of this research is to investigate how new keyboard-based gestural interfaces can enable musicians to control and transform live piano sound through the gestural nuances embedded in their technique. Specifically, how keyboard interfaces with nuanced gestural control can extend the creative possibilities available to classically trained pianists, thus stimulating new approaches to build intuitive interfaces for musical expression, and new ways of learning and playing digital instruments. Towards this goal, interviews, user tests and case studies were conducted with a range of pianists coming from different musical backgrounds, and Reach, an augmented instrument for live sound modulation controlled by gestural nuances embedded in the pianistic technique was developed.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This thesis focuses on the implementation of gestural nuances already present in the pianistic technique for the control of sound modulation parameters in new gestural musical interfaces. The gestural nuances at the core of this investigation are interwoven in the technique of the pianists as a result of years of instrumental practice and are responsible for their acoustic signature. The main objective of this dissertation is to observe whether classically trained pianists can take advantage of these ancillary gestures in gestural musical interfaces and successfully control sound modulation parameters. Specifically, this research has the aim of analysing if ancillary gestures can improve the learnability of gestural musical interfaces and if this implementation could favour the rise of new creative paradigms and extend the potential sonic output available to the musicians. Towards this goal, user tests with pianists of different musical background and semi-longitudinal case studies have been conducted and Reach, the augmented instrument for live piano sound modulation used in these tests, was created.

1.1 Motivation and Background

This research project is rooted in my personal development as a pianist and musician, and in my late but fervent passion for music technology. The creative barriers I encountered when approaching the technological world with a traditional musical background, and the experiences I had with a number of artists and collaborators, brought to light the questions at the core of this thesis.

Having spent almost the entirety of my first 20 years of life studying classical piano, and having approached technology only at a later stage of my musical and academic career, I believe that interfaces tailored to performance with technology present a barrier to musical creativity when approached by traditionally trained musicians. Most of the time, this barrier is posed by the steep gestural learning curve that these interfaces pose: musicians are required to learn a new bespoke gestural vocabulary to operate these interfaces. As noted by Nicolls (2010) in the case of keyboard based instruments, trying to learn and apply these new gestures whilst retaining instrumental technique and freedom resulted in being a hard task to perform. Thus, there is the need to improve gestural musical interfaces by keeping traditionally-trained musicians as the main focus and involving the musicians themselves in the development stages to make these interfaces more accessible. Accessibility—intended as a set of affordances provided to the user to facilitate their initial approach with the instrument—should be a core aspect when designing new musical instruments and interfaces for musicians, focusing more on the first encounter with the digital environment itself. Understanding how traditionally trained musicians approach new musical interfaces helps make these interfaces more accessible and potentially overcome the longevity issues typical of these novel creations (Morreale & McPherson, 2017).

The strong gestural component in performances with traditional musical instruments has led researchers to investigate the topic of music interaction and develop gesturally controlled musical instruments. As noted by Paradiso (2004) although many different alternative controllers have been developed in the past 100 years, apart from a few rare exceptions such as the Theremin, most of them have been rapidly replaced by new devices. For Paradiso this was mainly due to two separate factors. First, the majority of the controllers were not playable beyond a basic level and did not allow any sort of virtuosistic playing. Second, due to the instability of the gestural interaction and the interface, these instruments were unable to support the increase in ability of the performers and to provide precise means of interaction.

In the 1980s the initial implementation of the Musical Instrument Digital Interface (MIDI) communication protocol (Smith & Wood, 1981) provided a standard for musical interfaces and a more stable ground for the development of gestural musical interfaces. The concept of gestural musical interfaces, in the context of this dissertation, is going to be used to group all of the musical interfaces containing a gestural component to control and generate sound. From the implementation of the MIDI protocol onwards, there have been a number of gestural musical interfaces that explored how different gestures could control musical parameters. An initial review of these interfaces was conducted by Wanderley & Battier (2000) and since then most of the gestural interfaces have been published at the New Interfaces for Musical Expression (NIME) conference.

Gestural musical interfaces have been around for almost a century (Crab, 2019), starting with the Theremin in 1928 (Ssergejewitsch, 1928). Only in the past few decades though, musical gestures have started to be interpreted as something more than a simple movement: an expressive gesture underlying a special meaning (Christopoulos, 2014). Lewis & Pestova (2012) use the term *sounding gestures* to describe aspects of sound tied to a physical movement. This change of perspective helps define and understand the impact of ancillary gestures on the expression of the musical performance. As explained by Lewis and Pestova (2012: 3): "In the case of singing lines, for example, the sounding gesture will have a lot more to do with the production and control of breath than with the striking of strings".

Since this realisation, with the increasing number of gestural musical interfaces being developed, the focus started shifting towards the richness of traditional musical gestures and their complexity. The investigation of *natural gestures* and the development of user-oriented gestural musical interfaces with intuitive gesture-sound mapping strategies (Maes et al., 2010) became a goal for researchers and developers. This focus pushed the development towards gestural musical interfaces that were meant to be a creative extension of ones natural and acquired gestural proficiency, rather than a difficulty to overcome (Rebelo, 2006).

However, despite the shift in paradigm of gestural musical interface development, these interfaces—whether traditional or novel—can often have an effect on the technique of the musician that approaches them. Tied to this phenomenon is the concept of interface theorised by Norman (1990). Norman states that every interface is unnatural because of the very essence of it. An interface *interfaces* a subject with reality; a musical instrument interfaces a musician with the resulting music. Acoustic musical instruments are interfaces that—supported by heritage—culture and years of practice become embodied and thus transparent (Leman, 2008). For example, the piano represents the music it produces, and is not an interface anymore. The shift in role of the piano is a result of the embodiment concept: whilst the piano interfaces the musician with the music, in practice the embodiment of the instrument makes the instrument itself part of the performer removing any interface or barrier between the musician and the music. The embodiment of the acoustic instrument is what makes this research question important: any interruption in the transparent relationship established between the performer and the acoustic instrument can affect the performance of the classically trained pianist. Pianists in particular, develop through many years of practice a gestural technique that encompasses both sound producing and ancillary gestures (Cadoz & Wanderley, 2000). The combination of the years of practice and the development of a complex gestural technique, result in a personal pianistic fingerprint (Caramiaux et al., 2017). However, in gestural musical interfaces tailored to pianists—usually referred to as digital musical interfaces (DMI) if based on a keyboard design and augmented musical interfaces (AMI) if mounted on a traditional musical instrument—this technique is rarely taken into consideration and used to control sound modulation parameters. In many cases, these interfaces require the performer to learn a new gestural vocabulary. As noted by Nicolls (2010) being able to learn and apply a new gestural vocabulary whilst trying to retain pianistic freedom and control is one of the hardest tasks when performing with gestural musical interfaces. For this reason, this research investigates the possibility of creating interfaces for musicians that have already honed and crafted their technique over many years of practice and should not be required to disregard it to explore new sonic environments.

1.2 Aims and Objectives

The aim of this dissertation is to examine how ancillary gestural nuances used within traditional pianistic technique can be applied to an expressive keyboard-based gestural musical interface. The research questions of this project are:

1. Can classically trained pianists use ancillary gestures to control and transform the sound of their instruments and how?
2. Can ancillary gestures improve the learnability of gestural musical interfaces?
3. Does the implementation of ancillary gestures have an impact on the creation of new paradigms and extend the potential creative outputs available to the musician?

The first research question aims to explore and expand the concept of barriers to musical creativity when musicians are confronted by new technologies (Bullock et al., 2016). The concept of barriers to musical creativity, also referred to as entry barriers by Bullock et al. (2011), are defined as the off-putting amount of knowledge and skills required to use a new piece of software. This concept is borrowed and adapted with regard to pianists approaching new keyboard-based gestural musical interfaces for the first time. As noted by Nicolls (2010), the requirement for the pianist to learn a new gestural vocabulary to operate a gesturally controlled instrument is the most disruptive element. The ability for musicians to control a gestural musical interface using the technique they already possess would remove the creative barrier and learning curve usually involved with approaching a new interface and let musicians focus solely on the sound exploration and the investigation of the mappings between gestures and sounds produced.

The second question aims to explore the concepts of learnability in gestural musical interfaces taking advantage of ancillary gestures. Specifically, this discussion seeks to explore and understand how the implementation of ancillary gestures already present in the technique of the pianist can be used to control sound modulation parameters thus removing the learning curve usually required when approaching gestural musical interfaces. This exploration sought to explore how a higher focus on the gestural accessibility of the interface could help soften the steep learning curve usually associated with gestural musical interfaces.

The third and last question seeks to investigate through the development and implementation of the Reach system how it could impact the creation of new musical paradigms. As previously stated, the development of Reach aims to lay the grounds for new gestural musical interface development following the same path of employing ancillary gestures as control for sound modulation. Using the results of the first two research questions, this focus aims to take the knowledge acquired and apply it towards the creative aspect of music making. This investigation has the goal of understanding if the approach of making digital environments more accessible to classically trained musicians will lead to making digital music making more creative and expressive.

In the context of this research, the subdivision of musical gestures adopted is the one first identified by Delalande (1988) that subdivided musical gestures into effective, accompanist and symbolic gestures. The gestural nuances analysed in this context are to be categorised as accompanist gestures. These gestures, while not being part of the sound producing gestures, are inseparable from them (Hemery et al., 2015). Gestural nuances in this context are to be considered movements happening before or after the main sound producing movement, the striking of a key with one or more fingers. The word ancillary relates specifically to the role of these gestural nuances in the act of playing the piano: ancillary gestures are gestures not pertaining to the sound producing gesture. These gestures are called *sound facilitating* if they happen before the sound producing gesture, or *accompanying* if they happen after (Godøy & Leman, 2010). The scope of this research was limited to accompanying gestures, specifically the last version of the Reach system was able to recognise three gestures: the lateral swaying of the hands once one or more keys had been pressed, the palm height once one or more keys had been pressed, and the distance of the hands from the keyboard when leaping from one side of the keyboard to the other.

This specific gestural focus was chosen because of its ability to influence expressiveness in pianistic musical performance, and the increased presence of these gestures in expert pianists (Tits et al., 2015) and expressive performances (Winters & Wanderley, 2012). This research aims to investigate whether this layer of mute gesturality, or gestures not directly related to the production of sound, could be recognised and implemented into an augmented acoustic piano to allow the pianist to achieve an intuitive interaction for live sound processing and facilitate new modes of creative expression. The goal was to investigate if control via ancillary gestures could increase the level of expressivity of the performer, the learnability of the interface and the creative output.

It became clear that the transparency of the interface was not defined solely by the means of interaction, but by a broader set of elements. This steered the research to question and investigate more broadly what different aspects of Human Computer Interaction (HCI) and interface interaction could improve and or influence expressivity, learnability, accessibility and creative output in a keyboard-based AMI.

The realisation of the complexity behind the musician-instrument interaction, opened the research towards a more holistic approach, and made its accessibility a priority. This is the reason why this dissertation is being submitted as part written thesis, part augmented instrument. Reach, the augmented piano instrument developed throughout this research project, was created and used as a probing tool to investigate the subject matter and the relationship pianists have with gesturally controlled instruments. While the Reach system is to be considered a proof of concept, it was essential for the empirical tests carried throughout the three years of research. Being able use the Reach system to test research findings in real world scenario with the end users of this research project—the pianists—was essential to keep the research user centred. This approach also enabled to test the findings and compare them to existing keyboard interfaces that approached similar problems in different ways. This allowed to gain insight into the differences between not only different kinds of musical interfaces—AMIs and DMIs—but also on how different means of interaction affect the relationship between the musician and the interface. The revealed complexity of the interaction between the musician and the traditional musical instruments led me to explore more broadly the concept of gestural musical interface.

Exploring the implementation of ancillary gestures to control gestural musical interfaces shed light on the benefits of a novel design paradigm based on pre-existing instrumental technique. These novel interfaces would not require musicians to learn a new set of gestures, and would enable them to use the pre-existing gestural nuances in their instrumental technique. The research conducted in this instance on pianists could be generalised and abstracted to other kinds of musicians. For example, research with similar goals is being conducted on the picking technique of guitarists by Morreale et al. (2019), and on the remapping of existing violin technique gestures by Thorn (2018). The findings of this research will provide an insight on gestural nuances in pianistic instrumental technique and how these can be mapped to sound modulation control parameters.

1.3 Thesis Overview

The thesis, as shown in Figure 1.1, is subdivided in three main parts: an introductory one (Part 1), the empirical studies (Part 2) and the conclusion (Part 3). The first part includes the introduction (Chapter 1) and a literature review of the concepts of interface transparency, gestural sensing modalities, existing keyboard interfaces for live sound modulation and evaluation of the methods employed (Chapter 2). The second part consists in an overview of the first case study and rationale for the choice of the creative practitioners (Chapter 3); the first user test (Chapter 4); the case studies conducted with three different creative practitioners (Chapter 5); and the comparative user testing (Chapter 6). In the third part, conclusions are presented alongside a reflection of future works stemming from this research project including longitudinal studies and exploration of other gestural vocabularies pertaining to other musical instruments (Chapter 7).

The empirical methodology is represented by five empirical studies and the iterative development of the Reach system (Chapters 3 - 6). The choice of an approach based on field explorations, case studies and practice based methods was driven by the lack of existing research on the topic. While the broader field of augmented piano performance or gestural controlled keyboard-based interfaces for live piano sound modulation has been explored by a number of researchers—discussed in detail in Chapter 2—none of the referenced projects address the issue with the same approach of this research project. The studies mentioned focus on reliability of motion tracking devices, upper torso and head movement in piano playing or focus on hand

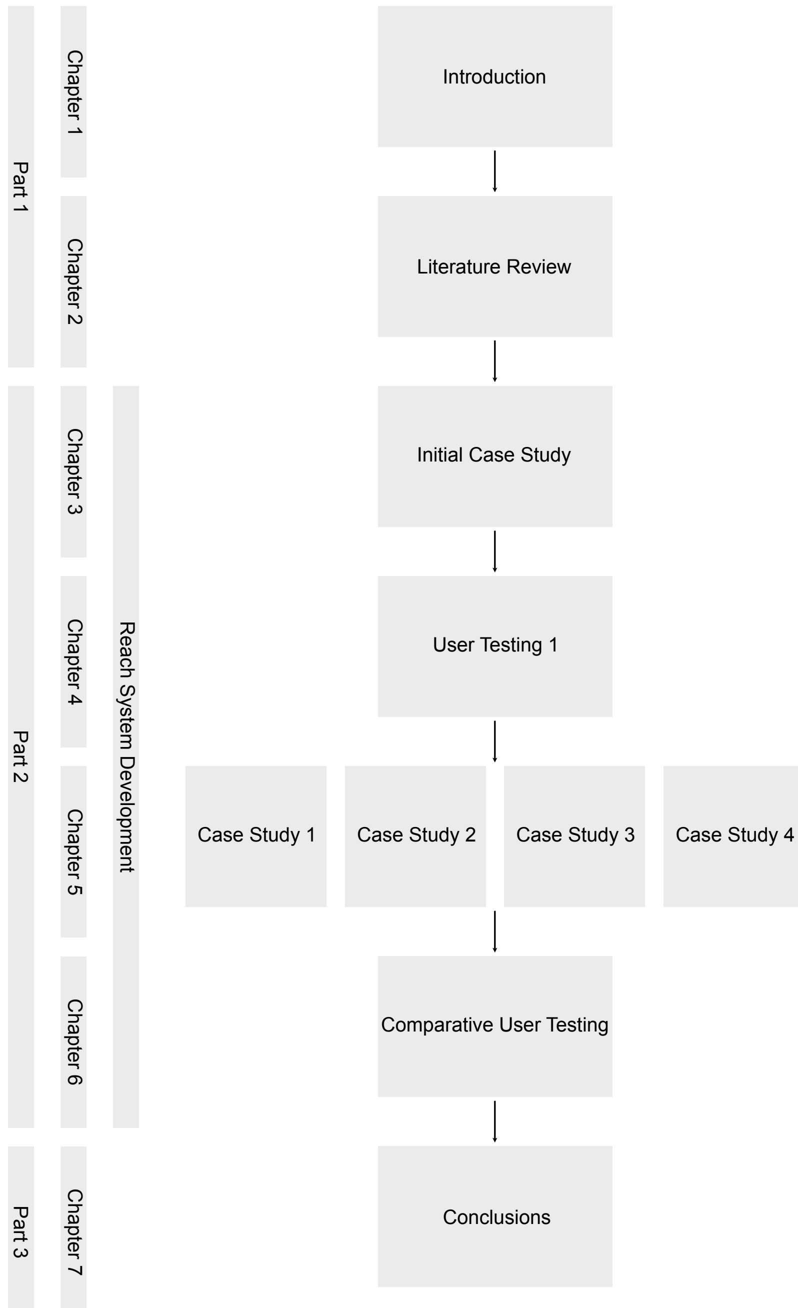

Figure 1.1: Flowchart of the thesis structure.

gestures but they all require the pianist to learn a new gestural vocabulary and play a new instrument rather than an augmented piano. The empirical approach chosen investigates systems tailored to pianists that foster pre-existing gestural nuances present in their technique with the goal to provide a more accessible instrument. The empirical and test driven methodology was useful to gather different sets of data. Quantitative data was gathered through initial task driven testing and the results of the User Experience Questionnaire (UEQ), and qualitative data was gathered through interviews with the participants and creative practitioners as well as through the recording and analysis of public performances, compositions and improvisations with the system.

The final chapter brings together the results from the empirical methods answering the research questions and providing a better understanding on the research field. However, while in the conclusion chapter (Chapter 7) the research is summed up, the same concepts are weaved in the discussion sections of each empirical test (Chapters 3 - 6).

1.3.1 Part 1

The literature review (Chapter 2) discusses three core topics surrounding this research project. The first, relates to the concept of *interface transparency* and sets the foundation for the exploration conducted in the empirical studies. An understanding of the concept of transparency in musical interfaces and its importance is key to frame the research project and to understand the reasoning behind the investigation around non-invasive musical interfaces. With regard to the field exploration and the development of Reach, a thorough review of the current sensing technologies is undertaken. The evidence gained supports the choice of using camera vision in the form of the Leap Motion[1] sensor to build the Reach system. The third section of the chapter reviews current keyboard based interfaces that implement gestures to control sound modulation parameters. While this review tries to be as thorough as possible covering both commercial products and prototypes, the focus on real world solutions helps frame the choice of the interfaces used as mean of comparison in the comparative user testing (Chapter 6). The final section of the chapter reviews the methods used for the evaluation of the hypothesis and Reach system itself. Alongside the review of the methods, this section grounds them in the surrounding literature providing context and reasoning behind the interdisciplinary approach used.

1.3.2 Part 2

The Action Research Spiral (Hayes, 2011) is the underlying framework on which the bespoke research methodology is built. The detailed description of the methodology is split into each chapter providing the reader with an overview of the methods used. In a similar manner, the subsequent developments of the Reach system are explained at the beginning of each chapter.

All the results and feedback gathered from the five empirical tests (Chapters 3 - 6) contribute to a better understanding of the implementation of gestural nuances for the control of live audio parameters in keyboard-based digital and augmented instruments.

[1] https://www.leapmotion.com

The initial case study (Chapter 3) explains how the first iteration of the Reach prototype was tested in a real world scenario. Rationale and results are discussed including the possible bias of the author performing the test himself.

Both the user testing and the Comparative User Testing (Chapter 4 & Chapter 6) put the Reach system in a lab scenario to test theories and hypotheses with traditionally trained pianists.

All of the case studies (Chapter 5) had the goal of investigating the system being used in a real world scenario through public performances and recordings, and exploring potential users of the system by approaching different creative practitioners in the field.

1.3.3 Part 3

In the conclusion (Chapter 7) the results of the research are collected, summarised and discussed, while the practices, limitations and successes of the chosen approach and methods are reviewed. This research project comprises three parts: the written thesis, the results of the empirical studies (public performances, recordings and compositions) and the Reach system.

The written thesis contributes to the knowledge around gesturally controlled musical instruments that make use of the gestural nuances present in acoustic instrumental technique. Using Reach as an instrumental probe, the thesis provides insight on advantages and disadvantages of these gestural interfaces in comparison with commercial systems already available to the general public. It is expected that this research will contribute to further developments in the field bringing attention and focus on the need to build digital environments with musicians in mind.

The results of the empirical studies give proof of the effectiveness of such systems and provide visual material for future research projects to draw upon.

The Reach system is a concrete outcome that can benefit musicians in the immediate future and a stepping stone for other researchers to keep investigating the field.

Chapter 2

Literature Review

This chapter analyses the conceptual and philosophical underpinning of interface transparency, which is defined as the ability of an interface to disappear to the user and enable an embodied interaction. It also investigates the purely technical aspect of sensing modalities, or technological means to detect gestures, and the existing pragmatic solutions to the researched problem of existing interfaces. An overview of the methods used and the grounding in the field is also provided as a final section to the chapter.

Section 2.1 focuses on the concept of interface transparency, and the correlation with the ability of musicians to become proficient and convey musical expressiveness through the instrument itself. Instrument transparency is analysed alongside the concept of instrument embodiment providing insight on the ability of musicians to become one with an instrument. Section 2.2 explores the existing keyboard interfaces that enable live sound modulation through gestural interaction. Section 2.3 explores the technical means through which the hand gestures of the pianist can be tracked and recognised maintaining a transparent interface. Section 2.4 analyses the combination of qualitative and quantitative data gathering methods used in this research project, providing contextualisation and references in the field of research.

2.1 Transparency in Musical Instruments

The concept of transparency applied to digital interfaces has been defined in a number of ways, from being a measure for the usability of computer systems (Wixon & Good, 1987; Sturman, 1992) to being considered as a design goal by Bolter & Gromala (2003). The concept of transparency applied to musical interfaces has been thoroughly investigated by researchers in the field of music technology (Fels et al., 2002; Jordà, 2005; Nijs et al., 2009; Magnusson, 2010; Jack et al., 2018) who have tried to define which aspects of the instrument—or of the relationship between the musician and the interface—make it transparent.

The instrument transparency explored in relation to a musical instrument, can be analysed and explained in parallel with the concept of tool transparency used by Leman (2008) in the context of embodied music cognition. Leman explains that in order to have a high level of transparency, an instrument must enable the musician to focus on higher-level musical intentions instead of individual operations of manipulating the instrument itself. This, however, does not exclude non transparent relationships in the context of creative musical explorations. An instrument that does not provide a transparent relationship could be a source of inspiration. The element that in the context of this research is seen as an obstacle could be seen as a creative spark enabling a dialogue between the musician and the instrument. Contextualised in this research project, the concept of transparency is key to enabling the musician to focus on high level musical intentions and explore the musical potential that, traditionally, the instrument has to offer. In this context, transparency is the ability of the instrument to *disappear* in the interaction providing the musician with the ability to interact directly with the music. This interaction can be considered as a direct connection between the intention of the musician and the resulting music: a transparent instrument will enable this direct transfer of information without interfering. Nijs et al. (2009) formulates the concept of transparency by visualising it as the merging of the musician and the musical instrument. More specifically, Nijs applies the concept of interiorisation of the instrument (Nardi, 1996; Nosulenko et al., 2005) to musical instruments, describing the role of the transparent musical instrument as an *organic component of the performer's body*.

The explanation given by Nijs (2017) refers to how Rabardel (1995) breaks down interface transparency into two distinct factors: *relational transparency* and *functional transparency*. Rabardel speaks about *the operator* of the tool, but in this case the subject will be intended as the musician. Relational transparency enables the musician to interact directly with the musical environment without any conscious cognitive effort, relying only on the previously acquired and now embodied skills. This is because, as noted by Pestova (2008: 15): "(...) musicians interact with each other, their instruments, the audience and the acoustic properties of the hall". Functional transparency instead happens when the musical instrument does not interfere with the relationship between the musician and the musical environment and the musician has the perception of interacting with the music itself. The manifestation of these two different kinds of transparency can be observed in the traditional performance scenario. When the performer is able to influence the performance by reacting to the musical environment (i.e. the presence of the audience, the feeling of the instrument) focusing solely on the musical outcome of the performance, we are witnessing a fully transparent relationship taking place. The musician is not only able to interact with the musical environment without interference from the instrument, but is also able to use specific situations (e.g., the concert hall atmosphere, audience feedback and presence) to inspire and affect the performance. The instrument does not prevent the musician from reacting to particular cues in an expressive way. However, the ability of focusing on higher-level musical intentions instead of individual operations or to react to particular cues is not always enabled solely by the transparency of the instrument. The transparency of the musical instrument is enabled by the combination of the features of the instrument itself and the familiarity that the musician has with the instrument. The ability for a tool to achieve a high level of transparency requires familiarity with the tool, and in the case of musical instruments, familiarity can be achieved only through years of study. The fact that musical instruments are embedded in musical culture justifies the amount of time spent practicing by the performer to effectively master and embody the instrument thus rendering it transparent.

The importance of instrument transparency tied with the historical and cultural background of the instrument itself can be explained with Magnusson's interpretation of Heidegger's concept of tools considered 'ready-to-hand' (Heidegger, 1962; Dourish, 2004) applied to musical instruments. Magnusson ties the Heideggerian concept of 'ready-to-hand' tool with the embodied nature of skill acquisition related to learning an acoustic instrument. When the focus can be directed towards the action that the tool enables—and not the use of the tool itself—the instrument becomes 'ready-to-hand'. The 'ready-to-hand' status for Heidegger can be achieved only when the user is familiar with the tool. Magnusson (2009) states that musicians gradually learn instrumental technique through persistent practice and thus transforming the tools in question—the musical instruments—into 'ready-to-hand' phenomena. In the eyes of the performer, the transformed instrument is not based on symbolic or theoretical relations any more but becomes embodied in the action of playing. The musicians do not think about all the elements that composed their instrumental education throughout the years of study, but focus solely on the act of playing that is inherently a result of the years of study. In a similar way, Leman (2008) proposes that a good technique allows musicians to concentrate on *musicality* thus making the musical instrument transparent. Combining the two definitions, Leman's *good technique* transforms the instrument into a 'ready-to-hand' phenomenon enabling the musician to perceive the instrument as transparent and focus solely on the musicality of the performance.

Gesturally controlled digital and electronic keyboard instruments instead provide an experience that is more 'present-at-hand' (Dourish, 2004) than 'ready-to-hand'. The 'present-at-hand' experience happens when the musician is presented with a new interface, that requires a new set of skills to be acquired and is ever-changing for the very nature of it: mappings and sounds are usually dynamic and not tied to the instrument. While the keyboard technique makes a traditional acoustic piano 'ready-to-hand' and enables the musician to focus on higher-level musical intentions, gesturally controlled interfaces make the musician aware of the interface and bring the focus back to the simple interaction with it, due to the unfamiliarity of the technique required (Dourish, 2004). While theoretically this status could progress to a 'ready-to-hand' relationship through years of practice, the lack of longevity of these instruments (Morreale & McPherson, 2017) combined with the particular nature of the interaction applicable only to one specific instrument, makes it harder for musicians to achieve the same level of familiarity of an acoustic instrument. The more we dive deeper in the concept of instrument transparency, the more the concept seems to be bound to a relation and not to the object itself. While it is true that instruments can be transparent in their physicality, without a transparent relationship with the musician this physical transparency is not enough to provide a transparent experience.

The concept of instrument transparency delineated until this point and the consequently developed tight relationship with the instrument seems to draw firmly from the concept of control intimacy theorised by Moore (1988). Moore sees control intimacy as the match between the multiplicity of musical sounds produced by the instrument and the performer's psychophysiological capabilities to perform those sounds. For the author, control intimacy is a delicate balance between the variety of possible sounds that the instrument can produce and the potential of the musicians themselves to perform said sounds with their own abilities aided by the feedback produced by the instrument itself. This intimacy can be achieved through years of instrumental practice. As noted by Nijs et al. (2009) to achieve a transparent relationship, musicians must overcome the limitations and constraints of the instrument. With keyboard-based digital or electronic musical instruments it is harder to obtain control intimacy because they provide the user with an infinite amount of possible sounds. Magnusson (2009)—bringing as an example *The Hands* (Waisvisz, 1985)—explains how a change in the algorithms that were part of the sound engine resulted in the requirement for the creator/performer to learn a new instrument, having to re-incorporate the conceptual understanding of the system's functionality into his bodily memory even though the interface itself remained the same.

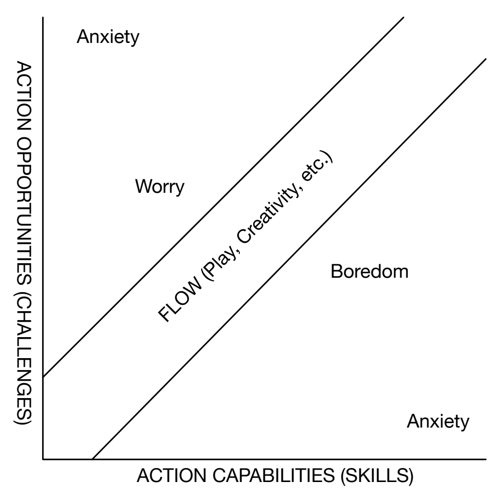

Figure 2.1: FLOW: Action Opportunities v Action Capabilities (Csikszentmihalyi, 2014b: p.147)

The concept of control intimacy applied to a musical instrument is generalised and applied in various situations by Csikszentmihalyi (2014a) as *Flow* status. Csikszentmihalyi depicts the concept theorised by Moore as shown in graph 2.1 (Csikszentmihalyi, 2014b: p.147). This graph places on the horizontal axis the skills involved in a certain task, also called action capabilities, and on the vertical axis the challenges that the task involves, also called action opportunities. The control intimacy discussed by Moore, *flow* for Csikszentmihalyi, is the area in which these two values are balanced or reach a point that is very close to an equilibrium. This state of flow results in being a desirable state for expert performers, in which the instrument becomes completely transparent and only the action of playing becomes apparent. The instrument becomes 'ready-to-hand', enabling the pianist through embodiment to use the piano as an extension of their own body.

The time spent working within the limitations of their instrument brings pianists to refine more and more their ability to control the sonic output of the piano and develop a relationship with the instrument itself. Due to the interplay of sound facilitating, producing and accompanying movements (Kanga, Nunn 2015), pianists are able to develop a proper and unique acoustic signature on the instrument (Chadefaux et al., 2010). This has proven to be recognisable and categorisable both by humans and machines with a high level of accuracy. Personal pianistic fingerprint, in the same study, is directly related to the amount of years of practice (Caramiaux et al., 2017). The combination of years of practice and technique development create a unique relationship between the instrumentalist and the instrument.

The same cannot be said for DMIs. The strong bond between physical and aural identity of traditional musical instruments does not apply to digital interfaces whose gestural interaction and timbral identity are not defined by their physical properties.

Magnusson (2009), starting from the concept formulated by Ihde (1979) of acoustic instruments being an *illustrative example* of technology that enables embodied relations to the world, follows describing DMIs as more of an extension of the mind rather than the body. A digital musical interface tends to be very versatile with no fixed sound, and this versatility makes digital musical interfaces hard to embody. But the problem is rooted even deeper in the concept of interface. As Norman (1990) states in his publication with the self-explanatory title '*Why Interfaces Don't Work*', interfaces get in the way of interaction by definition: their role is to interface two ends of an interaction.

It is possible to apply Norman's concept to the acoustic piano, as it can be seen as an interface. Acoustic pianos—and more generally traditional musical instruments—present the musician with an interface through which they can interact with a musical environment. The piano, for example, enables the pianist to interact with the musical environment through an interface composed by 88 weighted keys. The pianist, by understanding how the interface works and becoming proficient in the use of it, is able to interact with the musical environment in an expressive manner and produce a musical outcome. However, musical instruments avoid this looming categorisation because they are engraved in musical and historical culture. The musical culture in which traditional musical instrument are engraved, makes the years of practice on them justifiable, transforming them for the performers from interfaces to tools. As explained by Magnusson (2009) borrowing the concept of object naturalisation from Latour (1987) and Bowker & Star (2000), musical instruments are to be considered objects that through heavy use loose the connection with their origins. This process of naturalisation, makes the instrument sink in the community's *routinely forgotten memory*. This does not happen with digital and electronic musical instruments because—as found by Morreale & McPherson (2017)— these instruments rarely establish themselves in the musical community thus preventing a naturalisation process to commence.

When talking about DMIs instead, Norman's definition of interfaces is quite appropriate. The problem with interfaces is that they are interfaces, as stated by Norman. They interface a subject, the musician, with an object, the musical interface. Digital musical instruments (*or interfaces*), in this specific case digital keyboard instruments, add a digital layer between the action and the sound production. This added step, at varying degrees, can be disruptive for performers. For example, if this digital layer required the performer to reinvent their instrumental technique, as found by Nicolls (2010), trying at the same time to retain pianistic control and freedom, would result in a disruptive experience. Grandhi et al. (2011) redefine the concept of naturalness of an interface shifting the focus from the system to the interaction. In doing so they define the naturalness of a gesturally controlled interface as *'marked by spontaneity'*. This ties with Nicolls' findings that found the learning of a new gestural vocabulary—or lack of spontaneity in the gestures required—as a disruptive element when trying to play the piano retaining instrumental virtuosity. Viewing the same problem from a different angle, Aguilera et al. (2006) note that unnatural interfaces are usually the result of bad implementation and design. This could be caused by a series of different factors, one of which is the building of the interface around what the designer *thinks the user wants* rather than focusing on the effective needs of the end user. This has been found to cause a fabrication of a new type of performer that combines the creator, the composer and the performer (Michailidis, 2016). Morreale et al. (2018) in a paper survey around new interfaces for musical expression (NIME) titled '*NIME Identity from the Performer's Perspective*' found that 78% of the 61 performers taken into consideration designed the NIME that they play, 97% of which had been involved in the process of instrument making. This new figure that combines the composer, the performer and the instrument creator seems to be prominent in the NIME community and responsible for a large part of the most recent NIMEs. However, the end result of these practices ultimately leads to the DMI being intuitive only to the creator.

Having analysed the concept of transparency in musical instruments—both acoustic and digital—the features of an instrument that enable its transparent relationship with the musician cannot be limited to the interaction. Through the works analysed, it emerged that musicians are able to achieve a transparent relationship and interaction with a musical instrument due to different factors. While the affordances of the musical instrument itself contribute to the transparency of the resulting relationship, the social and historical background of the instrument plays a major part in enabling a transparent interaction because it justifies the years of practice required to become proficient on the instrument. The interplay of these three factors—the interface, the historical background and the proficiency of the musician—provide a fertile ground for a transparent relationship to arise. If one of these factors is missing, it is harder for a musician to establish a transparent relationship with the instrument. In the following section, existing musical interfaces will be analysed to explore how the aspects of transparency just explored have been taken into consideration in the development of gesturally controlled keyboard-instruments.

2.2 Existing Interfaces

A multiplicity of keyboard systems that enable various kinds of sound modulation through the use of different kinds of gestures are available. Since the creation of the MIDI aftertouch technology in the 1980s, there have been multiple attempts to add different expressive layers to keyboard instruments. The following is a critical list of all the major systems—taking into consideration both research prototypes and commercial systems—that will help frame the scope of the research. It is worth stating that, to limit the scope of this study, only non-wearable means of interaction with sound modulation parameters and systems obtainable by the general public will be considered. This choice was made to maintain a consistent level of interface transparency throughout the chosen keyboards and to avoid introducing another variable——the wearability aspect of the interface——that could have affected the relationship between the musician, the instrument and the resulting music. The restriction regarding the obtainability of the systems removed some existing solutions from the scope of this research, such as the Magnetic Resonator Piano (McPherson, 2010) that enables gestural control of the sound by analysing the movement of the piano keys.

2.2.1 Keyboards

The Haken Continuum Fingerboard (Haken et al., 1992) and the Rolky (Johnstone, 1985) can be considered pioneering devices in this field. The two interfaces approached the gestural sound modulations in two different ways. The Haken Contin-

Figure 2.2: Haken Continuum

uum Fingerboard, as shown in 2.2, presented the artist with a keyboard-like pattern drawn on a continuous fabric surface. On this surface, the x, y and z coordinates of 10 independent fingers are tracked and used to control the pitch, loudness and timbre of the notes played. Alternatively the Rolky, used a transparent glass slate and a light to detect the number and position of fingers to enable polyphonic interaction.[1]

Both these interfaces had a limited amount of tactile information regarding the location of the fingers, and did not manage to provide an intuitive way to enable polyphonic pitch-bending while also enabling effective tuned playing (Lamb & Robertson, 2011). In the first instance, the Continuum Fingerboard had no moving keys. This lack of movement failed to provide useful tactile feedback to the user that would

[1] The Rolky was only ever released as a research paper (Johnstone, 1985) and no physical prototype was ever documented.

have enabled the musician to know the position of the fingers without constant observation of the instrument. On the other hand the Rolky, while conceived with the polyphonic touch concept in mind, was mostly seen as a controller, not an instrument. Neither of these interfaces provided a musically intuitive way of playing and performing polyphonic pitch-bending while also enabling the musician to play in tune (Lamb & Robertson, 2011). Both interfaces require users to learn a completely new gestural vocabulary due to the unfamiliarity of the musical interface. Whilst the Rolky remained in a prototype stage, the Haken Continuum Fingerboard was produced commercially and is still being sold to this day. However, when looking at musicians using the instrument in their musical practice, it is clear that the instrument is used more as an expressive sampler, or sound bank, rather than a creative iteration of a keyboard instrument (Tobin, 2018).

In the past twenty years, with the rise of computing power and the lowering of the prices for the average consumer, sensor technologies became more and more affordable and easy to interface with the everyday computing machine. As a result, more augmented piano performances using depth sensor cameras, such as the work from Yang & Essl (2012) or Nicolls & Gillian (2012), and interfaces such as the Elbow Piano (Hadjakos et al., 2008) were developed. While the Elbow Piano, and the successive studies from Hadjakos such as the Vibrato Piano (ViP) (Hadjakos & Waloschek, 2014) have an element of instrument transparency, both present intrusive elements that make them incomparable to other interfaces that strive to be completely transparent. To use the Elbow Piano, the performer must attach a goniometers to each arm. The elbow angle provided by the goniometers combined with a webcam placed perpendicularly to the keyboard enable the user to have control over the sound produced by the keyboard. Similarly, the ViP provides control over the vibrato of the notes played on a MIDI enabled keyboard through two wrist sensors worn by the pianist on each hand. While both these DMIs are worth mentioning when discussing gestural musical interface advancements in the world of keyboard instruments, they require a wearable element that makes them harder to use when trying to retain pianistic freedom and proficiency. The systems built by Yang & Essl (2012) and Nicolls & Gillian (2012) instead are closely related to this research. While the aims of the research projects might differ drastically, the research focuses on the problem of manipulating sound parameters on keyboard interfaces.

In the first case, Yang and Essl describe traditional keyboard controls—such as knobs and faders—as unintuitive and different, and thus the approach of using gesture recognition was chosen. The choice of augmenting the piano keyboard, rather than building an instrument from scratch, was approached as an "easy-to-learn pitch control over a purely gesture-based interface" (Yang & Essl, 2012: 1). While these premises align with this dissertation, there is a difference between these two research projects. All of the gestures required to control the system created by Yang and Essl are extra musical. While the research is based on the notion that a keyboard provides an easy way of playing notes with a precise pitch, it does not seem to be tailored towards instrumentalists that want to port their pianistic technique to a new instrument.

In the second case, Nicolls & Gillian (2012) propose a system for live piano sound modulation that enables the performer to control sound modulation parameters effecting the live piano sound through the use of skeletal tracking and body-gesture recognition. This system provides some interesting insights both on the aural aspect and on the gestural component. Aurally, the system focuses solely on the sound of an acoustic piano. This choice, as explained by the authors, enables the musician to perform a *totally instinctive improvisation* with the system. The authors, while not explicitly stating it in the paper, developed a system that enabled the performer to establish a transparent and embodied relationship with the sound. In a similar way, when approaching the skeletal tracking the authors decided to focus solely on ancillary and *realistic* gestures. They justify the choice of these gestures by referring to two key aspects: the desire to avoid the need to learn a new set of gestures and the ability of the pianist to move effortlessly between instrumental gestures and ancillary gestures without having to make any conscious change in her physical state. The system created and performed by Nicolls and Gillian at the 2011 International Conference for New Interfaces for Musical Expression (NIME) laid the ground for Brent's Gesturally Extended Piano (GEP) (Brent, 2012).

The GEP is a system that enables gestural interaction with audio-visuals by analysing the data coming from a high speed USB camera, an array of IR sensors and some trackers placed on the hands of the pianist. However, the hardware used needed hand markers in order to have stable hand tracking. The study and interface never made it to a more advanced stage, but what had been published in the initial paper laid the grounds for the present research project. The prototype developed by Brent (2012) had four key characteristics. The system presented was both minimally intrusive, requiring just a small strip of IR sensors and four trackers placed on the hands and arms of the pianist, and enabled the pianist to play on an acoustic piano. The augmented nature of this system enabled pianists to use their existing technique to control the sound modulation parameters. Moreover, the portable lightweight hardware made the system ready for a performance scenario and the low cost of the components resulted in a highly accessible instrument for pianists looking to embed it in their performance.

Brent (2012) illustrated the conclusion of his work and the next steps for his research, underlining how the above mentioned characteristics would have remained a central focus of the newer iterations of the system. While focusing on improving the accuracy, tracking and general performance of the system, Brent stated that improvements in ease of setup and operation would have continued to be prioritised, underlining the importance of such a feature in a system that focuses on performance. Future tests of the system were meant to be based on performances to shift the focus more on the sound and movements associated with the real-world scenario of an acoustic piano performance. All the mentioned characteristics appear to be essential to reduce the learning curve of an AMI, and improving the accessibility and creative output for a performer. Without performing a user study specifically tailoring pianists and performers, the research that Brent conducted lay the ground for keyboard-based instrument development in the following eight years.

Having said this, the two most popular gesturally controlled keyboard-based interfaces do not use the principles followed by Brent or Nicolls & Gillian (2012). Both the ROLI Seaboard (Lamb & Robertson, 2011) and TouchKeys (McPherson, 2017), require the user to learn and utilise additional techniques.

Figure 2.3: Seaboard Rise 49

The Seaboard (Figure 2.3), as described by its creator Roland Lamb, is a musical instrument that enables real time continuous polyphonic control of pitch, amplitude and timbre. Based on the concept of a keyboard, the Seaboard shapes an entire slate of silicone to resemble a traditional keyboard. Alongside traditional "*note-on*" and "*note-off*" MIDI capabilities, on this keyboard-shaped continuous plane the position, pressure, and movement of the fingers can be tracked and mapped to control individual sound parameters. The sound can be generated through *Equator*, the provided software, or any MIDI polyphonic expression (MPE)[2] compatible software.

Similarly the TouchKeys (Figure 2.4) provides a sort of key-coating that can be applied to keyboards, with a touch capacitive sleeve that enables the individual detection of the fingers along the length of the keys, enabling the control of different parameters passing through the provided software into any musical environment.

[2]MPE - https://www.midi.org/articles-old/midi-polyphonic-expression-mpe

Figure 2.4: TouchKeys mounted on a Novation Impulse 49

Even though the Seaboard changes the keyboard entirely providing a different feel to the musician, both keyboards present the musician with a similar kind of gestural interaction. Of the traditional pianistic technique, only the sound producing gesture is used to play the instrument. This main gesture is enhanced by implementing individual note pitch bending capabilities, and multiple other sound modulations, all controlled by the information gathered from the fingers in contact with the keyboard. At the same time, both instruments tend to disassemble a familiar pianistic technique into time-dependent gestures. This means that they isolate the sound producing gesture of pressing a key, and build a new set of gestures around that movement. The process of learning a new gesture vocabulary, whilst trying to retain pianistic technique and freedom of playing, can result in a disruptive experience for the expert pianist (Nicolls, 2010). TouchKeys is seen and used as a keyboard-based instrument, while the Seaboard has been described as an instrument not aimed at keyboard players due to its wedge-shaped key design that prevents any advanced keyboard technique to be used (Dahlstedt, 2017). Whilst the implementation of the technology in these innovative interfaces must be acknowledged, this research aims to address the steep learning curve that pianists are faced with when approaching these musical interfaces.

This section explored a series of existing interfaces that enable gestural sound modulation on keyboard-based instruments. Having explored both prototypes and commercial systems, the technical ways in which these interfaces enabled the control over sound modulation parameters are multiple and unexplored until this point. The following section, will provide an in-depth analysis about the sensing modalities available helping the reader understand the different ways in which gestures can be tracked and recognised.

2.3 Sensing Modalities

There are different ways in which gestures can be sensed for musical applications. The sensors can be grouped in different categories depending on the technology supporting the gesture recognition and the most suitable application for them. In this section, a categorisation and analysis of different modalities and the main sensors representing each category will be conducted. The problems of existing sensors and devices will be analysed alongside the role of the interface for musical applications in support of the choice made for the interface developed. The sensors will be initially divided by sensing modality in macro-categories, and then analysed taking into consideration different aspects that would potentially affect the development

Figure 2.5: Graph representation of the sensing modalities reviewed.

of the interface, longevity and availability of the resulting interactive instrument. The macro-categories used to define the sensing modalities in this chapter will be non-camera-based (Section 2.3.1) and camera based (Section 2.3.2). Camera based devices use a video feed from a camera to detect gestures while non-camera-based devices rely on other sensing technologies (e.g., radar, WiFi).

2.3.1 Non-Camera-Based Devices

Non-camera-based devices rely on sensing technologies such as high frequency radars, WiFi signals and infrared light to analyse movement and recognise gestures. Recently, with the increasing demand for smaller devices capable to recognise gestures, this category has been led by the development of the Google Soli sensor. However, due to the limited availability of the Soli, in the academic environment this category is led by gesture recognition through the analysis of WiFi signals, GSM disturbances and the use of passive infrared (PIR) sensors.

Google's Soli, a radar-based sensor for gesture detection, uses millimetre-wave radar to detect fine grain and microscopic gestures with modulated pulses emitted at frequencies between 1-10 kHz. Lien et al. (2016) provides a thorough technical description of the Soli examining its hardware, software, and design. The strength of a radar-based signal lies in its ability to offer a high temporal resolution, the ability to work through specific materials such as cloth and plastic and to perform independently of environmental lighting conditions (Bernardo et al., 2017). One significant feature of the Soli is the highly optimised hardware and software that prioritises motion over spatial or static poses. In addition, the compact size makes it an excellent choice for musical purposes that require non-invasive interaction. Soli uses high-frequency radar with a central frequency of 60 GHz that increases the device's level of accuracy, making it suitable for fine-grain gesture sensing (Wang et al., 2016). The development of the project Soli was blocked in the early stages of its alpha developer-only release, making it a niche sensor that only a few people in the world owned. After a hiatus of two years, the Soli sensor has been speculated to be in the latest release of Google's Pixel phone[3] possibly providing users in the future to interact with the sensor through their phones. Even if these pragmatic issues were not to be considered, initial research with the Soli sensor brought a number of different issues to the surface. One of the biggest issues with the sensor was its ability to detect movement happening behind physical barriers. While radar detection has

[3] https://www.extremetech.com/mobile/295413-pixel-4-might-support-soli-touchless-gestures

the ability to work through specific materials, it does not have the ability to work through body parts. This meant that initial tests done with the Soli placed in a hidden spot on the right side of the keyboard, resulted in the users having control over the system only with their right hand, as seen in Figure 2.6.

Figure 2.6: Eagle-eye view of the piano keyboard with Google's Soli hidden under a piece of cloth on the right hand side.

Other devices that enable gestural input using radar-like detection are the SideSwipe (Zhao et al., 2014), which analyses disturbances of GSM signals and the WiSee (Pu et al., 2013) that analyses existing WiFi signals and their perturbances to recognise human gestures. Both are not suitable for tackling musical gestures due to the lack of precision when trying to analyse small hand movements.

In this same macro-category we can also find a series of non-camera-based sensors that, with the help of an embedded board, can be used to feed gestural data into any program dedicated to live sound modulation. These sensors can be divided into two categories: ultrasonic range detectors and PIR sensors. While the first category had been discarded immediately due to the lack of ability to recognise gestures, and effectively being able to only measure the distance between an object and the sensor itself within a very restricted view angle, PIR sensors result in being more suitable for gesture recognition. A PIR sensor measures infrared light radiating from objects in its field of view. While Raykov et al. (2016) found that through the use of machine learning the data coming from PIR sensors can be parsed to recognise complex motion data—specifically detecting and counting humans passing through a door—these sensors have been mainly employed in alarm systems, lighting systems and hand dryers. The limited use of these sensors is due to the highly accurate detection of movement but lack of specific data regarding the moving object because of the inner workings of the circuitry. PIR sensors work by analysing the difference in infrared light between two small sensors placed close to each other. When one sensor detects a change in quantity of infrared light, the sensor detects that an object has placed itself in the sensors field of view thus detecting movement. Informal testing found that due to the lack of ability to track hand movement the sensors were not suitable for gestural detection.

Overall, while there are other means of recognising hand gestural nuances, with the price and availability of camera-based peripherals and considered the precision of these sensors, there is almost no reason why one would choose the first over the latter if the goal is to perform hand tracking.

2.3.2 Camera-Based Devices

Camera-based devices rely on computer vision to collect data. As found out by Wilson (2010) while they provide less precision for touch based events than physical sensors, they have been found to have other advantages. Wilson found that alongside the possibility of having an non-instrumented and non-flat surface to detect movement on, camera based systems provided the ability of using information from the hands to detect hovering states, multiple touches and assigning actions to specific users.

Initially, due to the environmental limitations that the available technologies had (light-dependant tracking, cost, limited vision of non-wide angle lenses) after initial informal testing, they were considered not suitable for the purpose of this research.

In the past three years, with the advent of virtual reality (VR) gaming and the necessity of tracking hands to interact into virtual spaces, these technologies have had a massive spike in development. Currently this category is led by three main peripherals: Microsoft Kinect, Leap Motion and Intel RealSense.

2.3.2.1 Microsoft Kinect

The Microsoft Kinect[4] has been discontinued in 2017, making it virtually impossible to connect it to newer consoles, but it is still used in the world of audio-visual art for projects requiring skeletal and face tracking. Composed by 3D sensing arrays of cameras (Han et al., 2013), the Microsoft Kinect provided a technology to consumers which had previously been unavailable or unrealistic (Sung, 2011). The system was

[4] https://developer.microsoft.com/en-us/windows/kinect

built to recognise both full body figures and facial traits. Even not taking into consideration the fact that the peripheral is now discontinued, one major drawback of the system has always been its lack of resolution. While the technology was able to provide full body and skeletal tracking, it has proven to only have a 5% success rate of recognising hand or finger gestures (Cronin, 2014).

2.3.2.2 Leap Motion

The Leap Motion[5] was initially built specifically for hand gesture recognition. However, informal tests found that its seemingly obsolete hardware combined with an old VR software development kit (SDK), made the sensor not suitable for recognising gestures on a piano keyboard. The field of view of the peripheral was limited to a small box above the sensor between three and eight cubic feet (Cronin, 2014) that made it hard to use with the full span of an 88-key keyboard.

In 2016 Leap Motion (now Ultrahaptics) released a new version of their SDK, called Orion. The new SDK was limited to a software update, and used the same hardware Leap Motion had released in 2013. The SDK provided greater robustness, reduction of ambient light interference, lower latency, and better tracking on the edge of the field of view among many others features.[6] The same hardware with the new Orion SDK, built purposely for VR gaming, had improved performance, stability and tracking as promised. The fact that environmental related factors did not seem to be a problem anymore, re-opened the door towards camera-based systems in general, being able to move away from the constriction of having only one possible choice.

[5] https://www.leapmotion.com

[6] https://developer.leapmotion.com/orion#105

The inconsistent data output caused by specific scenarios—obstruction of the field of view, the close proximity of the hands or different orientation of the hands—that Cronin (2014) found in his analysis was not as prominent anymore, making the peripheral much more usable. And last but not least, the inability of complex gesture recognition (Cronin, 2014) could now be fixed through software analysis of the data that was much more precise and accurate.

2.3.2.3 Intel RealSense

The Intel RealSense[7] was released in 2015, as a VR Leap Motion competitor. While it promised very good hand tracking, from the specifications and the multiple comparisons online, the RealSense was purposely built for a VR gaming environment, thus making it less versatile in other situations. The combination of camera-based technology, IR depth sensing and microphones (for the base models) including voice recognition and IMU sensors (for the top models) identified these Intel peripherals to be purposely built for a multimodal VR environment. Moreover, informal testing and research on online VR gaming forums showed that the LeapMotion was always the preferred peripheral when referring to hand tracking.

Other systems are also able to identify gestures. This includes colour detection from 2D RGB cameras (Erol et al., 2007). However, such technologies often lack in precision when aimed for fine grain hand-gesture detection.

Having explored the different kind of gesture-sensing modalities available, camera based systems have been identified as the strongest elements for a combination

[7] https://software.intel.com/en-us/realsense

of availability of the technology, sensing precision and low sensitivity to ambient factors. Thanks to the improvement of the tracking algorithms, the effect of these external factors on the performance of camera based tracking systems has been reduced considerably. Having analysed the philosophical underpinning (Section 2.1), existing gesturally-controlled keyboard-based interfaces (Section 2.2) and the sensing modalities responsible for the gesture tracking (Section 2.3), the following section will present an overview of the methods employed and their grounding in the field of music technology.

2.4 Evaluation of Methodology

An interdisciplinary approach that combines quantitative and qualitative data gathering through the implementation of an activity-led methodology (Gay & Hembrooke, 2004) is common practice in the field of music technology and more specifically in the creation of new interfaces for musical expression. The reasoning behind the establishment of mixed research methodologies can be found in the works of Stowell et al. (2009) that explains how the act of live music-making that use interactive systems cannot be limited to a typical HCI methodological evaluation such as task-completion rates. This is because the very nature of music making contains aspects that cannot be quantified and counted numerically, and thus require a combination of different techniques to be evaluated. Stowell et al. (2009) state that the creation of sonic interactive systems does not begin until the systems come into use. It is the use of the interactive system and the interaction with musicians and their experience that creates the affordances and helps contextualise the artefact. The findings of Stowell et al. can be broadened to encompass other artistic practices. In a similar way—to explore the relevance of practice-led research in this context—Bolt (2007: 30) explains David Hockney's exploration of Ingres' painting technique as if the "engagement with the tools and technologies of drawing practice produced its own kind of sight or logic". The authors link the action of acquiring knowledge through practical experience to Heidegger's concept of *circumspection*. In the same way that Heidegger affirms that we can understand the world theoretically only after we have experienced it through practice, Hockney discovers that the only way to understand Ingres' painting techniques was to experience them first-hand. While in the context of this research the goal is not to understand the technique of a specific composer, the investigation of the potential creative output of novel gestural musical interfaces can be seen as the object to be experienced. To experience the musical environment the musician needs to experience first-hand the tools—the gestural musical interface—feeding back to the researcher with the acquired theoretical knowledge.

2.4.1 Commissioned Works and Public Facing Performances

As shown by Green (2014), the NIME community has contributed greatly to enriching practice-led research. These interdisciplinary approaches help the researchers move beyond technical concerns—typical of HCI driven evaluations—to develop a firmer and more nuanced understanding of complex questions of musicality. In this interdisciplinary field lying in the midst of music technology, HCI, and music composition, the use of interdisciplinary methodologies employing musical composition and performance as a method of evaluation is common practice (O'Modhrain, 2011; Taylor et al., 2013). The analysis of HCI issues through the framing of digital art practice allowed the researchers to interrogate participant experience by focusing on creativity and aesthetic sensitivity. Moreover, O'Modhrain (2011) explores through the proposal of a new methodological framework for the evaluation of digital musical instruments how the term itself of *evaluation* should be considered in a much broader sense due to the number of stakeholders involved in the creation of said instruments.

Regarding the combination of practice-led research with more conventional research methods and practices, Green (2016) highlights how these methods offer a valuable contribution when it comes to the topic of new instruments for musical expression. Bolt (2007), Green (2014) and Barrett & Bolt (2014) describe the ability of conducting practice-led research alongside quantitative and controlled methods as a way to obtain more holistic insight around the investigated topic. In an effort to combine these methods, an iterative approach has been explored by Brown & Sorensen (2009) that have found that the *feedback loop* between speculation and experimentation is fundamental for research in the creative arts. This iterative cycle— constituting the core of activity theory 3.1— is a key method of exploration in the field of activity-led artistic research. Gaver (2012) in his work exploring the concept of research through design, lays the ground for how in practice-led research, the design of artefacts can come prior to the formulation of an issue and thus carry the rest of the exploration forward. In a similar way, artists and researchers pertaining to the world of NIME, use the instruments developed to scope the field and through a series of iterative designs, performances and reflection on their artistic practice are able to fuse the different elements in one cohesive methodology.

2.4.2 User Testing

As described by Johnston (2011), the mixture between quantitative and qualitative research through the implementation of user studies is a powerful tool that enables both detailed reflection on the creative affordances provided by the design of the musical interface and documentation of the creative outcome itself. In a study dedicated to the exploration of video analysis as a means of evaluation, Xambo et al. (2013) illustrated how the use of video analysis can be a valuable asset for a number of different aspects: from qualitative to quantitative analysis as well as improvisational and participatory perspectives.

2.4.2.1 User Experience Questionnaire

While the use of the User Experience Questionnaire (UEQ) (Laugwitz et al., 2008) has been identified as an invaluable tool to scope and analyse product user experience, both Schrepp et al. (2014) and Di Donato (2018) have successfully explored how the UEQ could be implemented in scenarios outside the evaluation of a commercial product. Di Donato in particular, explored how the UEQ could have been applied in the evaluation of an interactive music system to probe the musicians about the experience they had at a first encounter with the musical environment and gather alongside the qualitative data from the performances, quantitative data. In the original study, (Laugwitz et al., 2008) defined the UEQ as a valid measure for user experience in general, particularly when used to complement data from other more subjective evaluation methods.

2.4.3 Collaborative Design Workshops

The shift from a user-centred research project, to a user-involved one (Gay & Hembrooke, 2004: XVIII) has been identified as one of the key aspects of activity theory. For Gay and Hembrooke, the difference between user-involved and user-centred lies in the focus of the analysis. While both methods keep the user as the main subject of the research, a user centred approach focuses on what the user should do, while a user involved approach tries to understand what the user already does. Framing these definitions in the context of gesturally controlled musical interfaces, we could argue that interfaces that require the musicians to learn a gestural vocabulary in order to control them could be defined as user centred, while interfaces that use what the user already does could be defined as user involved. In the context of this chapter, the focus on the practices of composers and artist in a real world scenario such as a public facing performance, brings attention to what the creative practitioners already do and focuses on how the system impacts on these actions.

In the workshops the creative practitioners are asked to take part in the design of the gestural musical instrument. The insight is initially given as an external observer and then—thanks to the iterative nature of the subsequent workshops—the practitioner becomes acquainted with the interface and thus able to become involved in the design process. The longitudinal nature of the workshops combined with the collaborative design element is an efficient way to gather data on how musicians would ideally create an instrument to suit their needs. The inclusion of the musicians' point of view and direct experience infused the development of the system they would have used to compose and perform. As seen by (Fischer, 1998) to include important perspectives in the process of design, all stakeholders in question should participate in the design process. In the specific case of these case studies, both composers and performers were identified as stakeholders in the project, thus able to provide essential feedback to the development process. While not necessarily able to frame problems, end-users have been seen to bring perspective to collaborative design that are essential to the development and, as Fisher (1998: 6) states: "For most design problems (...) the knowledge to understand, frame, and solve the problems does not exist; rather it is constructed and evolved during the process of solving these problems (...)".

2.5 Chapter Summary

Examining existing musical interfaces that enable live piano sound modulation using a keyboard (or similar), there seems to be a lack of focus and interest in making these interfaces accessible and available to the trained musician. The reason driving this trend is hard to identify considering that the development of these interface is possible, however one could hypothesise that the thriving market of interfaces for studio production is broader than the one devoted to create new musical interfaces for seasoned musicians. While the GEP (Brent, 2012) and the system proposed by Nicolls & Gillian (2012) was a promising step in the right direction, all other trails seem to have departed from that idea towards other solutions. The current state of commercial keyboard controllers is represented by the ROLI Seaboard and TouchKeys. The ROLI Seaboard is not marketed as an instrument for pianists and the TouchKeys succeeds in being the most familiar interface amongst the ones proposed, but fails in overall accessibility and practical elements. Even though the objective of this research is not to undermine the “cutting-edge” keyboard interfaces already available, an investigation and comparison of these interfaces showed a gap in knowledge and availability of a system tailored to pianists.

Having analysed also the current state of gesture tracking sensing modalities, it has been uncovered that there are reliable ways of tracking hands in a non-invasive way that could provide a transparent experience to the musician. Among all those reviewed, the Leap Motion seemed to be the best for this specific research, thanks to its balance between development, availability and overall reliability.

Having reviewed the hardware and technical aspects underpinning this research, there is already in place a research framework that supports the case for transparency in human computer interaction. This specific aspect of HCI can be linked to transparency and non-invasiveness in AMIs and DMIs through the concepts of control intimacy theorised by Moore (1988) and flow (Csikszentmihalyi & Csikszentmihalyi, 1975) when engaging with a musical instrument.

Chapter 3

Initial Case Study and Targeting Users

This chapter aims to justify the choice of the case studies and the motivation behind the tailored user bases through an overview of the initial case study. Section 3.1 provides a brief overview of the methodology used in the preliminary case study. Section 3.2 presents the early prototype of the Reach system. Section 3.3 describes the preliminary case study and discusses the questions around potential user base of the system and the musical repertoire that arose from it. Section 3.4 covers the chosen categories of musicians for further creative explorations. A chapter summary is provided in Section 3.5.

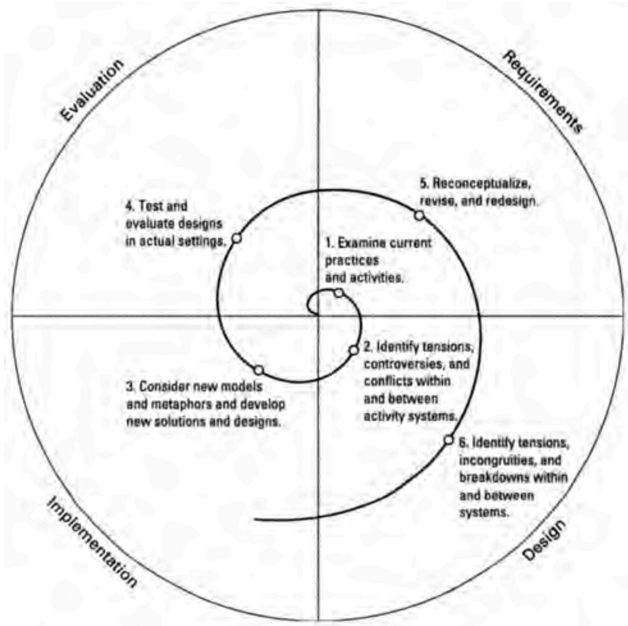

Figure 3.1: Iterative Design Cycle as imagined by Gay & Hembrooke (2004).

3.1 Methodology

The choice of performing an initial case study on my own practice was made knowing its limitations and the possible testing bias. However, while being aware of the limitations, testing the first iteration of the system and gathering the initial research findings on my personal practice was expected to be the fastest and best way to approach this project phase. Referring to Figure 3.1 this initial case study is to be placed at the centre of the spiral. After reviewing the existing literature and developing an initial prototype, the case study on my own practice was identified as a way to further explore the relationship between pianistic gestures and sound modulation, in a real world scenario, before testing the theories with external musicians. The case study could be synthesised in three parts that were later developed using in-depth methods throughout the three years of research. This study required the composition of the piece for the system and piano solo and the public-facing performance of the piece at the Beyond Borders Conference,[1]. This pilot case study carried out on my own practice, was also chosen to see how the ideal candidate would use of the system. The interaction between the system and a musician that was not only a pianist, but also a technologist would have provided a valid interaction baseline to compare the experiences of other musicians with the system.

This initial case study was also used as a first public-facing event in which Reach, and the initial findings of the research project, were presented to the general public. While limited in scope, this initial case study was essential for the refinement of the first iteration of the Reach system in light of the first user testing cycle and provided important insight on the further testing methods used. The composition of the piece for Reach and piano solo uncovered the importance of the point of view of the composer for such a system. While the goal was always to position the pianist at the centre of the research, the interest around how a composer would score for the system arose. The performance made me experience first hand how a system based on microgestural control of sound manipulation could be used in a real world scenario as well as uncovering software flaws and false positives that informed the development of the system in light of the upcoming tests. The experience was extremely helpful from both a practical and conceptual point of view. Thanks to this real-world testing, major software flaws were uncovered and fixed. On a more conceptual level, the case

[1] https://www.bcu.ac.uk/media/research/news-and-events/beyond-borders-event

study provided some useful insight on how the system performed in a real world scenario and how the interaction between the gestures of the pianist and the sound modulation could be used and fostered in a live performance.

3.2 The Early Prototype

Figure 3.2: Overview of Prototype 1

Figure 3.2 presents an overview of the first iteration of the Reach system. This early prototype was able to recognise only one gesture, namely the lateral swaying of the hands after a key was pressed (Figure 3.3). The use of Google's Soli sensor provided precise, complex and continuous data that could not be directly mapped to the pitch shifting effect. For this reason, to perform the microgesture classification, the *random forests* algorithm was used within the Gesture Recognition Toolkit (Gillian & Paradiso, 2014) that was bundled with the Google Soli. Random forests is the name of a machine learning algorithm that operates by combining multiple decision trees into a group (Breiman, 2001). Decision trees are algorithms based on tree-like model where every node represents a test, every branch the outcome of said test, and every terminal node the defined class label. The random forests algorithm takes the prediction from each individual tree and uses the predictions that occur throughout the most trees as the prediction of the model. The resulting combined classification is more accurate than the individual ones because the individual prediction errors of the single trees are not taken into consideration. This phenomenon happens because different trees from the forest will have different errors, when all the classification results from the individual trees are combined, the outliers—individual tree errors—are disregarded keeping only the most common and voted predictions. The random forests algorithm solves the problem of overfitting—issue that arises when the model is trained too specifically to a certain set of data and is not able to generalise—without increasing error. Ultimately, the choice of this specific machine learning algorithm was due to two main factors: prediction speed and generalisation. The robustness of the random forests model lies also in the precision of annotating the data: the better the data is tagged, the more robust the resulting model. For this reason, the gestures recorded and used to train the model were taken from different hands, at multiple distances from the sensor. The random forests algorithm is extremely fast in real-time prediction scenarios and has been successfully implemented in real-time musical applications before (Kursa et al., 2009). For this reason it was the first machine learning algorithm to be taken into consideration due to the live sound modulation nature of the task. Moreover, the generalisation feature makes this kind of algorithm as a good choice when dealing with complex classification tasks despite it being a fairly simple algorithm. While this first prototype was not required to recognise complex classes, the choice was made to classify multiple gestures in future iterations of the system. With the foreseen development of the system moving towards the recognition of multiple gestural nuances with minute differences in data, the random forests seemed to be the right choice. The forest parameters input in the model were the following:

- Forest Size = 5: provided a good outcome for this initial single gesture task. For multiple gestures the size would have been larger to increase the ability of the model to generalise.
- Maximum Depth = 10: the maximum depth of the tree was set to avoid over-fitting.
- Minimum Number Sample = 50: the minimum number of samples allowed per node.

Due to the proprietary nature of Google Soli's SDK, the codebase written using the OpenFrameworks toolkit can't be shared publicly. The software was programmed to obtain relevant data from the gesture identified—the lateral swaying of the hand showed in Figure 3.3—and to send through the Open Sound Control (OSC) protocol to the Pure Data patch, responsible for the sound modulation.

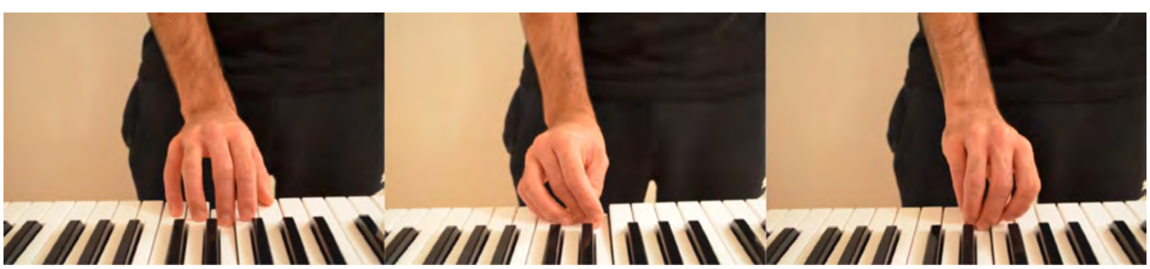

Figure 3.3: Gesture: lateral swaying of the hands

Figure 3.2 depicts two different streams of data—represented by the arrows—going from the sensor to the machine learning algorithm. These two streams of

Figure 3.4: UI of the first version of the Reach prototype.

data were intentionally separated to gate the outgoing data to the sound modulation patch: the first stream of data was the one fed into the machine learning algorithm. This data was used to detect if the gesture on which the algorithm was trained was being recognised. If the algorithm recognised the gesture, then the data was sent through to the sound modulation environment, otherwise it was blocked. The gesture recognition used a selection of Core Features from the Soli SDK to learn the gesture performed. The features taken into consideration were the following:

- *Acceleration* - time derivative of the Velocity core feature.
- *Energy moving* - non-static energy as measured by the sensor.
- *Energy Total* - measured energy as sum of the received radar signal.

- *Movement Index* - level of movement activity of the most dominant object.
- *Range* - distance of the most dominant target component from the sensor.
- *Velocity* - measured energy of the target component.
- *Velocity Dispersion* - measure of the spread of the energy across various velocities.

Figure 3.5: Flow chart of the data affecting the pitch shifting effect.

The second stream of data was composed by three Core Features that, combined together, controlled the pitch shifting algorithm: Energy Moving, Velocity Centroid (weighted average of velocity within a single frame) and Fine Displacement (very small displacements of any moving target component). The value resulting from the combination of these three Core Features—as shown in Figure 3.5—was then mapped to the modulation of the pitch shifting effect. While the Energy Moving was used as a main modulator, the other two parameters multiplied the first one to determine the intensity of the modulation. The choice of having these three values interact in the modulation of the sound was to remove the *one-to-one* mapping *feeling* that would have occurred if the distance from the sensor was directly mapped to the amount of pitch shift. The effect of this parameter interaction caused the hand movement to be associated to the sound modulation, but not in a direct manner: the pitch shift increased and oscillated when the gesture was recognised without directly mapping the position of the hand. Without the use of the gesture recognition algorithm, the direct mapping of the values to the pitch shifting algorithm caused the system to mis-trigger, applying a constant pitch-shifting effect to the live piano feed. For this reason, the example codebase was adapted and used the machine learning outcome to gate the data flow to Pure Data. When the *random forests* recognised the lateral swaying of the right hand after a key had been pressed, it forwarded the positional data to the Pure Data patch in order to perform a live pitch shifting of the incoming signal. If the gesture is not recognised, the data is gated and the pitch shifting algorithm has no effect on the live sound.

The Pure Data patch responsible for the sound modulation, was composed by different subpatches that had been built to perform separate tasks. The OSC receiver for data handling and parsing, an onset detector to reduce false positives incoming

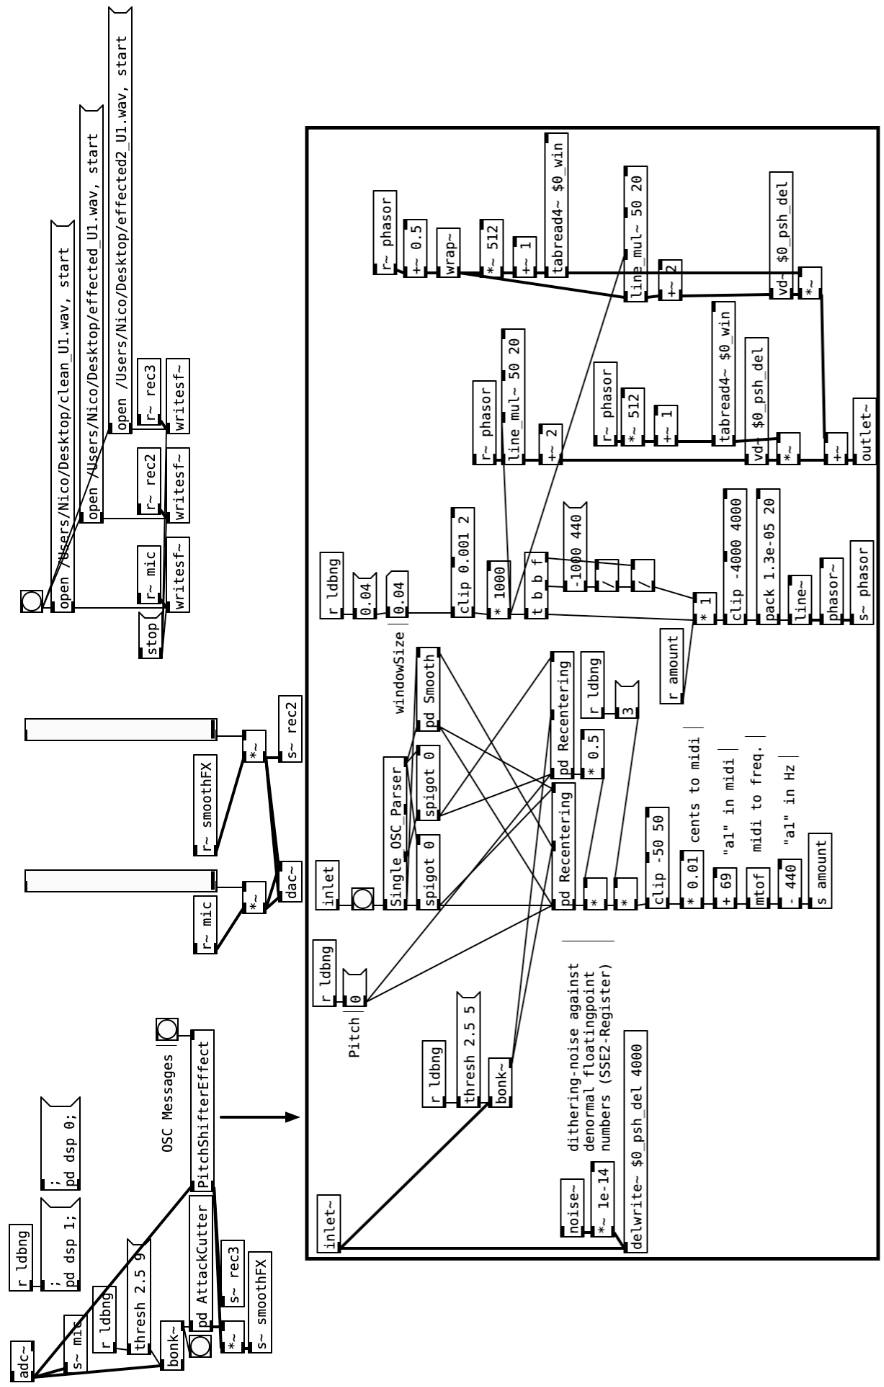

Figure 3.6: Overview of the Pure Data patch responsible for the sound modulation.

from the *random forests* algorithm and the pitch shifting effect, responsible for the sound modulation. The onset detector used was the *bonk*~ object bundled with the standard Pure Data library. The onset detector algorithm implemented uses bounded-Q analysis instead of a more traditional envelope follower, to detect onsets in percussive sounds (Puckette et al., 1998). The choice of this onset detector was due to the ability to detect new attacks which appear not only in the amplitude of the sound, but also in the spectrum without necessarily being accompanied by changes in power. This resulted in being an efficient way of recognising transients in piano notes that—most of the times— appear within the ringing of a previous note or chord.

The OSC receiver opened a communication line with the Soli software application—developed using the OpenFrameworks language—in order to receive the data when the machine learning algorithm opened the gate. The onset detector acted as a fail-safe for the machine learning result. Every time the pianist played a note, the onset detector would raise the volume of the effect and 're-centre' the pitch shifting values, meaning that no audio modulation would occur unless the hand moved away from the current position and the trained gesture was detected. However, if the gesture was erroneously recognised while the pianist was moving hands on the keyboard or performing a quick scale or virtuosic passage, the continuous re-centering caused by the onset detection eliminated any unwanted effect.

Because the focus of the case study was not centred on the originality of the sound effects employed by the system, a simple implementation of a delay based pitch shifting algorithm capped at $\pm$ 50 cents of a tone was implemented. An adapted implementation of the pitch shifting algorithm used in Integra Live[2] was adopted.

3.3 Preliminary Case Study: Beyond Borders

The preliminary case study was presented in a public performance at the Beyond Borders conference,[3] hosted by Birmingham City University. The aim of this preliminary case study was to identify the limitations and constraints in regards to the application of microgestural control of the system in light of the usability test.

The piece was a rearrangement of the intro of "Spain" by Chick Corea (Corea, 1971) for piano solo and wooden bars (Figure 3.7). In the key of D major, the piece explored the landscape of the tonal key through chords, voicings and different melodic lines superimposed. The use of the pedal was essential in creating an extended and continuous layer of sound harmonics. These harmonics created by the continuous sustain of the notes were purposefully notated to provide sonic space for the pitch-shifting effect to be heard and noticed. Considering that the Reach system was—and still is—entirely based on the live sound of an acoustic piano, continuous sound from the piano is required in order for the effects to be heard.

The structure of the piece was the most relevant aspect. The three parts in which the piece was subdivided aimed to highlight the differences of piano gestures and gestural nuances, and their effect on the live sound of the piano, guiding the

[2] www.integra.io/IntegraLive
[3] https://www.bcu.ac.uk/media/research/news-and-events/beyond-borders-event

Figure 3.7: One of the wooden sticks crafted purposely for the performance.

audience through the journey of the pianist that slowly regained control of their pianistic technique.

3.3.1 The Piece

3.3.1.1 Part 1

In this first part, the pianist used custom built wooden sticks (Figure 3.7) of different sizes that allow the playing of chords otherwise impossible to play. This section highlights the non-expressive element of a binary-like action on the piano keys. The pianist, limited only to the mechanical motion of pressing the wooden sticks on the keyboard to play a chord, is deprived of the pianistic hand and finger technique. The resulting performance is meant to be considered *almost MIDI-like* due to the *note-on, note-off* state of the piano keys activated by the wooden sticks. Confining the pianist to the sole use of wooden contraptions instead of fingers results in a binary action that lacks expressivity and musicality.

3.3.1.2 Part 2

The second part fo the piece, acts as a musical bridge between the first binary part and the last final section. In part 2, the pianist slowly starts abandoning the wooden contraptions physically dropping them to the ground. The result of this progressive shift opens the possibility for the pianist to start exploring the piano keyboard with the hands rather than the wooden sticks. In this section the hands of the pianist still perform binary movements to highlight the slow process from the completely binary performance of the first part to a new kind of playing. Because of the lack of expression and ancillary movements coming from the pianist, the system does not activate and the sound is still composed solely by the acoustic piano.

3.3.1.3 Part 3

In the third and final part of the piece, the pianist starts exploring chords and melodic lines using their pianistic technique and their ancillary gesture—the lateral swaying of the hand—to the full extent. Because of the expressive microgesture involved in the performance of this part, the system activates and enables the pianist to modulate sound using the gestural nuances embedded in the pianistic playing technique. This exploration of the newly encountered sonic world continues as the pianist becomes more and more aware of the connection between hand gesturality and sound modulation. The piece concludes when, at the pinnacle of exploration, the pianist picks up from the ground one of the wooden contraptions used to play in the first part and strikes a chord on the piano, returning to a binary playing of the instrument and stopping the interaction with the system and the sound modulation.

3.4 Investigating the Target User Base

Following this preliminary case study an effort was made to tailor all the following tests to different kinds of keyboard players. While the preliminary case study proved that the system was ready to be tested in a real world environment following minor adjustments, and could be used by an ideal musician—a music technology student with a background as a pianist and an in-depth understanding of the system—it did not provide any information on the possible user base. While the system was tailored towards pianists with an already existing piano technique, no tests had been conducted to support any specific musical background of the pianists and the ultimate use of the system.

For these reasons, one of the goals of the tests and case studies performed with the Reach system was to probe different user bases and analyse the responses and feedback to identify the ideal pianist for the system. To achieve this three broad pianist categories were identified to obtain a defined number of groups that were different enough in technique, approach to the piano and style. Moreover, the number of categories was chosen arbitrarily considering the timeframe allowed by PhD studies.

The three categories of pianists identified were: classical pianists, jazz pianists and songwriters that used the piano together with their voice in their practice. These three categories represented three kinds of pianists that used the piano in three different ways. An extra case study with a composer and a classical pianist was carried out to investigate how, from a creative and compositional point of view, an interface like Reach could have been used. Differently from the user testing (Chapter 4) and the comparative user testing (Chapter 6), in this scenario no other characteristics were identified in the participants of the case studies, apart from the requirement of being either formally trained or experienced in their craft.

3.4.1 Classical Pianists

Classical pianists were chosen as representatives of the traditional piano technique and repertoire. With a consolidated classical technique devoted mainly to the performance of scored pieces, they were identified as the base technique.

3.4.2 Jazz Pianists

Jazz pianists were chosen almost in counterposition to classical pianists. Most of the time using adapted techniques in order to perform mostly improvised repertoire, they were identified as holders of a pianistic extended technique that would have possibly triggered the system in a different way.

3.4.3 Singer-Songwriters

Singer-songwriters were chosen to represent a category of pianist outside of the instrumental tradition. Being the piano only part of the performance, singer-songwriters were identified as more willing to bend their technique in order to achieve a sonic result that the other two macro-categories chosen.

3.4.4 Composers

The category of the composers was chosen to insert variables in the case studies conducted with classical pianists. As explained in detail in chapter 5.2 the figure of the composer was used to investigate both how a composer would score a piece taking into consideration the pianist's gestural nuances and how the pianist would react to the scoring or non scoring of the gestural nuances to perform.

Due to the restricted time allowed by a doctoral degree and the will to explore the broadest possible set of creative practitioners, the choice of the different keyboard players was made so as to gain information on a broad range of playing styles. Moreover, the composer was included to explore how the system could have been used creatively from a perspective other than that of the pianist.

A figure that wasn't included is the figure of the contemporary pianist, accustomed to playing using extended technique and technology in performance. The lack of inclusion of this relevant figure was driven by two factors. The first was purely time-based. Considering that the overarching theme of the research was to bridge the gap between traditional instrumental technique and technology, the inclusion of instrumentalists that are by definition comfortable with technology was seen as not relevant for the initial investigation. The focus was purposely brought towards musicians that were not comfortable when using technology to study and understands what aspects of the technology were causing discomfort. The second reason was that the initial focus of the research was purely based around transparency and interaction: sound manipulation and audio effect design were minimal. This meant that what was seen as a novelty for the inexperienced pianist, could have potentially been shallow for a musician with experience in contemporary repertoire and live electronics. For these reasons, the figure of the contemporary pianist, experienced in the use of technology in performance, was seen as more of a future area of exploration. The knowledge gathered from this study as a whole makes a good platform for further research where experienced pianists could be included in the tests and thus expand the knowledge base (see Chapter 7.3.2).

3.5 Chapter Summary

Having pursued a research project revolving around the pianistic instrumental technique, and having focused the design around user centered and user involved approaches, the target musician—the holder of the instrumental technique—had to be thoroughly investigated.

While it was out of this research's scope to investigate every single kind of pianist to pinpoint the ideal user, it would also have been counterproductive to develop an interface for such a specific user. For this reason, the investigation around the concept of *traditionally trained pianist* using the three broad categories explained above resulted in being the best approach to scope the user base.

As explained in depth in the following chapters exploring all the individual tests carried out, the choice of opting for a heterogeneous style of testing brought different results throughout the spectrum of users. All the conducted tests have shown how different practitioners would take advantage of the system in different ways, and how all the pianists approached were able to easily become acquainted with the systems gestural interface at the first encounter.

Chapter 4

User Testing

Following the initial case study discussed in Chapter 3 the system was used for the first round of user testing to be carried out with professional pianists. This test aimed to investigate the nature of the interface interaction and the design of the interface itself. Considering the reciprocity of the relationship between this two elements, the user testing sought to investigate both how the interface interaction could affect further designs of the system as well as how the design of the system itself could affect the interaction of the pianists with the Reach system at a first encounter.

For the first round of user testing, pianists were asked to interact with the first prototype of the Reach system. The main aim of this first user testing was to investigate the interaction between formally-trained pianists and an interface controlled by nuanced gestures already present in their pianistic technique. The different aspects of this interaction that were taken into consideration were three as seen in Figure 4.1: the interaction between professional pianists and the system, the effect of the musical background of different pianists and the system, and finally the different performances and the effect on the interaction with the systems.

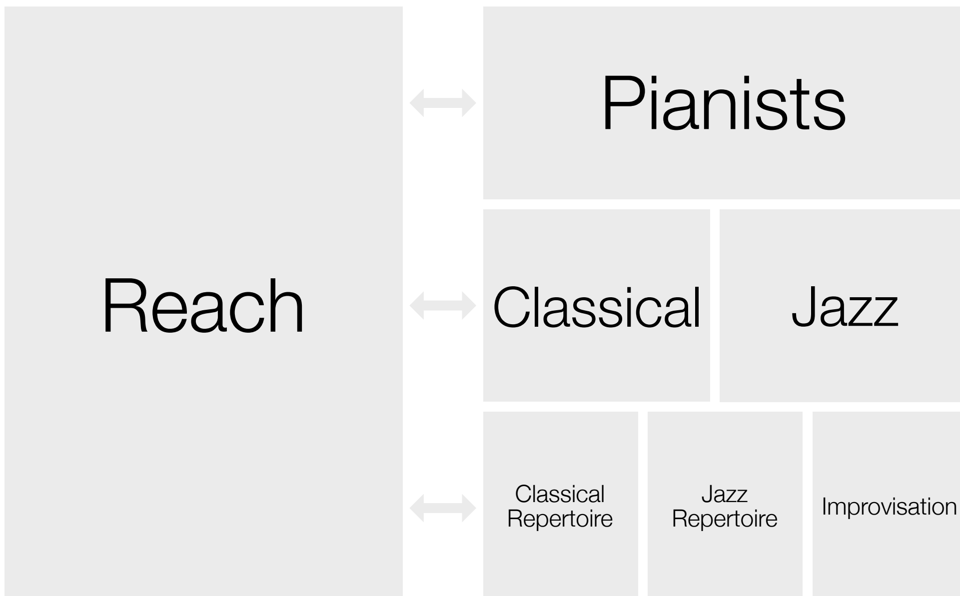

Figure 4.1: Interaction graph between the system and the test focuses.

To achieve these goals, the test was split into 3 main practical parts interacting with the system (excluding an introductory 5 minutes of exploration) that differed substantially from each other. The first part was task based, the second one focused on the performance of a repertoire piece, and the third and optional one was improvisational. The improvisational part of the test was optional due to the lack of experience that classical pianists usually have with this practice. In order to make sure that the pianists were always comfortable with the tasks proposed, in the initial interview each pianist was asked if they felt comfortable to improvise. If the pianists said they were not comfortable, that part of the test was skipped.

4.1 Setup

Twelve piano students from the Royal Birmingham Conservatoire six male and six female, took part in the user testing. The pool of participants took into account pianists of different age and at different stages in their studies, whereas the musical focus was equally split between classical and jazz. The musicians were also asked to consider themselves professional pianists in order to take part in the test. This didn't necessarily entail any specific musician trait (e.g., classical performer, jazz improviser, concert pianist) but rather was meant to highlight confidence in the musicianship and technical proficiency.

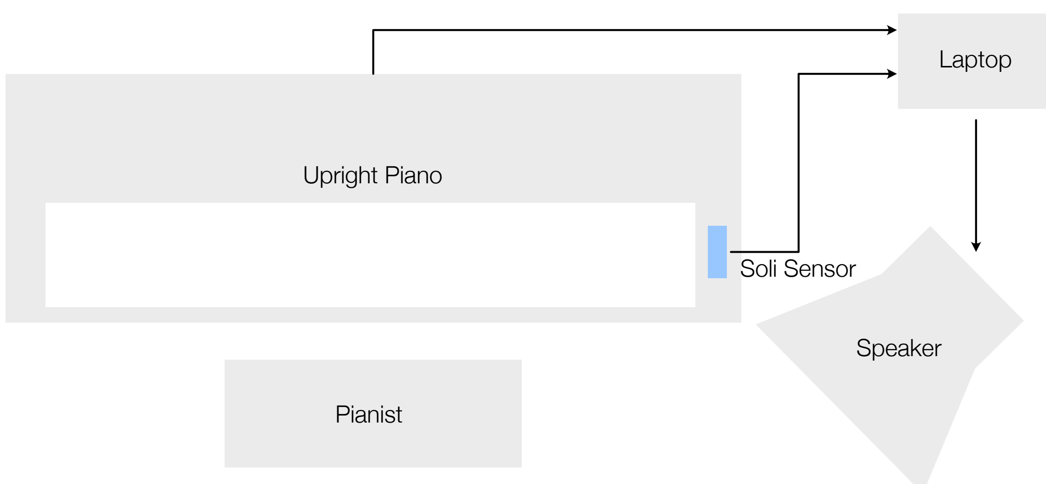

Figure 4.2: Physical layout of the user test.

Eight tests out of twelve were conducted in a recording studio at Birmingham City University. The remaining four tests were conducted in a piano rehearsal room in the Royal Birmingham Conservatoire due to room availability. The two environments used were similar enough not to affect the test outcome or introduce environmental bias. Both environments provided a small space with only an acoustic piano, and some sound absorbing material on the wall. An Audio Technica AT4040 condenser microphone was the only microphone used to capture the sound of the piano. The microphone was amplified through a TC Electronics Impact Twin audio interface, and the sound was then recorded through Logic Pro X. This choice of this specific setup was suited to the upright piano and the small test room. The effects were played through a single Behringer B2031A active studio monitor placed on the right side of the piano facing the pianist. All the gestural recognition was performed through the use of a single Google Soli sensor positioned at the right end of the keyboard.

4.1.1 Reach v0.0.1

As for the initial case study discussed in Chapter 3, the system was set up to detect and recognise only one gesture: the lateral swaying of the hands. This gesture, when recognised by the machine learning algorithm (Section 3.2), was then mapped to the pitch shifting effect (Section 3.2) created in Pure Data, which was capped at a maximum of half a tone above and below the note played. To avoid false positives, the *bonk~* object was used for onset detection as explained in the previous chapter (Section 3.2). This threshold helped limit the triggering of the effect while the user was performing a fast scale or arpeggio, where the fast and continuous movement of the hand would have caused some gestural misinterpretations in the algorithm. No major changes were made since the previous prototype apart from minor adjustments.

4.2 Methodology

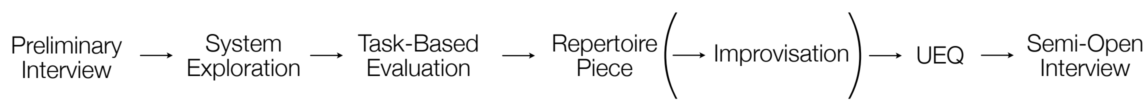

Figure 4.3: Overview of the test procedure.

The test lasted an average of 41 minutes per participant with a standard deviation of 8 minutes. Due to the early stage of the research and the ongoing development of the prototype system, the structure of the test followed an informal formative method (Bowman et al., 2002) to test and assess the interaction. Bowman et al. (2002) describe this method as an observational approach where the participants are put iteratively into situations where they are asked to interact with the interface through a task based approach. The informal side of this approach enables the researcher to gather qualitative data, critical incidents or system failures and user comments. In the conceptual stage of the testing, it became clear that a formal user-testing method including error counting and timed tasks would have been less useful to inform the further development of the system at this stage. Moreover, the formal approach wasn't suited to this particular situation because of the lack of previous reference; a simple empirical test, with a number of questions during a short interview regarding the system, would have resulted in the best outcome.

4.2.1 Preliminary Interview

Subjects were briefly interviewed about their pianistic backgrounds and current knowledge and experience with electronic music and digital instruments. After a brief introduction about the scope of the research and some generic questions aimed at making the user feel comfortable (full outline of the questions asked can be seen in the Appendix A.1), the main questions posed to the participants were the following.

1. *How long have you been playing the piano? What is your musical background?*

   The goal of these initial questions was to gather information about the musical education and background of the pianist. While the categorisation between classical and jazz pianists had been done in advance, a more thorough investigation of the musical experience and interest of the individual musicians aimed to uncover potential experiences outside the generic categorisation.

2. *Would you say you are comfortable with electronic instruments and systems? If yes, which instruments specifically?*

   The second set of questions aimed to investigate the possible experience with electronic instruments. If there was an interest in electronic instruments, this question aimed to elicit what kind electronic instruments the user had experience with. This led to a deeper discussion around the type of electronic instrument, the kind of interaction with it, and the scenarios in which these electronic instruments were played (e.g., in a private rehearsal scenario, in a public facing performance, in a music store). The outcome of this question would have helped frame the answers to the last part of the test, when the user was asked to discuss the experience with the Reach system, the interaction, and the possibility of performing with it.

3. *(If yes to the previous question) Which experience do you prefer: having somebody handling the live electronics for you, or having full control over the effects?*

   The third question, was aimed at that small group of pianists that could have experienced live electronics in a scenario where they weren't in control of the sound modulation directly. This question aimed to explore the concept of the *invisible technician* (Shapin, 1989) from the eyes of the performer and explore the importance of being in control of the sound modulation from a pianists point of view.

4. *Do you have any experience with improvisation?*

   The final question before the test had both a practical and a theoretical goal. The practical goal was to investigate if the pianist was comfortable at improvising in order to avoid creating an uncomfortable situation at the end of the test. Theoretically instead, the ability of the performer to improvise was meant to inform the kind of sonic exploration with the system. Prior to the test, the expectations were to see the classical pianist adhere very strictly to the score played, and the jazz pianists be more free to improvise and explore the sonic environment enabled by the Reach system.

4.2.2 Exploration of the System and Task Based Evaluation

After the short interview, the participants were given 5 to 10 minutes to explore the system and get used to the effected sound coming from the single speaker. The sound emitted contained no unprocessed piano signal: this choice was purely practical due to the loudness of the piano compared to the small size of the room. During the initial exploration, the system was calibrated adjusting to the gestural technique of the pianists. Participants, once comfortable, were then asked to perform a series of simple tasks to assess the precision and reliability of the system. The tasks were: play a note, play a chord, play a scale.

All the above tasks were to be performed twice; the first time, participants were asked to try not to activate the system, while the second time they were asked to purposely try to activate it. This exercise was done to make participants aware of the threshold in which the system would have recognised the gesture, and how it triggered the audio processing. This task-based part of the test was also used to test the reliability of the system and software. Later on, the data from these comparisons was also used to see when the participants were successful in completing the required task and how the exercise was reflected in the feedback given during the interview.

4.2.3 Repertoire Piece Performance and Improvisation

Subjects were then asked to either sight-read and perform a piece provided on the day,[1] or perform one from their repertoire.[2] Once the piece was chosen, participants were asked to perform the piece twice: the first time without the system active, as a frame of reference, and the second time with the system. Two participants chose to perform a piece from the provided repertoire, and both were coming from a classical background. The pieces chosen by these two performers were *September Chorale* by Gabriel Jackson and *Bells* by Simon Bainbridge. Finally if they felt comfortable enough with the system to improvise, they were asked to do so. Once all the tasks on the piano were complete, the participants were given 5 minutes to fill in the UEQ. Each subject was then asked in a brief final interview about the experience and the system.

4.2.3.1 Open Interview

The open interview method was employed at the end of the test to gather qualitative data, quotes and thoughts about the system, the experience and the interaction. An initial series of open interviews was considered as an ideal preliminary investigative

[1] The pieces provided on the day were: September Chorale (Jackson, Gabriel), Nocturne I (Harrison, Sadie), Nocturne II (Harrison, Sadie), Utrecht Chimes (Lange, Elena), Bells (Bainbridge, Simon), Yvaropera 5 (Finnissy, Michael).

[2] The pieces that the pianist played from their repertoire were: Nocturne in C-sharp minor, B.49 (Chopin, Frédéric), Piano Sonata No. 8 in C minor, Op. 13, Adagio cantabile (Beethoven, Ludwig van), Paraphrase de concert sur Rigoletto, S.434 (Liszt, Franz), Faschingsschwank aus Wien, Op. 26: I. Allegro (Schumann, Robert), personal arrangement of When You Wish Upon A Star (Edwards, Cliff), Nocturne, Op.32, Andante sostenuto in B major (Chopin, Frédéric), Piano Sonata No.10, Op.14, No.2 (Beethoven, Ludwig van), Prelude in G-Sharp Minor, Op. 32, N. 12 (Rachmaninoff, Sergei).

method that would have informed the semi-structured interview process of the next testing cycle (Wilson, 2013). The interview was led by generic question about the different sections of the test, then going in depth when the user explored any particular argument or unveiled any topic that might have been of interest. While at this stage of the research there were already some noticeable trends, an open approach to the first experience of the users resulted in being the best approach.

4.3 Discussion

The questions in the preliminary interview helped frame the feedback and the interviews with the different musical background of each user. However, no significant correlation was found between the result of the test and the musical provenance of the user. All of the participants, whether coming from a jazz or classical background, were able to perform with the system and expressed the potential of it being implemented in their personal performances. During the interview, one user said:

> *This is very diverse, can be applied to classical, jazz, anybody who plays the piano. It can be for anyone [...] it was really interesting to play on a real piano, in its natural form being able to effect sound is not something is possible without controls and effects.* (referring to knobs and effects on their keyboard)

This was a common thread between all of the participants: the ability of a pianist to sculpt the resonance of the sound on a traditional acoustic piano, was what fascinated and attracted the participants to the Reach system. While it was noted that all of the users found the research project exciting and fascinating, these comments were mostly due to the fact that the pianists had never experienced any gesture based interface before.

Another common topic between the users was the accessibility of the system in relation to both the use of an acoustic piano and the recognition of gestural nuances present in the pianistic technique. One user explained how the combination between these two core elements of the system made the experience accessible. When asked to describe the approach of the system towards traditionally trained musicians the pianist said:

> *You're playing an acoustic piano for one, so if that's what you know it's quite easy to step into it. Also it has a quite strong effect with a small movement. So it's immediately quite accessible.*

Without being prompted, two users discussed the gestures required to interact with Reach. One user, when asked if the system was what they expected said:

> *It wasn't what I expected, actually I didn't know what to expect before I came here. It was very interesting to see how my hands movements could change the sound. What I noticed is that if I'm very rigid, and If I don't try to express with music the sound is plain. When I try to communicate instead the sound changes.*

The found connection between gestural nuances and expressive playing on behalf of the user was very interesting. As noted by Tits et al. (2015), the amount of gestural nuances—named *eigenmovements* in the study—in the pianist's technique increase with technique proficiency and years of study. This connection was one of the reasons for which gestural nuances were chosen as control gestures for the system. The second user, when asked if they would use the Reach system in a performance, replied:

> *Yeah absolutely. Especially not having to modify technique, I think it's really important especially if you're not used to doing various things on hand on the piano and another hand fiddling around, that could be difficult. I would definitely see that as something I would use.*

This comment aligned with the findings of Nicolls (2010). The problem with interfaces that required pianists to learn additional gestural vocabulary was not the gestural vocabulary itself, but rather the ability to retain pianistic freedom and control at the same time. The participant in question saw interfaces that require to interact with knobs and faders as a problem when it came to play the piano at the same time. For this reason the use of pianistic gestural nuances of the Reach system was identified as a positive quality.

The data extracted and plotted from the UEQ, displayed high and positive marks regarding the attractiveness and hedonic quality of the system, with a lower average for the pragmatic section, as seen in Figure 4.4. These lower but still positive marks concerning the responsiveness and reliability of the system, were tightly connected to the prototyping stage of the project. During the series of simple tasks the participants were asked to perform, sometimes the system activated without the user purposely performing the gesture, or, in rare cases, did not activate at all. This was due to the nature of the radar sensor, sometimes unstable due to its early development stage, and overall the initial development stage of the software itself. This led to some participants not feeling completely in control of the system, some even thinking that they were performing the gesture wrong.

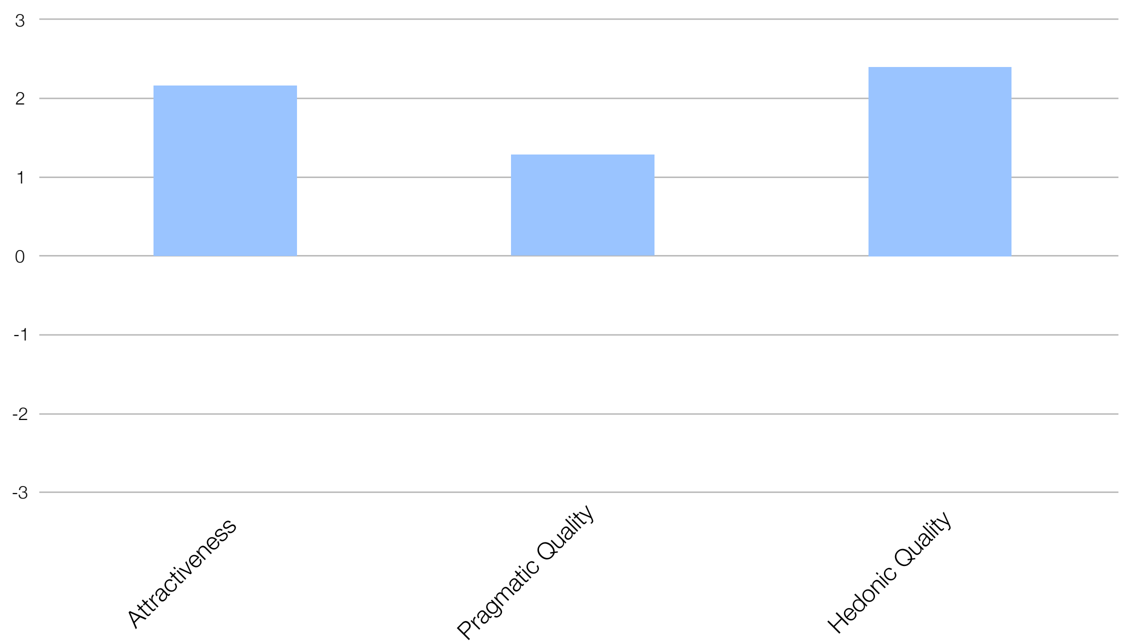

Figure 4.4: Bar graph showing values for the three broader qualities analysed by the UEQ.

> *I think that sometimes when I stay on a note, if I stay on the note a really short time it doesn't really react much. I wouldn't like it to react faster, it was just a consideration. Also, when I play a single note I feel like it reacts more than when I play a chord or an arpeggio. I think this is because of me doing it incorrectly.*

During the interviews, a relationship emerged between the pieces performed by the participants and the feedback given. Participants that chose to perform one of the pieces proposed on the day, tended to be more willing to adapt the performance and explore the system, interpreting the composition and the tempo to accommodate sound modulation through the system. The listed pieces were chosen together with a piano teacher from the Royal Birmingham Conservatoire because of their temporal and rhythmic freedom and long ringing chords. The teacher put together this collection of pieces after having tried the system, and realised that it would have been better used on these kind of slow and open compositions, rather than on classical repertoire pieces. This shows links to the findings of Winters & Wanderley (2012), who noted how the structural level of gesture is tied to the underlying music; in their study, easy passages and phrase boundaries saw an increase in movement. In fact, this turned out to be a key factor that encouraged pianists to take advantage of the system. In the majority of the tests, when participants chose to play a piece from their classical repertoire, the feedback was less encouraging. Participants didn't feel the need to add this additional expressive layer on a consolidated repertoire piece that they already knew how to play expressively to convey a certain emotion. While asking the pianists to use the Reach system—specifically with a pitch shifting effect—on a classical repertoire piece was expected to receive negative feedback, it was essential to provide the pianists with an alternative in case they were not comfortable in playing at first sight one of the proposed pieces. Overall, it was more important to make the users feel comfortable and make the ancillary gestural nuances surface naturally during the performance of a known piece than requiring the participants to sight-read a new piece possibly causing a less spontaneous performance. Moreover, the negative feedback received in these cases can be also related to an observation by McNutt (2003): performers need to have a reasonable mental representation of what sounds they will hear, and in this specific case, what sound their hands will produce.

In hindsight, rather than having the pianists play pieces from their repertoire, étude-like pieces could have been composed purposely for the test. These études would have enabled an exploration of the connection between music close to classical repertoire, without having the clear disconnect between tradition and modulated sound. Moreover, being études composed to practice particular aspects of the pianistic technique, they could be composed to relate to specific ancillary gestures that highlight the pianist's subconscious movements. This idea could be an implemented in future iterations of similar user tests building upon the results of these early stages of the research.

Reinforcing the observation by McNutt, all of the participants noticed and expressed a similar feeling after having tried the system. While the system was reducing the learning curve typical of interactive systems significantly, the strain had shifted instead to the ability of predicting and expecting the sound of the instruments. Five participants out of twelve clearly stated that the hardest element of the system to get accustomed to was the sound, and not the gestures used to obtain it. Only one user stated that they did not like the sound to be different from what they expected.

> *In this case I heard something I wasn't expecting, before I played I knew how the sound (of the acoustic piano) should have been, and when I played*

now I was like 'wow what is this' because it's something new, and I don't like the sound to be different to what I hear before (. . .)

Practice time was another key element that emerged frequently during the interviews. Even though the system did not require additional techniques, three out of twelve participants said they would have needed and wanted more time to practice playing with the system. This practice time would have helped them learn what their pianistic gestures would correspond to, from a sonic point of view. When one of the participants was asked if they felt the need to add additional technique to play the system, they replied:

The technique that's needed is the listening, as we say we pedal with our ears. It's really what it's about.

The response of the user underlines how closely tied the last two topics are. The fact that participants could not predict how the system would have sounded, meant that the system would have felt invasive from a sonic point of view when talking about a performance of an already known piece or, more in general, any piece of classical music. For pianists, classical music is expected to sound a certain way. Thus hearing a classical piece with the effect of the Reach system superimposed felt aurally invasive and wrong. This explains why even though some participants did not regard the system as invasive from a sonic point of view, they felt that they needed practice time to get used to the system itself.

The feedback received from the improvisation-driven section of the test, contains data from seven participants out of twelve. The limited amount of participants willing to partake in a small improvisation was expected, as this part of the test was mainly aimed towards the jazz pianists; however one classical pianist asked to try and improvise with the system. The results ended to be very similar to the previous part of the test. During the improvisations, participants were more willing to explore the gestural components of their playing and were keen to unexpected sounds and timbres. This brought to some really interesting improvisations that explored this new found sonic environment, and more than one user asked if they could keep the recording of their performance. When asked to compare the repertoire and the improvised musical moments, one user said:

> *I'd say they were different, I wouldn't say one was better than the other. The theme was less spontaneous, so you knew what was coming up, so I was able to pre-empt. Whereas the improvisation is spontaneous, so I would have to be actively putting it and using it.*

Another user said:

> *I guess someone could be inspired, and write a piece for it, or someone could use it to aid a performance. Not so sure about pre-existing composition, I am sure that for me if I wrote something I wouldn't want to mess around and perform it in a manner that's adding something that's not in the original scripting of my writing.*

As previously seen in Figure 4.4, before the final interview process, the participants were asked to complete a UEQ that enabled to evaluate the system regarding its efficiency, perspicuity, dependability as well as other aspects regarding the user experience such as originality and stimulation.

Figure 4.5 shows the mean and standard deviation of each individual parameter showing the maximum and minimum score on the Y-axis. In Figure 4.4, the attributes of attractiveness and hedonic quality were given the highest score in the test. These two values are closely related to the appreciation of the system itself, such as the system being interesting, fun to play with and enjoyable. The most likely responsible for these high marks is the non-invasive character of the system that gave the user an additional sonic element without needing extensive training. These two macro-categories can be seen in Figure 4.5 as well, represented by attractiveness, perspicuity, stimulation and novelty.

Efficiency and dependability instead—categories pertaining to the pragmatic macro-category—were the ones that received the lowest marks. While nine participants found the system really inspiring to play with, three participants didn't feel in control of the effects, and found the system not responsive enough. Mis-triggering was a major factor that caused this lack of confidence in the system.

On the other hand, when the system operated in an optimal manner without any major mis-trigger, the participants felt in control and were able to control the audio effects in an expressive manner through their own technique. In other occasions, participants managed to get used to the instrument and the interface during the short duration of the test. One user stated:

> *Yeah I felt mostly in control at some points maybe I was worried I wasn't doing it right. But especially once I got used to it, it felt a lot easier to*

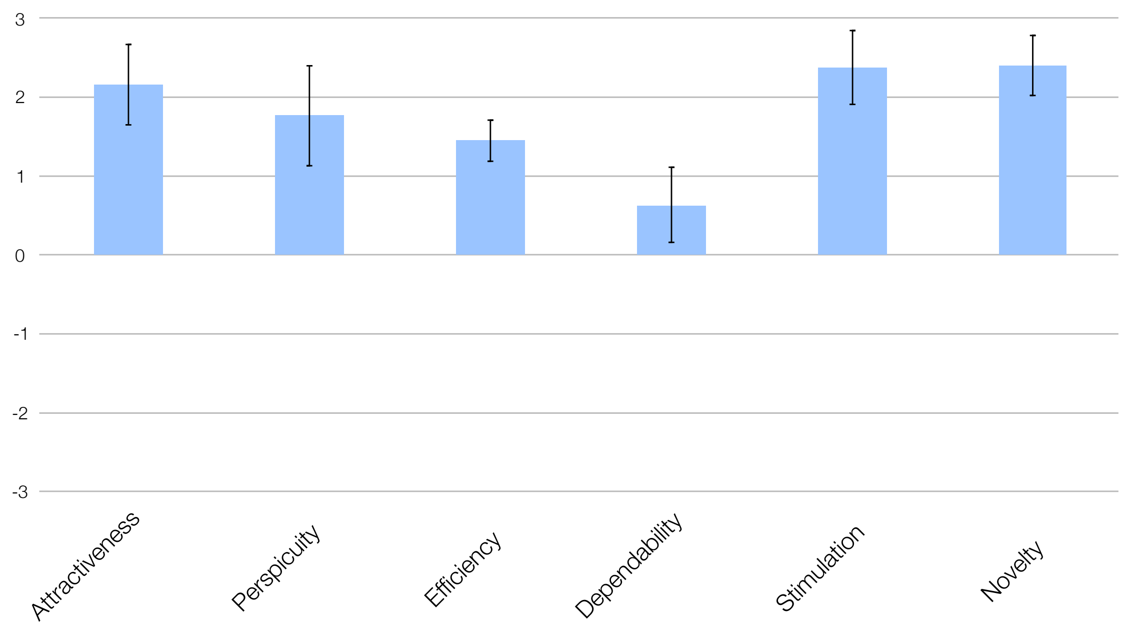

Figure 4.5: Graph showing mean and standard deviation for each value analysed in the UEQ.

> *control. There were a couple of points where I really was thinking If I was performing the gesture correctly, but I don't see it as a long-term issue because I played for a total of 15 minutes.*

Due to the prototyping stage of the system, these kinds of comments were expected and informed later stages of development to make the system more reliable.

4.4 Chapter Summary

The initial user testing phase, showed initial positive results in favour of the development of new interfaces for musical expression based on pre-existing gestural nuances in piano technique. It was uncovered that interfaces that build on existing instrumental technique and are considered non-invasive, such as the Reach v0.0.1 system, create less disruptive experiences of technology in performance. With a pool of 12 pianists ranging from two different musical backgrounds, the test showed how pianists from various genres may benefit from the development of such interfaces. Even though the low degree of invasiveness related to the interaction has shown promising outcomes, the sonic invasiveness of the system has revealed to be an obstacle. However, due to the nature of AMIs, where the sonic invasiveness is a default characteristic, a system in which the effected sound is mixed acoustically with the true sound of the piano is the best and least invasive way of handling this issue. Because of the complexity of the v0.0.1 system, and the lack of availability of the technologies used, a new topic emerged out of this user testing. Further tests started to question how to make the interface accessible to participants overall, not only in the interaction. For this reason, the technology used to develop Reach was re-evaluated to find more accessible solutions. With the development of new technologies and devices available, a more accessible way to achieve the same results needed to be explored in order to actually open new horizons to composers and musicians alikef

Chapter 5

Case Studies

The case studies conducted during the second year of research aimed at providing more insight around some of the questions surfaced during the previous studies. These questions revolved around further interaction possibilities, identifying a user base and finding the ideal musical application. Regarding the user base, apart from a brief encounter with a small pool of classical and jazz musicians during the first user testing (Chapter 4) up to this point, this topic had remained uninvestigated. Moreover, an interest around what range of interaction possibilities could be opened by investigating more ancillary gestures and more types of control over sound manipulation flourished. During the interviews of the first user testing, one question focused on the possible uses of the system. The answer to this question revealed quite a disparity in replies. While some users said the system would have been better tailored to an improvisation-based scenario, others felt it would have been a valuable addition to selected repertoire pieces as well, mostly referring to the pieces proposed during the test. Some users strongly disagreed with the last statement, and saw the system as more of a compositional tool that would have inspired composers to craft and compose new pieces. A similar situations can be found when trying to frame the most suitable musical genre for the system: while classical repertoire was identified, as expected, as the least suitable genre—as previously explained in Section 4.3—, jazz, free improvisation and experimental music were seen as appropriate candidates for the system. In the user test explored in Chapter 4, pianists seemed to agree that this system was not tailored to traditional piano repertoire, in which the introduction of a novel element not originally scored would have brought disruption to the composition. However the remaining musical spectrum had yet to be explored.

The present chapter is going to explore four different case studies conducted with three creative practitioners that investigated the systems in different musical scenarios. The aim of the case studies is to examine not only the user-base of the Reach system, but also the effects of different musical backgrounds on the first encounter with the musical interface. The guided and test-based nature of the previous explorations did not provide enough information on the relationship between musical context and relationship with the musical interface. For this reason, the following case studies adopt a less structured approach based on workshops, collaborative design and public facing performances.

5.1 Reach v0.1

Figure 5.1: Overview of the second prototype of the Reach system.

The user tests conducted in a controlled environment informed the development of the next iteration of the system, Reach v0.1 (Figures 5.1). This system changed not only the software itself, but switched to a more robust and easily obtainable hardware sensor (Section 5.1.1). This hardware change was driven by both the accessibility of the new hardware and the ability to track both hands simultaneously. This was the first iteration of the system that used the Leap Motion as main sensor, and is the one that forms the basis for the software and hardware used in the final iteration of the Reach system.

5.1.1 Camera-based tracking

After the first round of user testing, the main technical difficulty with Google's Soli sensor was the inability of detecting both hands on the keyboard without having a two-sensor system setup. At a conceptual stage this issue was disregarded, but the more the research advanced, the more the tests required the users to interact with a finished instrument. This limitation of the early prototype surfaced multiple times in the first user testing (Chapter 4) and was one of the most pressing issues that had to be fixed in the next development stage. Due to the early stage of development of the Soli sensor and the dissolution of the Soli R&D team (later linked to the FCC stalling the development (TechCruch, 2019)) the use of the Soli sensor was hindering the progress of this research project.

In the same year, Leap Motion released a new SDK oriented towards VR for the same hardware they had been selling and supporting for the previous five years. The low price of the old and consolidated hardware paired up with the improved performance seemed an appealing alternative to the radar-sensing technology used until then. Informal testing found that the Leap Motion could offer a more stable and supported platform to develop the system following the research and feedback gathered until that moment. A more detailed explanation of the sensors and the reasoning behind the choice can be read in Chapter 2.

5.1.2 System Structure

The second prototype was more streamlined than the initial one and, thanks to the improved data coming from the new sensor—now pre-processed by Leap Motion's SDK—the machine learning algorithm was removed. Only the onset detection algorithm was left to control and avoid false positives, as explained in Section 3.2. The onset detection, developed in the Pure Data environment, replaced the machine learning algorithm in gating the signal.

To take advantage of the VR oriented Orion SDK, the bespoke app was developed using the Unity environment. This choice was purely practical, as it would have

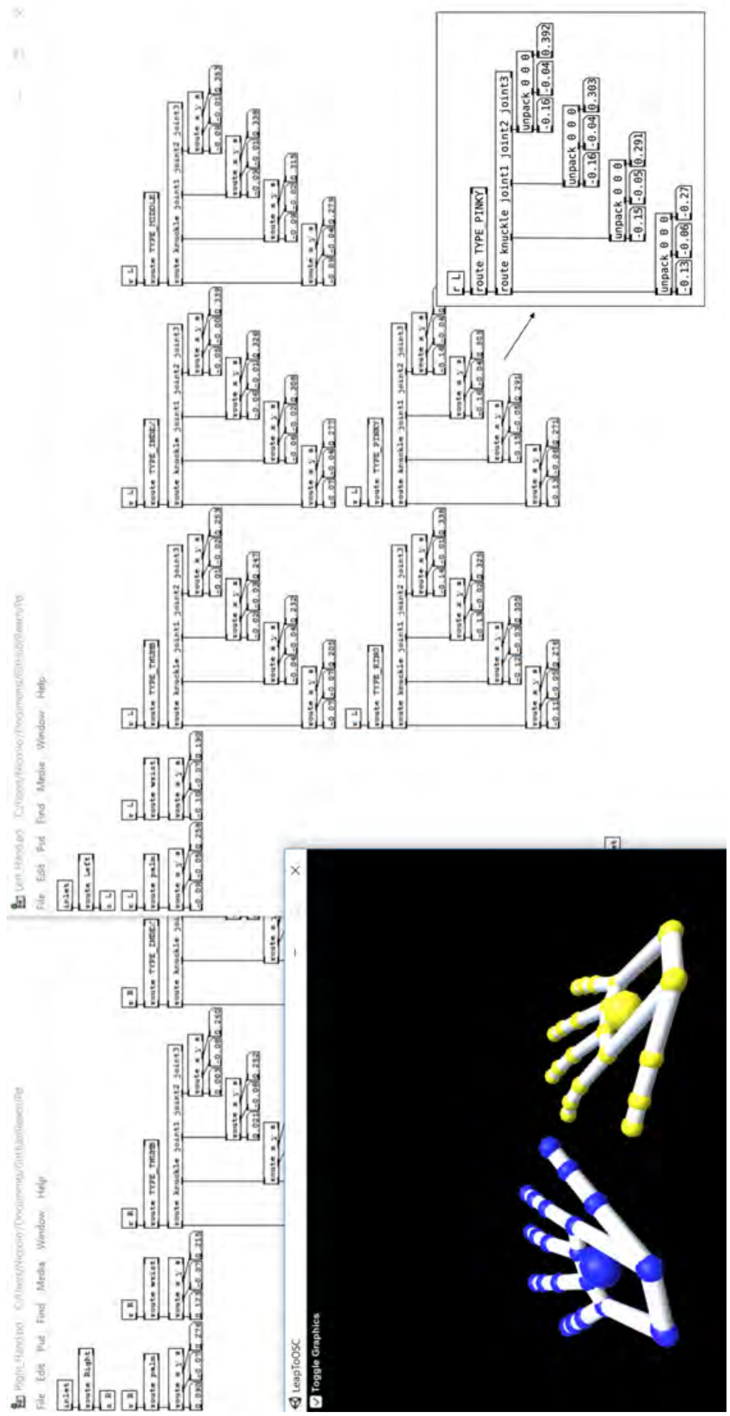

Figure 5.2: UI of the second iteration of the Reach prototype.

enabled short prototyping times thanks to the large support of the environment itself, the OSC libraries and the Leap Motion integration (Figure 5.2).

This version of the Reach system was built using free development platforms and released as an open source project to make it as widely available as possible.

Due to the shift to the Leap Motion hardware, and the positioning of the sensor above the keyboard rather than on either sides, this version of the Reach system was the first to track both hands and enable multiple gestures to be mapped to sound modulation. While every case study had a tailored version of the mappings and thus had different gesture-sound modulation pairs, the system was initially conceived to recognise three separate gestures: lateral swaying of the hand, wrist height when the hands are in contact with the keyboard and hand height when leaping away from the keyboard. The sound-gesture pairs specific to each case study will be explored later on in the dedicated sections (Sections 5.2.2, 5.3.2, 5.4.2, 5.5.2).

5.2 Case Study 1: Richard Stenton (Composer)

The first case study was conducted with Richard Stenton, post-graduate researcher and composer at the Royal Birmingham Conservatoire. The piece composed throughout this case study titled *//aria alla francese* (Stenton and Granieri, 2018), explored the relationship between the combination of traditional and graphical notation, gesturality and sonification through a piece for piano solo.

//aria alla francese had the goal to focus on the interaction between musical notation, pianistic gestural nuances and sound modulation. For this specific case

study, the composer decided to focus entirely on sound accompanying gestures that take place after the sound producing gesture has occurred. The composer's goal was to analyse how a combination of traditional and graphical notation, could inform and cause the surfacing or disappearing of gestural nuances. The hypothesis was that a specific kind of scoring would elicit gestural nuances to surface in the playing of the pianists. The compositional approach was derived from the philosophy of the system. The aim was to write music that would facilitate the pianist to deploy additional expression through their natural technique as opposed to directing the pianist to create a predetermined affectation. Thanks to the low degree of invasiveness of the system, and the acoustic grand piano, the composer aimed to provide a fertile environment for nuanced gesturality to surface in performance.

5.2.1 Methodology

To compose the piece and co-design the system, three composer-led workshops were organised where the composer, the pianist and I were present. My role in these workshops was purely technical as I was there to tailor the system to the needs of the composer, or to try and accommodate the requests of the pianist. This was necessary due to the prototyping stage of the system, not yet ready to be used without assistance.

5.2.1.1 Workshops

The first workshop was used to test different types of musical material to see to what extent they would firstly, elicit movement from the pianist and secondly, allow space for the pianist to make decisions rather than making gestures only be a direct result of her instinctual reaction. For this first workshop the system was limited to recognise one gesture only, to avoid overloading the pianists with new sonic information and to enable the composer to clearly see the relation between gesture and sound. The gesture chosen for this initial session was the lateral swaying of the hands, paired with the same pitch shifting effect implemented in the first early prototype of the system. This choice was made for two reasons: the successful test of the gesture-sound effect mapping in the previous user testing with multiple pianists, and the ease to approach this mapping due to the strong relationship with the gesture performed by string players to play vibrato. This concept can be linked to intuitive sound-gesture metaphors implemented by Mainsbridge & Beilharz (2014) in the *Gestate* project. While in that instance the concept of automatic knowing—called intuitive knowing by Antle et al. (2009)—was used to describe the incorporation of familiar associations and embodied understandings of physical phenomena, in the Reach system, the automatic knowing of the sound-gesture coupling related to the string instrument *vibrato-like* gesture and the embodied understanding of this musical phenomena through familiar association and musical culture.

The second workshop, conducted in a similar manner to the first one, presented the users of the Reach system with three sound-gesture pairings. The system was now able to recognise three gestures chosen by the composer. The gestures added were the vertical distance of the hands from the keyboard when performing leaps and the height of the palms of the hands when on the keyboard. The first gesture was paired with the wet/dry control of a reverberation effect: the more the hands were distant from the keyboard, the more the reverberation. The height of the palms when the hands were on the keyboard was mapped to the volume of the clean signal of the piano coming out of the speakers. The higher the palms, the louder the volume. This was meant to give the performer more control on the sound decay by amplifying the tail of a note or a chord, similarly to string instruments. These gesture-sound couplings were combined with a more structured and detailed score (Appendix A.4) to explore the connection between notation and gesture. The composer was thus able to analyse the relationship between the system, the gestures and the scored material.

The third and final workshop was intended as a rehearsal for the recorded performance, and was followed by a brief discussion between the performer and the composer to gather some feedback on the system, the piece and the interaction. Results are discussed in the discussion section of this chapter (5.2.4).

5.2.2 Sound-Gesture Mappings

An addition was made to the original gesture-sound coupling—previously used in the user testing (Chapter 4)—thanks to the new feature of the system, that was now able to track separately each hand. By using hand position on the keyboard, the pianists was able to modulate the bass spectrum with the left hand, and treble with the right. This effect was implemented by mapping hand position to the frequencies corresponding to the extremities of the keyboard. The left and right hand had control respectively over a low-pass and high pass filters with centre frequency mapped between 27.5Hz (A0, the lowest key on a piano) and 4186Hz (C8, frequency of the highest key on a piano) as illustrated in Figure 5.3. This effect filtered the modulated sound removing high-pitched frequencies from the notes played with the left hand, and low frequencies from notes played with the right. The resulting modulated sound gave the performer the impression to have control over the specific notes being played with each hand and thus was implemented even though the filtered modulated sound lost in timbre richness.

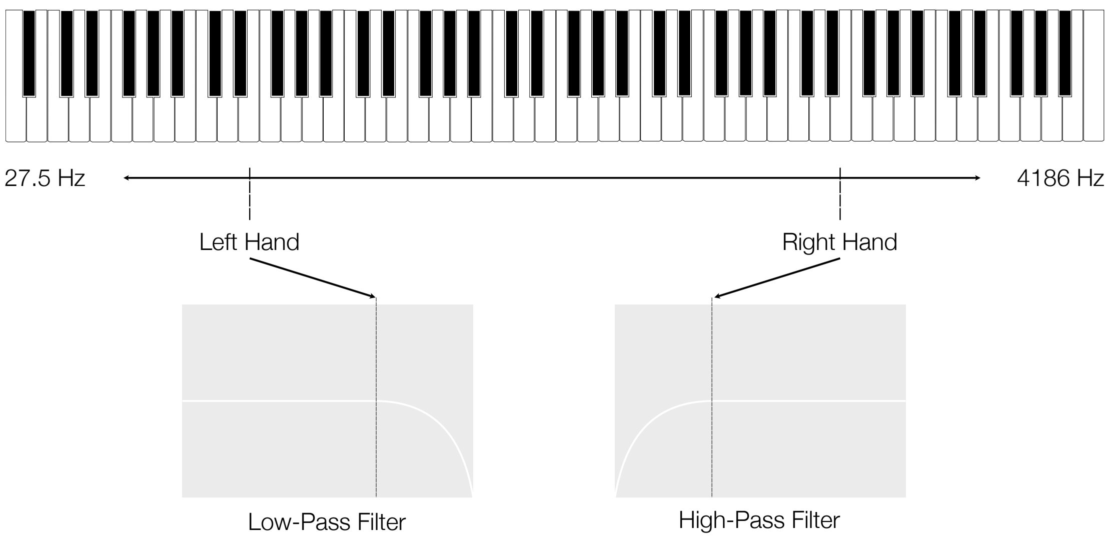

Figure 5.3: Outline of the hand position mapping to the low-pass and high-pass filters.

5.2.3 Initial Findings

The three workshops uncovered an interesting relationship between musical tempo and gestural content. Music at a medium tempo (*approximately 80bpm*) but with—using the words of the pianist—an *"unstable rhythmic feel"* elicited the most interesting results. As seen from the full score in Appendix A.5, the unstable rhythmic feel experienced by the pianist was due to the syncopated motion of the piece. This was later linked to the concepts of *ping* and *lilt* (Pierce, 2010), terms used to describe respectively the physical experience of beat, and everything in between beats. For Pierce, the musical rhythm, which is the combination of pings and lilts, is reflected in the musicians body in the form of sudden movements of the wrist and smooth swaying of the arm. During the workshops it was uncovered that with slower pings, and unstable rhythmic feel, the energy accumulated in the lilt to perform the scored notes produced predictable gestural nuances that could have been used to control the sound modulation parameters. This discovered correlation between rhythm and gestural nuances disappeared if the musical conditions were not optimal. For example, if the musical material was too fast, then there seemed to be no time for the pianist to perform the gestural nuances and although some effects would be activated they were purely accidental. The result was a piece for piano solo with sporadic gesturally controlled effects that were not predicted by the composer nor predictable and interfered with the performance. Alternatively, with a *Lento* pace, the pianist was faced with ample temporal spaces during which controlling the system through gestures was possible. However, it was also noticed that when the tempo was too slow, the pianist didn't perform as many gestural nuances as a result of the expressive playing and was left with a lack of direction on how to move or what to do. When this problem surfaced in the first workshop, an initial solution was to increase the amount of effects that the pianist controlled and score explicitly the gestures envisioned by the composer, thus making the pianist more aware of the connection between the gestures and the modulated sound. This explicit scoring of the gestures had been explored during the first workshop with the initial investigative score (Appendix

A.4) that was meant to help uncover the correlation between scored gesture, actual gesture and sound modulation. The implementation of an increased modulation of the effects combined with a more thorough scored notation for movement reduced the lack of direction that the pianist initially felt. However, when the pianist was required to play the piece and not only test some notes and gestures, the gestural score elicited non-natural movements. These movements were performed by the pianist only to comply with the score and thus affect the piano sound. This approach was briefly tested and abandoned due to the lack of expression in the playing caused by the rigid gestural scoring and the loss of focus on gestural nuances present in pianistic technique. The problem with the ability to play expressively was linked to the issue of a lack of unified notation protocol when focusing on non-traditional gestures and sounds found by (Nicolls, 2010: 38). Pestova (2017) described a similar issue regarding notation and non-traditional compositional approaches. The more the composer tried to score something out of the ordinary, the more this caused a stifled interpretation because the unorthodox scoring was unnatural for the performer. Due to the objective of using natural gestural nuances, a less scored and more traditional approach was taken.

Using a strict pulse at an *Andante* pace (*approximately 80bpm*), gave a context for the pianist to place gestures, and unpredictable rhythms gave energy and vitality to the movement. It was also noted that giving the pianist a sense of cadential pull from chord to chord gave them a more natural and emotional reason to use the system. This was in line with the findings of King (2006) that noted how in Beethoven's Sonata in C Minor, Op. 13 'Pathètique' (*Adagio cantabile*), hand gestures in proximity of the cadential statement where not only consistent but also caused gestural reverberations (Pierce, 2010) through the following quaver rest. Gestural reverberations are the gestural way in which performers show to be aware of the junctures between musical phrases: prolongations of musical gestures not necessarily related to the sound production help give rise to a less *stifled interpretation*. The three workshops strengthened the understanding of the relationship between musical phrasing and gesture. In this case study, it became clear how the pianist accentuated the harmonic direction of the music both with ancillary movements and stillness.

Regarding the effects, while all the gestures chosen were detected and mapped, the palm height paired with the quantity of dry signal played through the speakers was the least effective. This was due to a technical problem regarding the nature of the effect. The acoustic piano, being louder than the processed sound even when played softly, easily masked the sound of the clean audio output. To overcome this, the volume output of the speaker was brought up, inevitably causing feedback loops in the system. This problem was highly dependent on the performance space, the placement of the speaker, and the microphone choice. In this case study scenario, using a pair of Genelec 8010A 3-inch speakers close to the piano and microphones the effect disappeared in the sound of the acoustic piano most of the time. In later developments of the system, when tested in performance spaces with dedicated piano microphones (DPA 4011 MMP-ES), or by controlling different sounds, the coupling had a much bigger effect and impact (Section 5.3). The second case study revealed that in a performance space, controlling the amplitude of a fixed media using the same gesture was as effective as the other sound-gesture pairs. This was due to the different nature of the fixed media sound—a choir sample— and the distance between the sound system installed in the venue and the piano microphones. This increased the feedback threshold, meaning that the modulated sound could be played louder without inducing feedback. The different nature of the sound instead, reduced the volume required before the sound was loud enough to be distinguishable from the acoustic piano. Further explorations led to the realisation that the modulation of the amplitude of the piano sound was possible in large venues even though it remained the most subtle effect due to the similar spectral qualities with the direct sound from the acoustic piano.

5.2.4 Discussion

After establishing the content and character of the musical material, to make the end goal of the project clearer to the pianist, the composer decided to provide subjective instructions in the score as to elicit use of the system by proxy, using terms such as "expressive" or "rigid" to describe individual phrases. However, this was changed during the course of the workshops. It became apparent that it was more useful to compose in opportunities to exploit the system with the notes themselves. The pianist's willingness to be expressive through movement was not an issue. The difficulty revolved around the sort of situations that needed to arise in order to use the system effectively. For example, inserting a long chord in order to be able to bend the pitch or a large jump to activate effects in space above the keyboard.

In Figure 5.4, marked by +, there is a jump of an octave in the right hand with the pedal sustaining the notes until the end of the bar. The pianist naturally raises

Figure 5.4: Score extract showing multiple uses of the + symbol.

their hand thus entering the 3D space above the keyboard that increases the amount of reverberation on the live sound from the piano. Both hands are able to control the effect separately and combined.

As can also be seen in Figure 5.4 marked by the + a longer note duration has been inserted into what, in other repeats, would be a rest. The longer duration allows the pianist to bend the notes with sideways movement of the hand.

It should be noted that the system was being adapted to the piece during the process of development in reaction to the way in which the pianist was moving and the goals of the composition. An example of this adaptation would be the 3D space above the keyboard, which was slightly different between the two hands: after having gathered some feedback, it was noticed that in the leaps, the left hand of the pianist reached lower heights than the right one. This was part of the pianists set of pianistic expressive gestures, that saw a more pronounced movement with the right hand than the left hand when performing ample leaps on the keyboard. For this reason, the system was tailored to lower the threshold of the left hand in order to accommodate an easier—and more natural—triggering of the effect. Having started to define the style of the piece with baroque and renaissance *tints*, the pitch bend effect triggered by the lateral movement of the hands was made more musical and subtle. This was achieved by applying a low-pass and high-pass filter respectively to the left and right hand, which affected continuously the cut-off frequency in relation to their position on the keyboard as explained in Figure 5.3. This meant that the right hand could pitch bend only high frequencies, and the left hand the low frequencies, enabling the pianist to perform a vibrato gesture with the right hand to modulate a melodic note while retaining a stable bass line, or vice versa. Through observing the workshops, it was revealed that the pianist was merging her expression through the system with her expression through dynamics and articulation; the more familiar the pianist became with the system the more she was able to use it in harmony with the other elements. The score, therefore, ended up not containing dynamics or articulation to leave the pianist to adopt the means of expression at her disposal freely.

Having written out the score with these additional points, a decision was taken by both the composer and the pianist that no detail was to be present in the score, rather just opportunities for the pianist to express through gestures. In this light, the idea of the pianist being able to ornament the music as in the renaissance and baroque eras became an apparent analogy. The piece was then titled *//aria alla francese* and is made up of fragmented baroque style materials. The "+" symbols seen on Figures 5.4 were the only indications that remained in the final score as points of reference for the pianist. The symbol, used by some composers historically to highlight a moment at which an ornament might be appropriate, was meant in this case to highlight a musical moment during the workshops and rehearsals when an expressive gesture surfaced naturally and thus might be appropriate once again. The final score should be seen as a strategically composed base for the pianist to expressively ornament with a mixture of the system and other modes of expression such as dynamics and articulation. The piece was performed to be audio and video recorded (Stenton and Granieri, 2018) and later published online.

The case study confirmed, following the results of the initial user testing (Chapter 4), that limiting the amount of modifications applied to the instrument led to a more comfortable experience for the user, that was able to retain to the full extent of her gestural vocabulary and expressive nuances. The additions that were made to the instrument, such as placing the Leap Motion on the top of the piano keyboard, limited the amount of invasiveness and discomfort to the pianist. Even though the Leap Motion was clamped on the music stand potentially obstructing the field of view of the pianist, the plastic arm resulted in being thin enough to avoid causing problems when reading the score. However this was also attenuated by the short length of the score, with longer scores and page turning this was identified as a possible issue. For this reason, Reach v1.0 provided a custom 3D-printed mount that was mounted on the piano frame avoiding contact with the music stand. The comfortable environment provided by the system enabled a deeper exploration of the effect of notation, rhythm and melodic content on the gestural nuances of the pianist with little to no system invasiveness. By tailoring the system to the movements of the pianist, the accessibility of the instrument improved and the learning curve reduced, enabling the pianist to use to the full extent the system from the first encounter with it. The ability of the pianist to take full advantage of the system increased in each successive workshop even though no practice was involved in between these three sporadic events.

The compositional aspect of the workshops mutated throughout the length of the case study from a more detailed approach, that tried to inform the pianist through non-conventional notation and text about the effects and movements required, to a sparser and more traditional notation. This was due to the fact that the more the system was tailored to the movements of the pianist, the less directions the pianist required to take advantage of it. It became apparent that an overly detailed score had a negative affect on the ability of the pianist to play and retain the pianistic technique. An overly detailed approach to scoring caused the performer to overthink the technique, and the movements required, resulting in a performance that resulted in being unnatural. In this scenario, the gestural nuances previously seen in the pianistic technique did not surface as often or as consistently, making it difficult for the composer to compose for the system. This amount of detail had a similar effect to the implementation of slow musical tempo: the dilatation of the tempo gave too much time to the performer to think, resulting in an unnatural performance.

In conclusion, it was uncovered that sporadic hints on the score paired with a medium tempo with an unstable rhythmic feel produced the best results. Eliciting gestural nuances through the rhythm and melodic content of the piece gave the performer the freedom to play with the system whilst maintaining the pianistic fluency.

5.3 Case Study 2: Richard Stenton (Composer)

The second case study produced a piece for the Cobalt Duo, a piano duo formed by Kate Halsall and Fumiko Miyachi, titled *12.jpg* (Stenton and Granieri, 2018), that explored the improvisatory nature of a graphically scored piece, and the interaction between two musicians and the system in a live performance.

The experience gathered during the first case study, led the composer to approach the second piece in completely different way. Having had almost three months of experience with the system, and being aware of its limits and capabilities, the composer decided to work with a piano duo. The difference in the approach was that for this piece, the composer decided to work with the two pianists without the Reach system for the whole duration of the compositional process. This approach provided useful insight on the concept of first encounters with the system, that later on led to the exploration of this aspect in depth through the second round of user testing.

5.3.1 Methodology

This case study was based on five workshops. Three workshops were conducted with the pianists and the composer, and two workshops were conducted solely with the composer. In this way the composer could compose and gather feedback on the piece and the movements of the pianists and transfer it into the system without the pianists ever experiencing it. During the workshops with the two pianists, the composer explained the gestures and the relative sound modulation mapping in order for them to picture the interaction and the sonic outcome. The day of the performance, the pianists had the chance to rehearse the piece with the system only once, and then performed it in a public concert.

Thanks to the insight gathered during the composition of the first piece (Section 5.2.4), this second case study aimed to investigate the relationship between graphical notation and gestural nuances. Contrary to the goal of the first case study, this piece aimed to investigate how the freedom of interpretation of a graphical score and the lack of direction could elicit gestural nuances in a performance scenario.

5.3.2 Sound-Gesture Mappings

Two of the gestures analysed in this piece came from the first case study, and were paired with different sound modulations. The palm height of the hands on the keyboard, that resulted in being too subtle of an effect in the first case study, was in this case used to control the volume of a choir sound file playing in the background. The lateral swaying of the hands instead was still coupled with the same pitch shifting effect of the previous case study. The score was entirely graphical and gave the pianist complete freedom of interpretation regarding the gestural component of the performance. The choice of having only two gestures controlling sound modulation parameters was made not to overload the pianist. Keeping the gesture recognition simple was key in the positive outcome of this performance.

5.3.2.1 Setup

As explained in the instructions (Appendix A.6) the piece was composed for piano duo with one pianist playing on the keyboard using Reach, and one pianist playing inside the piano. The piano was amplified with a stereo pair of DPA 4011 placed around the performer inside the piano, and the sound was fed through the Reach system to the sound system of The Lab, Royal Birmingham Conservatoire. The system played both the effected piano sound coming from the Reach system, the choir sound file modulated by the pianist hands and a regular auditory cue that signalled a page turning point.

5.3.3 Discussion

12.jpg is a piece that inherits the form from the contemporary music tradition, and adapts the classic composition for *instrument and tape* to be for *4 hands 1 piano + Reach + .wav*. The score (keyboard part viewable in the Appendix A.6) except for rough indication of the note without the octave indication and timing determined by the 10 second long bars, is entirely graphical and lets the pianists interpret how to perform the pitches indicated. From the instructional page (Appendix A.6):

> *The octave of the pitches is optional but should be the same for the entire page. It is preferable to adhere to the order of the notes dictated with gaps roughly according the layout in the score.*

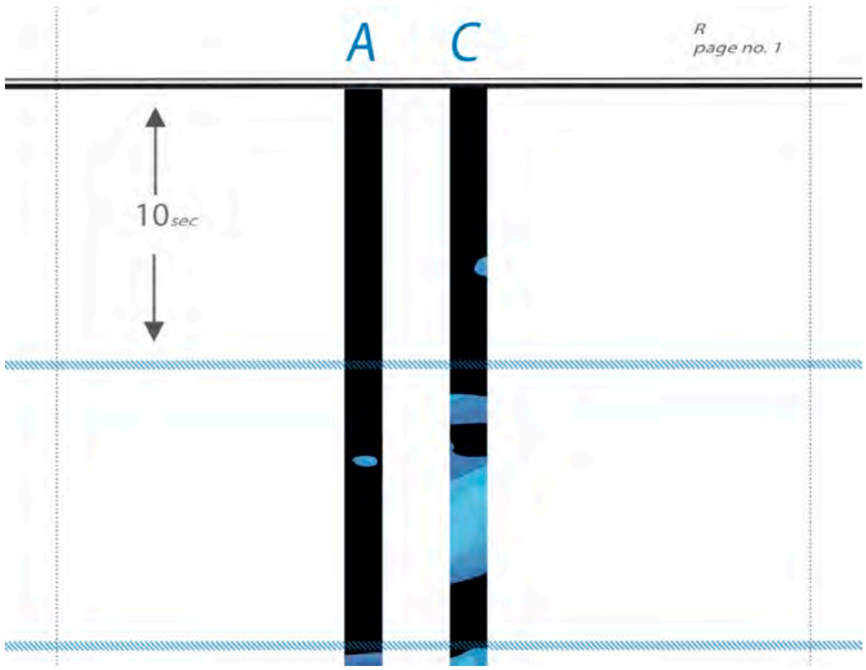

Figure 5.5: Score sample showing the graphical nature of the score. Light blue lines divide page in 10 second long sections, middle dark columns with blue graphics convey note to be played with the name at the top and expect the performer to interpret the pitches and expression.

> *In order to interpret the subtitles in the score we will use a sensor that tracks the movement of your hands. You will have the ability to manipulate the sound through sideways movement of the hand to bend the pitch and up and down movement of the wrist to fade in and out an additional sample (subject to change upon testing). Your movement will be tracked by a small sensor attached to the stand and will not impede on your movement in any way.*

5.3.3.1 The Performance

The methodology implied that the feedback was to be gathered from the composer throughout the workshops, and from the performers on the day of the performance during the only rehearsal before the concert. Due to time restrictions in the rehearsal before the concert the pianists had time to rehearse the piece with the Reach system only two times. During the rehearsal, the initial reaction of the pianist was one of surprise and confusion: the pianist was convinced that ample movements had to be made in order to trigger the effects. This caused the system to misinterpret the movements and to intermittently open and close the volume of the fixed media. Once the pianist understood that the system was tailored to the pianistic gestural nuances and did not require ample gestures, she was able to rehearse the piece a second time without any problems. Once finished rehearsing the pianist said:

> *That's much better. So now what you were doing in the other piece, is creeping back in because my movement when playing the notes is actually going with the palm moving.*

This reaction and the quick change of perspective from the point of view of the performer showcased how a system with a low degree of invasiveness reduced the time needed to become acquainted with it, thus making the first encounter more accessible. A short first encounter resulted in being enough for a professional pianist to perform a few hours later with the system in public. This shows how musical interfaces based on pre-existing instrumental technique and with a low degree of invasiveness can help overcome the barrier to musical creativity usually posed at first encounter with digital systems (Bullock et al., 2016).

The performance of the piece was successful and the pianists were able to perform the piece without problems. Moreover, an interesting dynamic between the performers arose throughout the performance. A gestural feedback loop between the pianist at the keyboard and the pianist playing inside the piano kept occurring throughout the piece. Because both the performers were *playing* the same piano, the pianist on the keyboard was able to pitch bend the sound of the piano being played from the inside. This was because the piano microphones placed inside amplified all the sound coming from the harmonic body of the instrument. This connection enabled the pianists to interact on a deeper musical level. Not only the actions of the pianist on the keyboard were affected by the one playing on the inside, but the pianist on the keyboard was able to manipulate the sound produced by the other performer through her gestures.

As noted by Jack et al. (2017) DMIs should be designed considering the body language of the performer, taking into consideration the intrinsic bottleneck and limitations that sensors add to the classical instrumental performance. In the case of Reach, AMIs should be designed considering the pianist's hand gestural nuances. What this case study helps uncover is that the even though the bottleneck that the sensor introduces in the interaction is still present, its importance can be reduced by limiting the amount of invasiveness of the system, using a touch-free mean of interaction, and encouraging the use of existing instrumental gestural nuances. This approach benefits both the performer, that will retain the full pianistic control and technique, and composers, that will be compelled to compose for a new and innovative system based on known instrumental dynamics and techniques.

Gestural nuances, as seen in the piece titled *//aria alla francese* can be encouraged by composing and scoring the musical piece without explicitly asking the performer to move in a certain way. By doing so, the composer is able to maintain the naturalness and flow of the performance whilst giving to the performer full control of the sound modulation parameters. A correct analysis and implementation of gestural nuances already present in the pianistic technique makes the interface easily approachable by performers with little or no prior experience. As seen in *12.jpg*, the two pianists were able to perform with the system after a brief first encounter that happened on the same day of the performance.

5.4 Case Study 3: Rosie Tee

Following the first two case studies carried out with a composer, two more case studies were conducted involving musicians with a different musical background. As explained in Chapter 3, professionally trained pianists, were taken into consideration with a focus on vocal performance, pianistic performance and composition. Following the first year of research, the user testing and the first two case studies, it became apparent that gestural nuances were surfacing more and more in a scenario where musical expressiveness was required. To further investigate the case, singer-songwriters were chosen because of their complete control over the artistic outcome of the piece. Not only they act as composers in the initial stages of the creative process, but they also end up embodying the emotional aspect of the piece when performing it. In this case it was hypothesised that gestural reverberations (Pierce, 2010), considered to give rise to a less stifled interpretation thanks to the prolongation of musical gestures, were going to be even more prominent that in the previously explored case studies.

For these two case studies, recent graduates from the composition department at the Royal Birmingham Conservatoire were asked to take part. The goal for them was to compose a piece following a similar approach to the previous case studies conducted with the composer. Following the concept of technology probes (Hutchinson et al., 2003) and iterative design processes through a series of individual meetings with the singer-songwriters, the system was tailored to their creative needs. This led to the exploration of a semi-longitudinal approach with the system. The singer-songwriters had multiple sessions to get acquainted with the system, see it develop according to their requests, and respond to their personal gesturality and needs in a live performance.

The first of the two case studies was conducted with singer-songwriter Rosie Tee, a singer-songwriter, composer and performer based in Birmingham. Rosie was the more "experimental" singer-songwriter between the two. Rosie describes her musical output as *between genres, fluctuating between classical orchestration, jazz-infused grooves and electronic samples.*[1] This led to a more open exploration of the system and its possibilities, even though she had never worked with gesturally controlled musical instruments before.

[1] https://www.rosietee.uk

Rosie received a formal education in composing, songwriting and piano that made her the ideal candidate for this study. While the case studies sought to explore creative practitioners with different musical backgrounds and interests, an advanced piano technique was a required criterion due to the research topic revolving around the use of pre-existing gestural nuances in piano technique. Moreover, this case study aimed to explore the use of the system with creative practitioners that were completely in control of their craft: from the compositional stage to the performance. In addition to all the required criteria, Rosie had previous experience with electronic instruments such as keyboards and she was acquainted with music software, digital sounds and effects and the overall music production workflow. This added experience led to an interesting exploration of the system and gave useful insight on the ability of understanding and co-developing a technology-based system with an artist with experience of technology.

5.4.1 Methodology

The case study developed over a total of three workshops, one rehearsal, a public concert and the recording of the piece titled *Wax & Wane* (Tee, 2018), in a live setting. As arranged in the initial meeting with the artist, both an audiovisual recording of the piece composed and a recording to be featured in the upcoming EP of the artist were going to be produced as a result of this collaboration.

The first meeting was an informal brief listening session of previously recorded material with the system, aimed to elicit the suggestions of the artist with regard to where the piece should have gone in terms of style, genre and effects. Being already familiar with audio effects and music technology from a production standpoint, this listening session was more tailored towards the analysis of audio effects on tracks that the artist liked and hypothesising potential gestures that could have controlled them.

Between the first informal meeting and the first workshop, the artist had already noted a series of gesture-sound pairs and ideas that she wanted to try out with the system. From the first instance, she was exploring extramusical gestures to control sound parameters. In an email communication she said:

> *Could we try a system that when my hands are raised it triggers some kind of heavy reverb to my vocal? It'd be nice if I could effect the piano output too but not sure if the same movement can trigger two outcomes?*

The development process was agile: after only two meetings the system was at a stage where the artist could have rehearsed with it. From that moment on the artist asked to rehearse with the system alone for sessions up to an hour, in order to further work on the compositional aspects of the piece. After every solo rehearsal, feedback would be gathered through a brief informal discussion in which notes about the rehearsal were shared. These notes were both informative about the compositional process, and technical regarding the system, the interaction and the effects.

5.4.1.1 Workshops

In the first workshop, the artist explored the system and the gestures defined following the first informal meeting. Having already defined the main gestural content and the themes of the piece, she used the time exploring musical ideas, and asking for the system to be tailored more to gestures she had envisioned. Once the first workshop was over, the artist requested for the system to be tailored in between workshops in order to use most of the workshop time for rehearsal and composition.

The second workshop developed as another explorative session with the same system setup defined at the end of the first workshop. A sonic and compositional exploration was conducted in a similar fashion to the first workshop, bringing to light some new ideas with regard to the gesture-sound mappings of the system. At the end of this second session, the artist expressed the desire to control two different effects with each hand performing the same gesture. This differentiation of effects between the two hands was similar to the different filters applied to each hand seen in the second case study.

Following these two guided workshops, the artist requested two exploratory and compositional rehearsals with the system alone. Considering that the system had been finalised, the artist was left alone for two one-hour session to rehearse, compose and refine the piece.

Wax & Wane was performed at the *Composer's Night* at the Royal Birmingham Conservatoire, and was later recorded as a live audio-visual performance for online release (Tee and Granieri, 2018).

5.4.2 Sound-Gesture Mappings

The final system was based on two main gesture-sound pairs. The two gestures chosen were the left and right hand vertical displacement from the keyboard shown in Figure 5.6. These ample gestures were discussed during the workshops and were seen by

Figure 5.6: Two frames from the *Wax & Wane* recording showing the left and right hand gestures chosen to control the sound modulation parameters.

the artist as part of her gestural vocabulary: the artist felt that she performed those gestures on the piano as part of her way to convey emotion. The vertical displacement of the right hand from the keyboard—after a note or chord was played—controlled the dry/wet parameter of a reverberation effect, similar to one of the effects used in the first case study. The reverberation in this case affected both the voice and the piano sound with separate parameter settings. This was due to the desire of the artist to create a homogeneous soundscape. The parameters of the reverberation were tailored to the sound source to replicate the original sonic idea of the artist. During the workshops, it was noticed that without having separate reverberations for the vocal and the piano, the resulting sound would loose definition and make the whole sonic experience *muddy*. However, the gestures chosen by Rosie had a peculiarity: the artist expressed the desire for the gestures to leave a *trail* even after their end (Figure 5.7). After a gesture had been performed, and the artist went back to playing the next section of the piece, the system was designed to keep effecting the sound and slowly fade away. This made the gesture reverberate in the sound and accentuated the expressive nature of the motions of Rosie's hands.

Figure 5.7: Overview of the gestural reverberation in relation to the amount of effect being controlled.

Due to the gesture being naturally present in the technique of the artist on both hands, a specular movement of the right hand was used on the left hand and mapped to a stereo panned low frequency oscillator (LFO) that modulated the sound of the unprocessed piano. This effect, originally discovered in one of the pieces used as inspiration in one of the workshops, was chosen by the artist to modulate the acoustic piano sound in an *unnatural* manner. For this effect the artist decided to control the amount of modulation with her gesture while the frequency of the LFO changed randomly with time. This meant that depending on when the artist performed the gesture, the LFO could have been a modulating wave of any frequency between 0.5Hz and 3Hz. The choice of keeping the LFO threshold low gave the best result, avoiding to generate the composite waveform typical of amplitude modulation with a carrier waveform in the audio frequency range. While the two hands had different effects paired with the same gesture, the range of the movement between the hands was different and tailored to the effect. The threshold of the right hand was much broader as the artist wanted to have a more ample movement controlling the wet/dry parameter of the reverberation. The ability to perform more easily an ample movement with the right hand was linked to the previously discussed embodiment of the gesture in the expressive gestural vocabulary of the artist. Considering that this gesture was performed more often with the right hand, it was more natural for her to control the reverberation parameter in a nuanced manner with this hand. Through the rehearsals it was uncovered that with more movement space she was able to select better the amount of reverberation to suit the musical moment. For the left hand instead, the movement was much shorter as the artist expressed the need to have immediate access to the full breadth of the LFO modulation. In performance, these nuanced differences in the gestures of the right and left hand enabled the artist to control better the sound effects.

Initially, the singer-songwriter was keen on implementing three effects in the piece, adding a distortion effect to the two already listed, but towards the end of the project she decided to limit the system to the two most effective ones. This was later discussed and reconnected to the extramusical nature of the gestures chosen. Because of the ample movements required to activate the first two effects, adding a third one would have required her to alter too heavily her musical performance to control confidently all the effects.

5.4.3 Discussion

The system developed through this case study explored different gestural elements from the previous ones. Instead of tailoring gestural nuances, the artist preferred utilising ample—almost theatrical—gestures to control the sound modulation parameters. While the learning and implementation of these gestures have the potential of interfering with the ability of a pianist to play freely, in this case the artist herself sought and chose these gestures. For that reason the amount of invasiveness did not have an impact on her ability to perform, or perform expressively. This case study shows that the presence of a new gestural vocabulary does not cause a disruption in the pianistic performance per se. If the system is created through a collaborative design process with the performer rather than being imposed, the gestures chosen will become completely integrated in the performance, whether pertaining to the pianistic technique or whether extramusical. It is possible however that the music technology background of the artist enabled her to be more open and accepting of both the interaction and the sonic result of the system.

Another factor that could have contributed to the acceptance of extramusical gestures and the lack of effect of these on the performance could be the small number of gestures chosen and the tight connection of the gesture with the concept of gestural reverberation (Pierce, 2010). The gesture-sound pairs chosen were based on the same gesture to control the sound modulation parameter: the vertical displacement of the hand from the keyboard after having played a note or a chord. The fact that only one gesture was selected to modulate the sound parameters was definitely a factor in limiting its impact on the ability to play fluently of the performer. This was also highlighted by the initial idea of the artist to implement a third effect that was then abandoned due to the risk of making the system *too hard to control*. While the extramusical chosen gesture did not seem to disrupt the performance flow of the artist, she seemed to be aware that adding more gestural control would have possibly caused problems in the performance.

Ultimately taking a closer look at the gesture chosen by the artist, while being categorised as an extramusical gesture, it could also be seen as an expressive gesture part of the personal technique of the pianist. By raising the hands after a note or chord is played, the artist is reverberating with her gestures the sound of the ringing chord, thus following corporally the sound envelope. Analysing further the sound-gesture pair, the artist specified in the development stage how she wanted the sound effect to continue after the gesture had ended. After a note had been played the artist desired her gestural reverberation to continue in the sound effect chosen while the hand went back to playing. This would help explain even more why the artist in question did not find such a wide and ample gesture as disruptive in performance.

5.5 Case Study 4: Chloe Knibbs

The fourth case study was conducted with singer-songwriter Chloe Knibbs. Similarly to what explained in Section 5.4, Chloe Knibbs was chosen for this case study due to her interdisciplinary studies in composition, piano and songwriting. Differently to Rosie Tee, Chloe had little prior experience with technology. This factor was key in the case study as it enabled to explore and analyse the first encounter of the artist with a new AMI, and investigate issues of accessibility for classically trained musicians with novel electronic interfaces.

It was expected that on the first encounter with the system the reaction of Chloe would have been similar to the classical pianists in the first user test (Chapter 4). However, results were different. The first problems started arising in the collaborative design aspect of the case study. When Chloe was asked to think about a possible setup for the system, a possible combination of effects, and a possible pairing with gestures, she struggled, faced with a wide number of possibilities and no prior experience with technology. Overall Chloe found it difficult to understand the interaction between her gestures and the sound modulation parameters. While the effects provided an intuitive way to interact with sound parameters and the initial explanation of the system was clear (i.e. the lateral swaying of the hands paired with the pitch shifting of the sound), proposing other gesture-sound pairs resulted in a slow and complex task.

Once the meetings became regular, and the performer became acquainted with the process of thinking about an idea, translating it into a possible outcome for her personal performance, and then trying it on the system, the system and piece started to develop.

5.5.1 Methodology

The methodology followed the same structure as in the third case study (Section 5.4.1). A series of three workshops were conducted to design the interface around the needs of the performer and were followed by a rehearsal, a public performance and an audio-visual recording. The piece, composed throughout the workshops and recorded in a live setting, was titled *Listen* (Knibbs, 2018).

5.5.1.1 Workshops

The first workshop was purely explorative. An informal conversation in front of the system was conducted with some brief practice time to explain and convey both the potential of the system and to align the goals of this collaboration. The system provided to the artist was similar to the one tested in the first user test (Chapter 4): the gesture recognised was the lateral swaying of the hand paired with the pitch shifting of the live piano sound.

Following the first workshop, the artist requested for previously written scores for the system to be viewed to see how other people had composed for such a system. In reply to this request, both the audiovisual recording and the scores produced from the case study conducted with the composer (Chapter 5) were shared. The artist showed particular interest towards the mapping of the lateral swaying of the hands with the pitch-shifting effect showcased in the video, described by her as *circular rotation of the hand on the key*.

The second workshop had been planned as an explorative and compositional workshop. An exploration of gesture-sound modulation couplings and overall piece structure was conducted. During the course of the workshop particular interest was brought towards what can be considered the inverse coupling of gesture-sound modulation: stillness-sound modulation. The artist asked if it could be possible to explore the ability of the system to detect stillness and lack of gestural nuances. A brief test was conducted using the reverberation effect used in previous case studies and inversely mapping it to the movement of the hands: after a note was played, the less the pianist moved, the more prominent the effect. Having observed the effectiveness of the novel mapping, this led to the exploration of the relation between stillness and audio distortion.

The artist consolidated her ideas into a final piece between the second and third workshop. The third workshop was used to test the final ideas—both musical and in terms of gesture-sound pairs—to make sure the system satisfied the needs of the artist. The final encounter was a rehearsal in light of the upcoming concert where the pianist rehearsed unsupervised with the system for two hours. This rehearsal was required by the artist to gain trust in the system even without the need of a technical supervision. While she was able to rehearse with the system alone, at this point of the development the systems still had to be set up for her by technicians.

5.5.2 Sound-Gesture Mappings

The system developed for the purpose of this case study used two of the previously used gestural nuances and one new gesture to modulate the sound of the voice and the piano. The three gestures analysed and recognised were the lateral swaying of

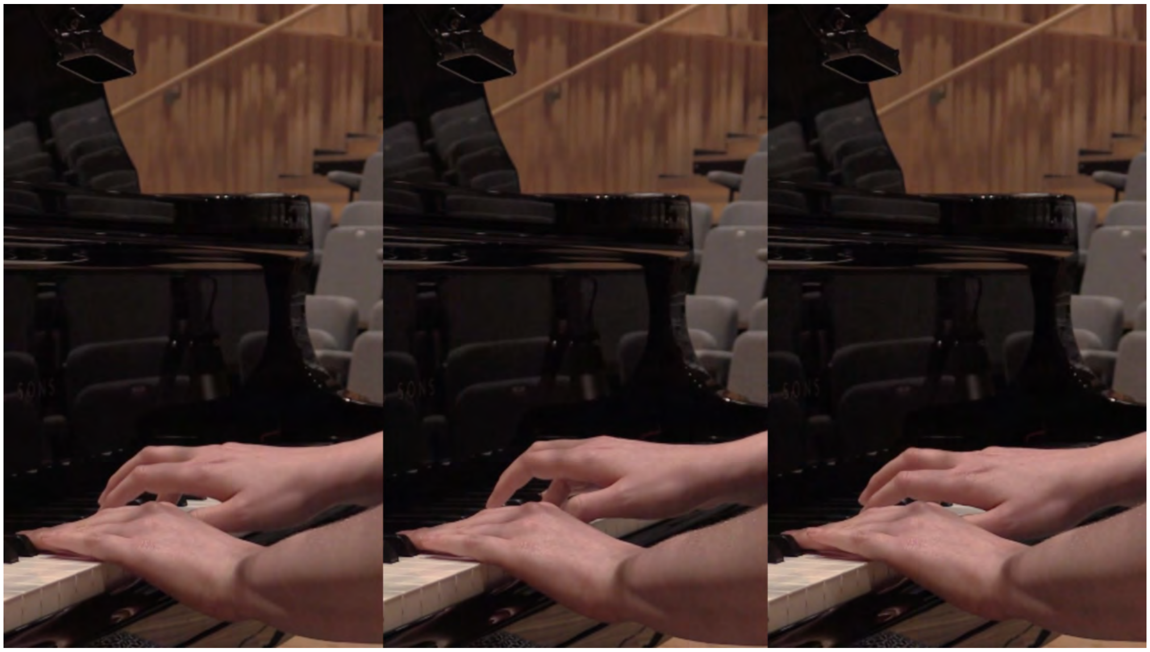

Figure 5.8: Consecutive frames from the recording of *Listen* showing the lateral swaying of the right hand used to control the pitch shift effect on the unprocessed piano sound.

the right hand, the heigh of the palms of both hands, and the stillness of the hands after a note or chord had been played.

As in previous instances of the system the swaying of the right hand (Figure 5.8) controlled a pitch shifting effect of the live piano sound capped at ± 50 cents of a tone. The palm height of both hands (Figure 5.9) controlled the amplitude of the unprocessed sound of the piano delayed by one second. The stillness of the hands detected after a note or chord had been played was mapped to a bespoke vocal distortion effect that combined classical distortion and a square wave LFO resulting in what can be described as a *distorted stutter* effect. In opposition to all the gesture-sound mappings explained until this moment, the absence of movement

Figure 5.9: Frame from the recording of *Listen* showing the palm height of the hands controlling the amplitude of the unprocessed piano sound delayed by one second.

detected caused the amount of the effect to increase over time. So rather than having a direct causal relationship—performing the lateral swaying of the hand caused the piano sound to be pitch-bent—the effect increased over time the more the hands were still. This was achieved by reversing the moving average of the data coming from the vertical axis of the palm of each hand to the amount of distortion. The more the hands moved over time, the higher the average and thus the lower the amount of distortion. This was developed in order to give Chloe control over the last part of the piece where she wanted the effect to gradually grow reaching a point where the distortion was so strong that the words would have become hard to understand. This meant that from approximately half way through the piece, with the increased stillness of the hands, the rolling average would have become gradually lower, increasing the amount of distortion on the sound.

5.5.3 Discussion

The system developed for the purpose of this case study differed from the previous ones because of the way the performer, Chloe, understood and embodied the nuance gestures used to activate the system. The most peculiar effect required by the performer was in complete opposition with what had been developed and required until that moment. A distortion effect applied to the voice was linked to the absence of nuanced gestures detected. This could be linked to the concept of negative gesture described by Lewis & Pestova (2012). In their paper, negative gestures are described as "gestures that are perceived through the absence of their opposite". Here the *sounding gesture* is heard through the distortion of the vocal and perceived through the absence of movement in the pianist's hands.

Instead of a typical gesture-sound mapping, a mapping of the average of the movement was used to detect the absence of gestural nuances. This effect was applied to the final section of the piece (Knibbs, 2018), in which the pianist kept playing the same chord progression and singing the refrain *The more I lose my voice*. The lack of gestural nuances distorted the sound, resulting in a very effective sound-lyrics coupling.

Along with the inverse gesture-sound mapping effect, the final patch and piece presented another peculiar aspect, that led to a non-typical exploration of the possibilities of the system. The palm height of the hands, was controlling a clean feed of the piano delayed by one second. This effect was discovered in one of the workshops and was originally caused by the system malfunctioning in a particular occasion. The singer-songwriter found the delay of the clean feed of the piano an interesting feature to implement. In complete opposition to the first case study, in this occasion the system was entirely tailored towards gestural nuances, and in one case on the absence of them. The same applied to the sound effects chosen, and the amount of effect applied. Everything was purposely chosen to be very subtle: from the amount of gestures required to activate a certain effect to the effects itself on the sound.

5.6 Chapter Summary

The four case studies explored in this chapter have provided insights into the development of the system. Moreover, the case studies contributed to an understanding of the original research question investigating not only the user base of the Reach system but also the impact of the musical background and experience of the musician on the relationship and interaction with the system.

One important finding is how practitioners coming from the same musical background with the same musical education approach the system differently depending on their previous relationship with technology. In case studies one and four the lack of direction caused the musician to lose control, not knowing what gestures to perform or finding it hard to compose for the system. Both pianists came from a classical background and had little experience in dealing directly with technology. The only experience with technology the first pianist had was in a live electronics scenario in which she was asked to follow a score and an *invisible technician* (Shapin, 1989) was in charge of the live electronics from behind the scenes.

Case studies two and three were conducted with performers that—while still coming from a classical background and with a traditional music education—were more acquainted with technology. In the first case, the Cobalt Duo, was able to utilise the system in an extreme context where almost no rehearsals were done with the system prior to the concert and the graphical score provided little to no strict guidance regarding the performance and the gestures. In a similar way, Rosie Tee, without any prior experience with gestural controllers but with a basic understanding of music technology with regard to music production and vocal processing, was able to immediately understand the capability of the system and radically disrupt the nuanced gestures used until that moment to create a bespoke gesture vocabulary that suited her personal playing style and performative act. Both these performances elicited a combination of gestural nuances and extramusical gestures that—from the point of view of the performers—had no disruptive effect on the performance or the piano technique.

These case studies have shown how the efforts to make the system accessible and non invasive for classically trained musicians resulted in different musicians exploiting these characteristics in different ways. Pianists with no prior experience with technology seemed to take advantage of these features to slowly transition from a purely instrumental type of playing to one intertwined with technology. The use of gestural nuances in these cases was appreciated and taken advantage of to aid the performance and make the overall experience less invasive. Pianists that were already acquainted with technology instead seemed to use the system's features to build upon the accessible entry point and explore their personal gestural vocabulary. These gestures—whether nuanced and embedded in their technique or extramusical— provided nevertheless the ability to perform with technology in a fluent and natural manner.

Chapter 6

Comparative User Testing

The second user testing was designed to investigate if and to what extent invasiveness in DMIs and AMIs affected the ability of pianists to freely improvise in a first encounter with the instrument. To do so the Reach v1.0 system, mounted on a grand piano, was compared with two existing keyboard interfaces, that enable real-time sound modulation providing different levels of invasiveness, TouchKeys (McPherson, 2012) and ROLI Seaboard (Lamb & Robertson, 2011). The test was conducted with six jazz pianists. The aim of this user testing was to investigate the different aspects of the relationship between pianists and the musical instruments they were presented with in an improvisational scenario. The focus of this study was the ability of users to freely improvise on a keyboard-based DMI or AMI at a first encounter—or without prior extensive experience with it—and the relevance of the instrument providing this affordance. Moreover, knowing that different gestural controllers came with varied levels of precision and control over sound, the test investigated through the final interviews with the participants the relevance of this element when compared to the affordance provided by the instrument itself. The ability of the pianists to freely improvise at a first encounter with the instrument was also compared to the degree of invasiveness.

The first two key points sought to investigate two concepts regarding instrument interaction: freedom of playing, learning curve and aural invasiveness. While usually these elements require longitudinal studies with tracked progression of the musicians, in this comparative user testing these concepts are to be applied solely to the first encounter that the musician has with the instrument.

In the context of this test, freedom of playing is interpreted as the amount of freedom that the user felt in the first encounter with each particular interface. The freedom of playing was also explored as an element that could facilitate the spontaneous exploration of extramusical gestures (as seen in Chapter 5.4). This element will be investigated by combining the video recordings of the improvisations and the responses to the questions in the semi-open interview.

Similarly, the concept of learning curve is to be interpreted as the amount of time needed for the user to get acquainted to the interface during the first encounter. This point will be investigated mainly throughout the interview and is closely tied to the freedom of playing.

The aural invasiveness aspect of this exploration follows the results of the first user testing (Chapter 4) that saw the classical pianists struggling more with the aural elements of the system than the gestural ones. In this comparative scenario a further exploration of the topic throughout all the proposed interfaces aimed to elicit the importance of aural invasiveness.

6.1 Setup

6.1.1 Instrumental Choice

The study compared three instruments that pertained to three distinct keyboard-based digital instruments categories. The Seaboard was chosen to represent what can still be considered a keyboard-based DMI with some major changes regarding both the hardware and the interaction. The instrument—apart from taking the keyboard layout as inspiration—modifies everything else: from the key spacing,[1] to the wedge shaped keys (Dahlstedt, 2017) all the way to the key interaction that shifts from the classical moving keys of a keyboard to a pressure and position sensitive silicone surface. With this different interaction plane, the Seaboard combined with a bespoke MIDI communication protocol now adopted as a standard, was chosen to be the instrument in the test that provided control over the piano sound with the steepest gestural learning curve. Being so different from a traditional keyboard, researchers like Dahlstedt (2017) noted that the technique required was not that of a pianist, making it an instrument tailored mainly to non-musicians.

The TouchKeys, an augmentation of an existing electronic keyboard, was chosen to represent the space in-between Reach, an augmentation of an acoustic piano,

[1] Smaller than a classical keyboard in all models except for the ROLI Seaboard Grand, gone out of production in 2018

and the ROLI Seaboard, a keyboard-based DMI. While both the Touchkeys and the Reach system could *potentially* be mounted both on a digital keyboard and an acoustic piano, the choice of having them mounted on two different instruments was purposeful. While TouchKeys is marketed as being able to transform any piano-style keyboard into an expressive multi-touch control surface, most of the use cases showcasing the system in action used the system applied to a MIDI keyboard.

As previously shown in Chapter 2 there are multiple prototype systems that provide the user with gestural control over sound modulation,[2] however the test was aimed at considering only instruments that could be used in a real world scenario. For this reason, prototypes not easily available to the general public were not taken into consideration and the TouchKeys used was the one sold pre-mounted on a Novation Impulse 49 electronic keyboard.[3]

6.1.2 The Layout

The three systems used in the testing, the Reach system mounted on a Steinway & Sons grand piano, the Seaboard and the TouchKeys, were placed in a triangular shape, in order to accommodate the pianist in the centre and comfortably switch from one system to the other as seen in Figure 6.1. All the three systems were active at all times, to limit the amount of setup time during the actual test. All three systems were being heard through two Genelec 8030A speakers placed at each end of all three keyboards. There were three speakers in total, with one speaker placed at each vertex of the triangle, meaning that each system shared one speaker that was

[2] i.e. Touch - Keith Hamel, Hayes (2011), Yang & Essl (2012) and Zandt-Escobar et al. (2014)
[3] http://touchkeys.co.uk/product/touchkeys-novation-impulse/

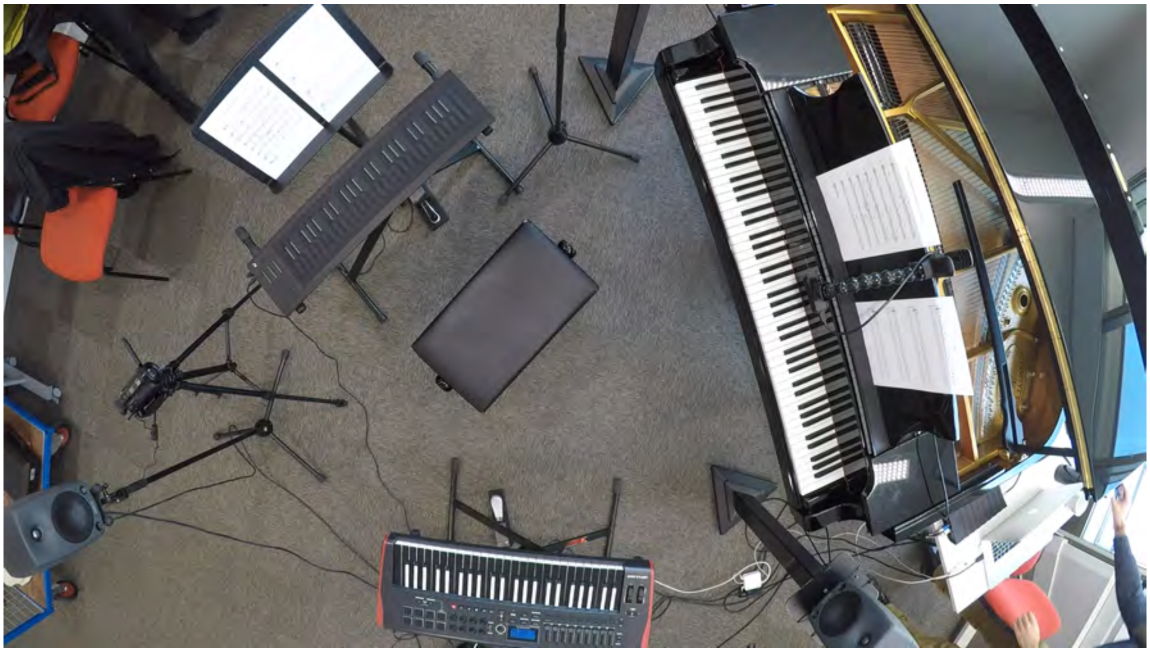

Figure 6.1: Comparative testing keyboard and speaker setup.

switched between one and the other via software routing in between the individual tests.

The sound fed into the Reach system was captured by an Audio Technica AT4040 cardiod microphone and fed into a Mac Mini 4,1 through an Edirol UA-25EX audio interface. The Leap Motion data, parsed through Reach on a Lenovo ThinkPad Yoga 260, was then fed in via Ethernet to modulate the sound parameters. To split the computing power between more machines, and to avoid any overloading issues, the audio, once processed, was sent to a Macbook Pro 11,4 through a TC Electronics Impact Twin audio interface to be recorded. This was solely a precautionary measure to avoid losing data during the tests; the system had been run through a single laptop setup before with success. The latency of the system was below the perceivable threshold of 15 milliseconds, and was thus disregarded; the destructive nature of the effects and the loudness of the acoustic piano itself had been observed to mask any delay occurring in previous tests and case studies.

Both the Seaboard and the Touchkeys were set up to play piano samples. The sample library used was the one included with the ROLI Equator software[4] because of the good quality of the samples, ease of setup and mapping of parameters and compatibility of the software itself with other MPE devices.

6.1.3 Effects

The effects used to modulate sound were coherent between the three systems with slight customisation to the each individual keyboard interface. This choice was made to make the user aware of the tradeoff between the three keyboard interfaces. The more invasive the interface, the more control and options. The Seaboard, considered the most invasive interface of the three because of the complete redesign of the keyboard itself, covered in touch and pressure sensitive material, was set up to enable sound modulation with four different gestures:

- Vibrato Effect - Lateral swaying of the finger(s)
- Amplitude modulation - Pressure of the finger(s) (if done fast enough, tremolo effect)
- Detuning Effect - Vertical slide of the finger(s)
- Continuous Glissando - Glissando on the bar above and below the keys

[4] https://roli.com/equator

The TouchKeys was set up to enable sound modulation with three different gestures:

- Vibrato Effect - Lateral swaying of the finger(s)
- Amplitude modulation - Pressure of the Finger(s) (if done fast enough, tremolo effect)
- Detuning Effect - Vertical Slide of the finger(s)

The Reach system instead, was set up to keep two gestures in common with the other systems, but offered a third gesture. This was done to make users aware that while touch-free interfaces lose control and precision when modulating effects—as shown by Wilson (2010) that explored the use of cameras as touch sensors—they are able to expand the gestural variety; in this case, the ability of tracking the natural leaps of the hand above the keyboard. The gestures implemented were the following:

- Vibrato Effect - Lateral swaying of the hand(s)
- Tremolo Effect - Height of the palm(s)
- Reverberation - Height of the hand(s) from the keyboard

6.2 Methodology

Each test lasted an average of 1 hour 18 minutes with a standard deviation of 29 minutes. Subjects were briefly interviewed about their pianistic background, current knowledge and experience with electronic music and digital instruments. To frame the results as accurately as possible, the users were also asked to describe their knowledge of the three systems, and if they had previously played or performed with them. This background information was later linked to the ease of use, learnability and approachability questions. Before the test, the users were also asked if they knew and had played or improvised already on the two proposed pieces: *Goodbye Pork Pie Hat* by Charles Mingus (Appendix A.7) and *Musica Ricercata n.7* by György Ligeti (Appendix A.8).

Considering the common background between all the users, the decision was made to push the two improvisations in two different directions to elicit from the musical performances different ways of playing. The first piece, *Goodbye Pork Pie Hat*, was chosen as representative of a jazz standard. Users were expected to be familiar with the piece, and thus brought to perform a more typical jazz improvisation. This was meant to provide not only confidence while improvising, but also draw out a series of jazz vocabulary automatisms that would have triggered the three systems in different ways, without asking the pianist to think about the gestures to perform.

On the other hand, *Musica Ricercata n.7*, was chosen not only because of the modal nature of the melody, but also because not all users were expected to be familiar with it. All the elements except for the melodic line, had been removed from the score, possibly laying out the ground for a freer improvisation, that explored all the encountered effects in the initial 10 minutes of playing on the instrument.

After the interview and a brief explanation of the system, the order in which the users would play the three systems was revealed, and the pianists were given 5 minutes to try the system and get comfortable with the sound and effects coming from the speakers. Having three separate instruments and six users, the order in which each pianist played the system was different every time. This randomisation had the aim of minimising the effects of one system on the perception of the next. When the users felt they were comfortable with the systems and all the gesture-effect couples, the pianists were asked to perform two five-minute improvisations.

The first improvisation was on *Goodbye Pork Pie Hat*. The pianists were asked to play the main theme and follow the score once, and then to improvise freely for the remaining time. All of the users were asked to play the main theme once in order to have a similar baseline.

The second improvisation was on *Musica Ricercata n.7*, and followed the same process as the first improvisation.

At the end of each section on one system, the users were asked to fill in a UEQ, the same used in the first round of user testing. Each subject was then asked in the final interview about the experience and the systems. The interview was meant to be open enough to enable the users to talk freely about their experience but had four questions revolving around specific topics, aligned with the aims of the comparative study, that were meant to investigate some key aspects of the interaction:

- Ability to play freely
- Learning curve
- Difference in precision when modulating sound
- Ability to transfer pianistic technique

After each question, the users were asked to rank the interfaces from best to worst regarding the current topic. This was useful to gather not only data thanks to the

UEQ, but also to compare clearly the ranking and comments gathered during the interviews and pinpoint which elements of a particular interface influenced which traits.

6.3 Reach v1.0

Figure 6.2: Overview of the Reach system.

The final iteration of the Reach system, named Reach v1.0, was built to gather the feedback and knowledge gained throughout the three years of research and make them available to the general public. This version of the Reach system is composed by a Leap Motion, the Reach App developed for Windows that gathers the data from the Leap Motion and makes it available via OSC communication protocol and a series of Pure Data patches responsible for the sound modulation.

Reach v1.0 was developed to work on a variety of computers, including embeddable development boards such as the LattePanda[5], to provide the broadest options to the users in terms of computational power. The work done to use the system with such development boards resulted in its optimisation. The system is now light enough to work on low powered laptops and has been further tested with inexpensive audio

[5] https://www.lattepanda.com

hardware (e.g., the Zoom U24 Audio interface and a general purpose microphone) in order to provide accessible options also regarding the price and implementation.

6.3.1 System & Data Handling

Figure 6.3: Picture of the the Reach system mounted on a Grand Piano using the custom 3D printed mount.

The Reach system uses a custom-made homonymous application, written in C++ using the JUCE framework and Leap Motion's VR oriented Orion SDK (Guna et al., 2014), to capture hand tracking data from the Leap Motion. The Leap Motion is placed approximately 30cm above the piano keyboard, mounted on a custom 3D printed piano stand that avoids interfering with the field of view of the pianist (Figure

6.3), where positional data from the joints of the fingers, palms and wrists is tracked by Reach, encoded and transmitted as OSC messages. This tracking data, is then received by a Pure Data patch, where it is mapped to sound modulation parameters as required.

The high level of precision of the Orion SDK negates the need to de-noise the incoming data. Peak detection audio analysis is performed on the incoming piano signal in Pure Data. When a note is struck, the positional data received from the Leap Motion is re-centred, meaning that no audio modulation would occur unless the hand moved away from the current position after striking a note. If the hand moved and a key was struck, the process repeated. This prevented constant and erratic behaviour in the audio processing.

6.4 Discussion

At a first approach, most of the pianists appreciated the ability to modulate sound parameters with their gestures on all three instruments. This happened consistently across all three keyboards, regardless of the order in which the pianists were asked to play them, and was also independent from the musical background of the each of the musicians: whether the user owned electronic gear, or was completely extraneous to the environment, this kind of gesture recognition technology remained fascinating. After the first five minutes of exploration with the individual interfaces, the users started commenting on the instrument and the experience. In order to gather an unbiased set of quantitative data through the UEQ, any intervention from the researcher was limited throughout the test, and the users were asked to keep the comments for the final interview when possible.

After every test, the pianists were asked to complete a UEQ that evaluated the user experience through efficiency, perspicuity, dependability, originality and stimulation. Figure 6.4 shows the average values from all the six users, for each keyboard, focusing on the attractiveness, perspicuity, efficiency, dependability, stimulation and novelty of the interface.

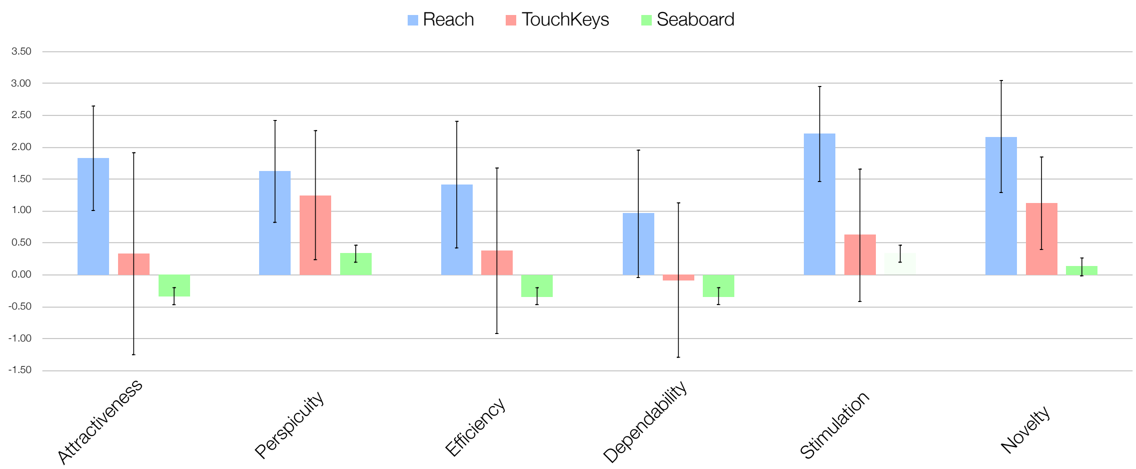

Figure 6.4: UEQ interface comparison showing mean and standard deviation for each value analysed.

In the same graph, as noted in the legend, the three interfaces are represented with three different colours: blue for the Reach v1.0 system, red for TouchKeys and green for Seaboard. Scoring an average above 1.50 in every category, with a peak of 3 in novelty, the Reach system was considered the one with the best interface, for playing and modulating in real time a piano sound. It is clear that, as for the first test iteration, the high values throughout the board were due to the non-invasive character of the system and the ability to play on an acoustic piano.

Interest			
Ranking	Reach	TouchKeys	Seaboard
First	5	1	2
Second	1	2	1
Third	0	3	3

Table 6.1: Table ranking user interest in the interfaces

Freedom of Playing			
Ranking	Reach	TouchKeys	Seaboard
First	6	0	0
Second	0	5	1
Third	0	1	5

Table 6.2: Table ranking user freedom of playing in the interfaces

Learning Curve			
Ranking	Reach	TouchKeys	Seaboard
First	4	3	0
Second	2	2	1
Third	0	1	5

Table 6.3: Table ranking user learning curve in the interfaces

The same trend can be seen in the rankings, shown above in Tables 6.1 , 6.2 and 6.3, of the three interfaces according to three factors: interest, freedom of playing and learning curve. The users were asked to rank the three interfaces according to these factors as a last question in the interview. While overall the Reach system was considered at the top in every ranking, in one case, one user considered it tied[6] with the TouchKeys both regarding the interest factor and the ease of learning. The same can be seen between the Seaboard and Reach where the two interfaces tied on the interest factor, while the Reach system was considered the easiest interface to learn and approach. It is clear that the data is non statistically valid due to the limited pool of users, but it definitely gives an idea of what the data from a large scale user testing might look like.

6.4.1 The Feel of an Acoustic Piano

During the interview process at the end of each test, multiple topics surfaced that shined a light on the results of the questionnaire. One user said:

> *The first thing I notice about an instrument is the feel (of the keys). That's the reason why I find many keyboard instruments really bothering to play. You don't have the same feel of a piano or a Rhodes for example, that has much more (tactile) feedback. I notice the difference of feel also from piano to piano.*

This appeared to be a recurring theme throughout all the tests and interviews where all of the users were drawn towards the Reach system because it allowed them to play on an acoustic piano. While this element could be seen as source of bias in

[6] In the tables, a tied ranking has been arbitrarily shown as the highest rank for both the tied interfaces. For example: in the mentioned case where Reach and TouchKeys were considered equal regarding the interest factor, a first mark has been assigned to both interfaces in the table. This is the reason why in some cases the sum of the rankings is seven or eight, even though the pianists taking part in the test were six.

the test results, it should instead inform the reader on the current state of keyboard-based AMIs. Considering the lack of keyboard-based AMIs that take full advantage of the acoustic piano sound and feel, providing a real-world option to the musician in terms availability and accessibility, the development of Reach is seen by most of the participants as a first step in the right direction. This kind of practice could be considered a *design exploration*.

> *Design exploration seeks to explore possibilities outside of current paradigms, to transcend and provoke.*
>
> *Dahl (2016)*

Needless to say, the provocation must not be read in negative terms. In this case it is meant and push the development of keyboard-based AMIs alongside the development of the existing paradigms: whether referring to innovative DMIs such a the Seaboard or more niche AMIs such as the TouchKeys. The point of the test was to look at keyboard-based DMIs and AMIs in a more holistic way. As discussed above in the introduction, one of the main goals was to analyse how performers would approach the proposed instruments at the first encounter and to what extent they would be able to improvise on two different pieces: a known jazz standard, and a modal melody. Having taken improvisation at a first encounter as the main focus of the comparative test, learning curve of the interfaces and precision in sound modulation were also considered and analysed throughout the semi-structured interviews.

For these reasons, it was important that every instrument chosen for the test was used as originally intended. The Reach system was limited to 49 keys to make the environment as similar as possible between interfaces, but was still mounted on an acoustic piano rather than a keyboard like TouchKeys.

In response to a question that sought to investigate why the Reach system was easier to use, compared to the other two interfaces, another user said:

> *Because the touch and everything is so familiar and feels good, the new elements are much easier to control because you don't have to think about that as well as the interface itself.*

This reply underlined how the familiar interface of the acoustic piano, helped the user focus on the novelty of the gesturally controlled interface without having to focus also on their pianistic technique.

6.4.2 Degree of Invasiveness, Learning Curve and Additional Technique

Closely tied to the feel of the instrument itself, another investigated aspect was how the invasiveness of the interfaces affected the ability to approach the instrument, learn the instrument and how much of the preexisting instrumental technique could be implemented. On multiple occasions, users commented on this aspect throughout the testing phase, and the same topics were then dealt more in depth in the interview. One user, after having described the TouchKeys as *really hard to control*, described their first approach with the Reach system saying *Oh loved it! So much easier!*. It is interesting how these comments could also be seen as a vocalisation of a state of flow. In fact, as noted by (Csikszentmihalyi, 2014b: 136): "the state of flow denotes the holistic sensation present when we act with total involvement. It is the kind of feeling after which one nostalgically says: *that was fun*, or *that was enjoyable*". The user in this specific case, found the interaction with the Reach system easier from a cognitive point of view. The new sonic elements provided by the Reach system were easier to control because the user felt that they did not have to think about the interface. The acoustic piano was an environment in which the user felt completely comfortable. In several cases the pianists seemed to find the Reach system more accessible because of its relation to an instrument they knew. It was as if they could just focus on exploring the system and its effects, forgetting about the actual pianistic technique that had been consolidated in many years of study. Regarding the additional technique usually required to play AMIs and DMIs, one user said:

> *My technique doesn't work there, I need to take one week with it, after one week I can go and play it. But I have one concert today, and if I need a piano, I don't take this and this* [pointing at the Seaboard and the Touchkeys]. *I prefer to play with my iPad, you have control there.*

The same comments were made by other users, who pointed out how, with some practice, they would have mastered the technique required to play the instruments. None of these comments were made regarding the Reach system. On the other hand, two users out of six spontaneously asked if they could employ extramusical technique on the Reach system to trigger it in various ways. One user later explained their experiments by saying:

> *(...) there were times when I was trying to do two gestures at once: I was trying to get the reverb to stay on, and put some pitch bend in there.*
> [he was covering one hand with the other]

This spontaneous research of additional technique to implement in the improvisation, is reflective of the level of accessibility, and ease of use of the Reach system. Similarly to what had previously happened in the case study with Rosie Tee (Chapter 5.4), the accessibility of the interface combined with the confidence of the user to approach the described interactions led to an exploration of extended techniques outside of the nuanced gestures proposed. These extramusical gestures—because they were sought by the participants themselves—did not provide a disruptive gestural element. The possibility of using these extramusical gestures provided by the instrument itself helped the user transition from the nuanced gestures embedded in the pianistic technique towards sound controlling gestures outside of the traditional movements of the performers.

Also verbally, the way in which the users approached the three systems was very diverse. All of the reactions of the users when trying the Reach system for the first time were positive and showed interest towards the sounds and processes behind it. Most of the immediate reactions of the same users when approaching the other two interfaces were referring to the difficulty to transfer their skills to that environment, or to the unconventional design of the interfaces themselves.

6.4.3 Control Precision

This first element of investigation of the test was closely tied to another question of the interview. Knowing the difference in precision between the three systems, with the Reach system being the least precise in terms of parameter modulation and control resolution, it was essential to investigate whether this difference in precision was an acceptable tradeoff considering the lower degree of invasiveness. Five users out of six stated that they would definitely prefer a less precise system that would give them more freedom of playing. A user specifically stated:

> *I would definitely go for the Reach system. I am all about the sound and feel. If we're speaking about piano sound, there are no keyboards that will hold up the comparison. If I play here* [indicates the piano] *I feel so much better. It's not only a question of control, but it's also a question of aesthetics and inspiration. How can you be inspired or play the piano on something like this* [indicates the Seaboard].

One user also stated:

> *I thought that precision matched the freedom on your system. I feel like it was easier to be precise on your system compared to the others.*

This quote could be linked to the test limitation discussed the next section (Section 6.4.4.1) with regard to the balance between wet and dry signal.

6.4.4 Test Limitations

Even though the test provided useful insights, it presented limitations that need to be addressed. The present section acknowledges the test limitations and aims to explain the reasoning behind certain structural choices.

6.4.4.1 Sound Balancing

The main limitation of the test concerned the balancing between the wet sound (or processed live piano sound) and the dry sound (or unprocessed live piano sound). When designing the test tasks and structure, a conscious decision was made to leave the digital piano sound to its default settings. This choice was made keeping in mind that the default state of the plugin for the piano sound would have been crafted and designed to work well with the Seaboard and any other MPE devices (ie. TouchKeys).

During the test the different sound balance when playing with the Reach system, with a prominent clean sound coming from the acoustic grand piano and a more nuanced effected sound coming through the stereo speaker combined with the natural ability to hear clearly both the wet and dry sound underlined prominently the difference in sound balance between the three systems.

In retrospect, to provide a more equal sonic experience to the users, an extra channel on the digital mixer for the clean piano sound without any modulation could have been created on the DAW providing the Seaboard and the TouchKeys with a similar sonic experience in terms of dry/wet balance.

6.5 Ideas for Future Development

A lot of feedback received from the users was directed not only to the investigated mean of interaction and invasiveness, but also to the most practical aspects of the instrument itself.

One user, towards the very end of the interview, asked about the actual setup of the system. After having explained the convoluted wirings, the user started making a point on how important it is for a touring musician to have an instrument that is accessible all round. As soon as the word *software* was mentioned, the user reacted by saying:

> *You still need a software? I am speaking as a musician. The cool thing about a contact mic is that it's immediate. You stick a microphone on the piano, plug it into a pedal, and you're set. So I'm wondering, how would you go and make it even more "compact"?*

The user then continued:

> *It would be really cool to have something really small without too many things. I hate computers in live settings. You are dependent to something that could crash every second.*

The elements that make an instrument accessible are not limited to the interaction with the instrument itself: accessibility is an all round feature. From the most practical elements, such as portability, size, setup time and reliability, to the most musical elements, such as interaction, and ability to port pre-existing instrumental technique.

This concept was later strengthened by every other participant in the test. As performing musicians, they all had personal aspects of the instrument that they found extremely important for it to be considered accessible. After, one user mentioned the possibility of having a portable version of the Reach system, comprised of free and open source software, cheap and replaceable parts and a mobile control interface, all of the users expressed their interest in owning one. Once the price of the individual pieces comprising Reach was broken down, one user said:

> *I'd be willing to pay way more than 150£ for something like that. I'd be willing to spend 3-400£ on it.*

This was emblematic for the purely practical side of the DMI, as most users valued the instrument more than the sum of the parts.

6.6 Chapter Summary

The results of the comparative testing showed a connection between the invasiveness of the digital system, confidence to approach the instrument and ability to freely improvise and transfer previously learned skills and technique. There have also been cases where this experience has led to a spontaneous exploration of extra-musical technique to control the system in different ways. Out of the investigated topics, the most surprising results were gathered when asking about the tradeoff between precision of modulation in relation to the invasiveness of the system. All of the jazz players that took part in the test seemed to prefer a less precise mean of modulation in favour of a less invasive interface. This was later linked to two main factors: the first one being that an overly precise mapping of the fingers disrupted the flow of the improvisation, and the second one that, not being able to predict precisely the outcome of the modulation, the pianists treated the system almost like another musician with which they could interact. Following the first user testing, it turned out that the users not only were keen to adapt and accept such a system building upon their existing technique; moreover, due to the proposed price range, they would be more willing to invest in an instrument like the Reach system rather than the other two tested and compared. In several instances this led to a discussion on the portability of the system that, at the time, was not regarded as compact and portable. From the comparison of the results from both the user testing sessions, it is clear that users are more likely to prefer less invasive AMIs or DMIs in favour of a less disruptive experience.

During the first round of user testing, when testing Reach alone, users noticed that the most invasive part of the system was its sonic aspect, making it difficult for users to predict the outcome of their gestures. Once it was compared with the other interfaces, the Reach system was seen as the least invasive, with an acceptable compromise between low degree of invasiveness and control precision.

Even though the first test showed that no particular musical genre was suited for the system, thanks to a deeper understanding of the user base, a more compositional and improvisatory approach has been identified as the one that would take advantage of the system the most. When compared to existing keyboard-based systems, the

Reach system has been identified as the easiest to approach and explore at a first encounter.

The UEQ data from both rounds of testing have shown how lowering the degree of invasiveness of the system by basing its means of interaction on pre-existing instrumental technique, is not only valid but inspiring and useful. Even when completely unexperienced, users were quickly able to take control of the system and control the sound with a minimal amount of guidance.

Thanks to the feedback gathered from both user tests, the latest version of the system was developed, trying to minimise the elements composing the system and porting the whole project to a more user friendly format.

Chapter 7

Conclusions

7.1 Summary of Dissertation

This dissertation focused on the application of ancillary gestural nuances in piano playing for the control of live sound modulation parameters. While the application of gestural control has been the subject of previous research projects, no such works have focused on the effect of these systems on the creative practice of a multiplicity of musicians coming from different musical backgrounds. The first part of the dissertation, after an introduction to the topic and motivation for the research (Chapter 1), focused on the concept of interface transparency for the development of musical instruments that provide an opportunity for embodiment, reviewed existing keyboard-based interfaces that enable gestural control over sound modulation parameters, and gave an overview of the sensing modalities for hand-based gestures. A contextualisation of the methods employed throughout the research in the relevant literature was also provided as a closing section of Chapter 2. The second part of the dissertation presented the seven empirical studies carried out as part of this research project. An initial case study culminating in a public performance with the first prototype of Reach, the system developed throughout this research project, was presented alongside an analysis of the profiles of creative practitioners targeted for the case studies (Chapter 3). Next, the first user testing conducted with classical and jazz pianists was presented (Chapter 4). A combination of task-based evaluation, UEQ, qualitative methods involving performance and open interviews were used to collect data in the process. Following the analysis of the target user base presented in Chapter 3 and the feedback gathered in the first user testing, an analysis and comparison of the four case studies was presented (Chapter 5). To conclude the second part of the dissertation, the comparative user testing was presented (Chapter 6). This test aimed to compare two commercially available keyboard interfaces—Seaboard and Touchkeys—with the Reach system. A different approach was taken compared to the first user testing: a combination of free improvisation, UEQ and semi-structured interviews were used to collect data during the process.

7.2 Thesis Contribution

While there have been research projects exploring gestural musical instruments based on existing instrumental technique, there seems to be a lack of focus currently on the topic both in the academic and the commercial world. Including the musician in the design process and analysing what was being performed helped examine what causes disruption in the musical flow when approaching new gestural musical interfaces. This research aimed to remove this disruption by utilising spare cognitive bandwidth: that which is intuitive through ancillary gestures that are redundant when it comes to sound production. By using these gestures, the creative control of the pianist over sound can be extended by utilising gestures that do not require *non-musical* thought.

The aim of this dissertation was to explore further the concept of instruments that exploit pianistic ancillary gestures as a means of controlling sound modulation parameters, focusing on the first encounter with the instrument, examining the initial learning process, and the potential extension of creative outputs that these musical instruments could provide. The four most important contributions of this work can be found in the results of the investigation: a review of the concepts supporting the development of new instruments that take advantage of pre-existing instrumental technique; the qualitative and quantitative data gathered through two formal user tests; a series of documented case studies culminating in public performances; and the creation of an augmented instrument that enables control over sound modulation through pianistic gestural nuances. The review of the philosophical concepts underpinning this dissertation brings attention to the concept of instrument transparency and uncovers the relationship between instrumental qualities and the history of traditional instruments. This relationship is explored and the concept of embodiment of a musical instrument is seen as a result of an interplay of factors.

Having explored embodiment and transparent relationship, the user tests provide empirical data around these phenomena exploring both prototype instruments and commercially available ones. Interviews performed with pianists from different musical backgrounds suggested that while the gestural interaction is not the only factor responsible for the transparency of the interface, being able to focus on the novelty of the sound rather than learning a new gestural vocabulary helped them approach the new instruments with more confidence. In a comparative situation, the element of the acoustic instrument—the piano—at the base of the Reach system, made the prototype more approachable providing a constraint-free environment. Alongside the technical evaluation conducted in the user tests, the more qualitative approach of the case studies provided a better understanding of the real world use of such instruments. Using a combination of participatory design and public performances these case studies uncovered the concept of *ideal instrument* for creative practitioners, exposing wishes and constraints in the development phases and providing useful real world scenarios in the performances. The ability of adapting the instrument according to the advances of the research was key in providing the tools to iteratively test the ideas in seven different empirical studies. The Reach system has ultimately been essential to explore the initial hypothesis but also to uncover new questions surrounding the topic that will lead to future explorations. In addition—being Reach open source and freely available online—it is hoped that the system would be useful not only to the research community but also to the community of musicians and pianists that seek an instrument to aid their exploration of gesturally controlled live piano sound modulation.

7.3 Future Work

7.3.1 Longitudinal Case Studies and Creative Explorations

In this dissertation the observation of the relationship between pianists and gesturally controlled systems, has been purposefully limited to a maximum of three gestures paired with three sound effects. This was done in order to keep the analysis focused and avoid falling into overstimulating experiences. The goal was to provide musicians with an instrument that presented novel characteristics paired with a new sonic environment that would have however a limited set of new elements: enough to keep the interest high, but low enough to provide a comfortable level of challenge for the musicians. This paradigm has been more flexible in the three case studies explored in Chapter 5, compared to the user testings presented in Chapter 4 and 6. In the case studies, thanks to the collaborative design workshops and a closer and longer collaboration with the artists, different gestural nuances paired with new sound effects were explored and implemented. The promising outcomes of all the four case studies, highlighted the importance of conducting longitudinal studies to explore more in depth the creative vision of creative practitioners.

Longitudinal case studies with pianists would provide a different point of view on the development process of the musical instrument regarding both the gestures used and sound effects. Providing creative practitioners with more time with the instrument will enable the analysis to shift from the first encounter with the musical instrument to the results of a longer and deeper exploration. Another novel aspect of this extended case study could be the effect on the materials produced. While the case studies carried out as part of this research were driven by a final composition and performance, an extended study could lead to different outcomes such as a series of pieces and the creation of a live performance based entirely on the novel interaction. Overall a longitudinal exploration would be interesting to provide one more point of view on the topic. As a first study, a collaboration with Dutch composer Raoul Van Herpen[1] has commenced exploring the potential use of nuanced gestural control in the artist's own work.

7.3.2 Experienced Technology Users and Electronic Music Repertoire

As previously explored in Chapter 3, pianists that are comfortable using gestural controllers in performance and generally more comfortable with technology were left out of this initial investigation. For this reason, further investigations could focus on analysing how musicians that are already combining instrumental practice and technology could take advantage of a system like Reach. Results from a study like that could provide useful data to compare to the original findings provided in this dissertation. If musicians that are not in need of this aid used such a system that was built to provide a seamless transition between instrumental technique and gestural control, it would provide insight that could help broaden the scope of this research reaching out to a more stylistically diverse array of musicians.

[1] https://raoulvanherpen.bandcamp.com

In the user testing in Chapter 4, testing the Reach system on classical repertoire caused issues with the users that were not accustomed to playing and hearing the piece with sound modulation applied to it. This hindered their experience, leading them to struggle during their exploration of the system as well as the mode of gestural interaction. Working with musicians accustomed to technology in live performance would have the benefit of being able to explore new repertoires. These pieces, written for piano and live electronics, would enable the musicians to play music from a repertoire they are comfortable with. Moreover, due to the fact that this repertoire contains modulated sound by default means that introducing the Reach system to it would not cause the kind of disruption experienced by musicians in the user testing in Chapter 4.

7.3.3 Application of Gestural Principles with other Instrumentalists

This dissertation has focused solely on pianists and the pianistic technique taking advantage of ancillary gestures for the control of sound modulation parameters. However pianists are not the only instrumentalists with a rich set of gestural nuances as part of their technique. As researched by Lähdeoja et al. (2009) and Brown et al. (2018), many other instrumentalists exhibit expressive gestural nuances as part of their instrumental technique. The knowledge acquired around the use of gestural nuances for the control of sound modulation parameters could be easily adapted for the analysis of a different set of gestural nuances on a different instrument providing fertile ground for further exploration.

Stemming from the work of Brown et al. (2018), the analysis of expressive hand gestures during a vocal performance has been one of the explored avenues. This project, following the third case study (Chapter 5.4), was conceived with Rosie Tee as a result of the performance with the Reach system. The artist noticed how in her own singing practice, expressive hand gestures would surface spontaneously during live performances as a means of expression and embodiment of the song and lyrics. For this reason, having had previous experience with the Reach system, she expressed the wish to try and develop a system using the core aspects of Reach to sonify her gestures during the vocal performance.

7.3.4 Hand-Gesture Sonic Performance System

Apart from gestures pertaining to the instrumental technique, several other kinds of gestures could be studied and implemented for sonic performances. Implementing symbolic hand and arm gestures in a novel instrument for musical expression, would be useful in learning how musicians would interact with a system that enables them to palpably sculpt sound and directly manipulate audio synthesis parameters with fine grain control. The knowledge gathered through this research project would allow the creation of a musical environment that utilises common everyday gestures (i.e. hand grab, hand release, simulation of object movement in space, etc.) and translate them to sound modulation parameters. Gestural and positional data could be used to create embodied interaction with immersive interactive sonic environments, changing the way a musician approaches a performance with technology.

The approach described will potentially lead to a change in the way musicians work with interactive-audio elements in their work, encouraging them to move away from workflows that respond to a fixed score and towards ones that facilitate greater creative freedom, similarly to the ones explored in this dissertation. While utilising the knowledge and experience gathered in the past three years, the application of gestural elements to a completely different sonic environment would broaden the scope and impact of this research.

7.3.5 Development and Consolidation of the Reach System

The Reach system v1.0 presented as part of this dissertation (Chapter 6.3) was developed to be used as a probing tool for this research project. For this reason, the system has been mostly limited to recognising a fixed number of gestures paired to specific sound effect used in the different studies. There are a number of improvements to be made to make the system more accessible and easy to use for musicians. Moreover multiple gesture-sound pairs are planned to be implemented to make the instrument flexible and usable in a real world scenario.

In the current state, users need to follow a two-step guide in order to set up the system on their computer. The two components that currently compose the Reach system, the Reach.app for Leap Motion data parsing and the Reach DSP patches responsible for the sound modulation, were initially kept separate to provide the user easy access to the sound modulation part of the software. This was done with the goal in mind to allow users to edit both the mapping between gestures and sound modulation parameters and the audio effects themselves. However, the current setup still requires the users to have knowledge of the Pure Data patching environment in order to fully customise the behaviour of the Reach system. Without any knowledge of Pure Data, the users would be limited to the predetermined selection of gesture-sound mappings provided by the patch itself.

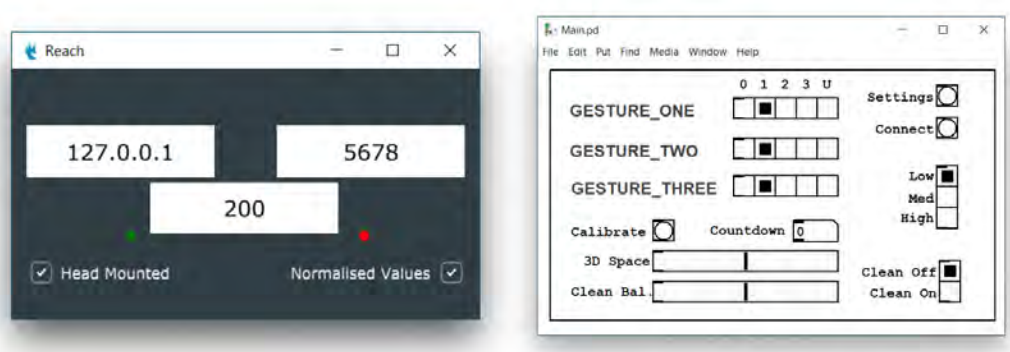

Figure 7.1: Reach v1.0 Setup: Reach.app and Reach DSP Pure Data main patch.

To overcome this usability issue, the plan is to develop the Reach system as a single app providing both Leap Motion data parsing and sound modulation capabilities within a single application. The aim is to embed the current Pure Data environment in the JUCE app through libpd—a library based on Pure Data that makes the environment embeddable in other applications—allowing end users to interact with the system through an accessible user interface. This would be an overall improvement providing smoother installation and set up process and easier interaction with the software whilst maintaining the customisation features. Thanks to the direct binding of the sound modulation patches inbuilt in this single application, there will be the potential to provide also complete access to the Leap Motion data to the unexperienced users, giving full control over the mappings not only to the expert Pure Data developers but also to the pianists approaching Reach from a musical standpoint.

Appendix A

Appendix Title

A.1 User Testing - Pre-Questionnaire

1. How long have you been playing the piano?
2. What is your musical background?
3. Is music currently your main occupation?
4. Do you know anything about electronic music or interactive music technology?
    (a) If yes to question 4,
        i. Do you have any experience with keyboard based electronic instruments?
            A. If yes to question (i),
                I. What instruments?
                II. In what occasion?
        ii. Have you ever had any experience in performing with technology?
            A. If yes to question (ii),
                I. In what occasion?
                II. Was it a solo performance where you were controlling the elctronics, a performance for instrument and tape, or was somebody live managing the electronics part?
5. Do you feel comfortable to do a solo improvisation?

A.2 Comparative User Testing - Pre-Questionnaire

1. Musical and performative background:
    (a) How long have you been playing?
    (b) Are you an active performer or composer that uses keyboard in his or her practice?
2. Relationship with electronics:
    (a) How would you describe your relationship with electronic music, and digital instruments?
        i. Expand further according to the answer given:
            A. How would you describe your knowledge of sound synthesis?
        ii. Do you have any experience with live electronics?
            A. If yes to question (ii)
                I. In what capacity?
                II. Can you describe your role in regards to the electronics?
3. Experience with the systems:
    (a) Have you ever heard of any of the systems you will be playing today?
    (b) Have you ever played any of the systems you will be playing today?
        i. If yes to question (b)
            A. In what context?
            B. For how long?
4. Experience with the pieces:
    (a) Have you ever played of C. Mingus' Goodbye Pork Pie Hat?
        i. If yes to question (a)
            A. Would you say you're familiar with the tune?
    (b) Have you ever played of G. Ligeti's Musica Ricercata No.7?
        i. If yes to question (b)
            A. Would you say you're familiar with the tune?

A.3 Comparative User Testing - Semi-Structured Interview

1. Generic question about the experience. To be expanded and explored depending on the user's reaction.

2. Now that you have tried the systems, would you say that the different levels of invasiveness affected your ability to play freely? Could you rank them and also try to explain why?

3. Regarding the learning curve instead, did you find one system harder to learn than another? Could you rank them and also try to explain why?

4. Did you notice a difference regarding precision of movements in relation to sound modulation? If so, could you rank them and also try to explain why?

    (a) If in question (4) user touches upon the preciseness of the three systems:

    i. Do you think a low level of precision is an acceptable trade-off when the degree of invasiveness and the learning curve are low, or is precision in modulation a priority when looking at keyboard based DMIs?

5. As a final thought, always taking into consideration the previous statements, would the difference in price between the three systems affect your choice if you were to buy one of these?

A.4 Notation Exploration

A.5 *//aria alla francese*

instrumentation: piano
+ reach system
developed by niccolò granieri

system settings:
sideways hand movement = pitch bend
wrist height = amplitude
hand above keyboard = reverb

composer: richard stenton

key:
"+" = opportunity for ornamentation


=80
5
10
15
20


25
30
35
Ped.
40
45
Ped.


50
Ped.
Ped.
55
60
Ped.
65
Ped.
Ped.


70
8va
8va
75
Ped.
80
85

blank

2 0 1 8

A.6 12.jpg

12.jpg

4 hands 1 piano + reach + .wav

richard stenton

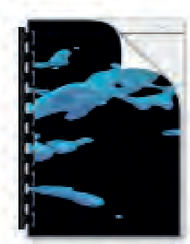

Score L is for player at keyboard.
Score R is for player at inside piano.

Both scores are read downwards with a duration of 10 seconds between blue lines (40 secs per page). Every page should be played by everyone (unlike in the workshop), consider them as part books. There will be an audio file which will mostly function to indicate 10 second intervals and page turns. Page turns happen at same time with both scores.

The octave of the pitches is optional but should be the same for the entire page. It is preferable to adhere to the order of the notes dictated with gaps roughly according the layout in the score. For example, if your two notes are "A" and "B" and there is a big gap between them you should play them at least one octave apart.

Keyboard (L)
In order to interpret the subtitles in the score we will use a sensor that tracks the movement of your hands. You will have the ability to manipulate the sound through sideways movement of the hand to bend the pitch and up and down movement of the wrist to fade in and out an additional sample (subject to change upon testing). Your movement will be tracked by a small sensor attached to the stand and will not impede on your movement in any way.

"Reach" is the PhD project of my college Niccolo Granieri who will be present for the whole day of rehearsals and the concert to run the technology.

Inside piano (R)
The inside player is free to express the score how they wish using multiple objects to make the sounds. The pitches should be adhered to as much as possible but neighbouring pitches could be used to express the shape. There are only ever two pitches at once for obvious reasons. The sound will be affected occasionally when playing at the same time as the keyboard player.

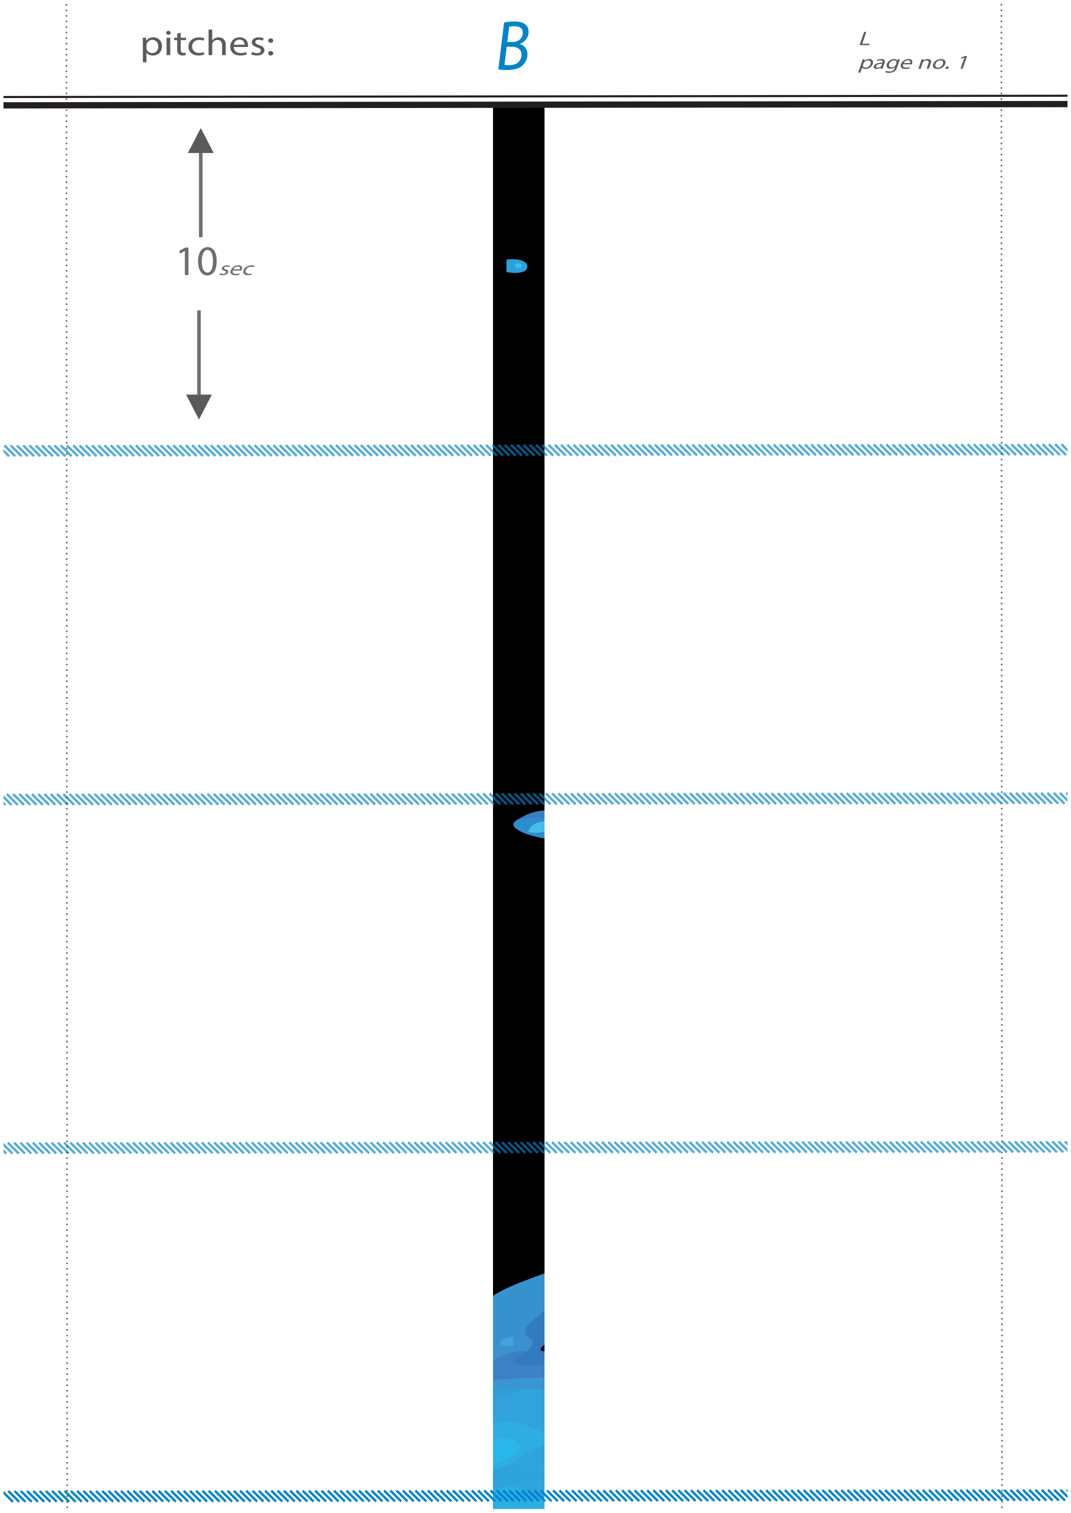
pitches:
B
L
page no. 1
10sec

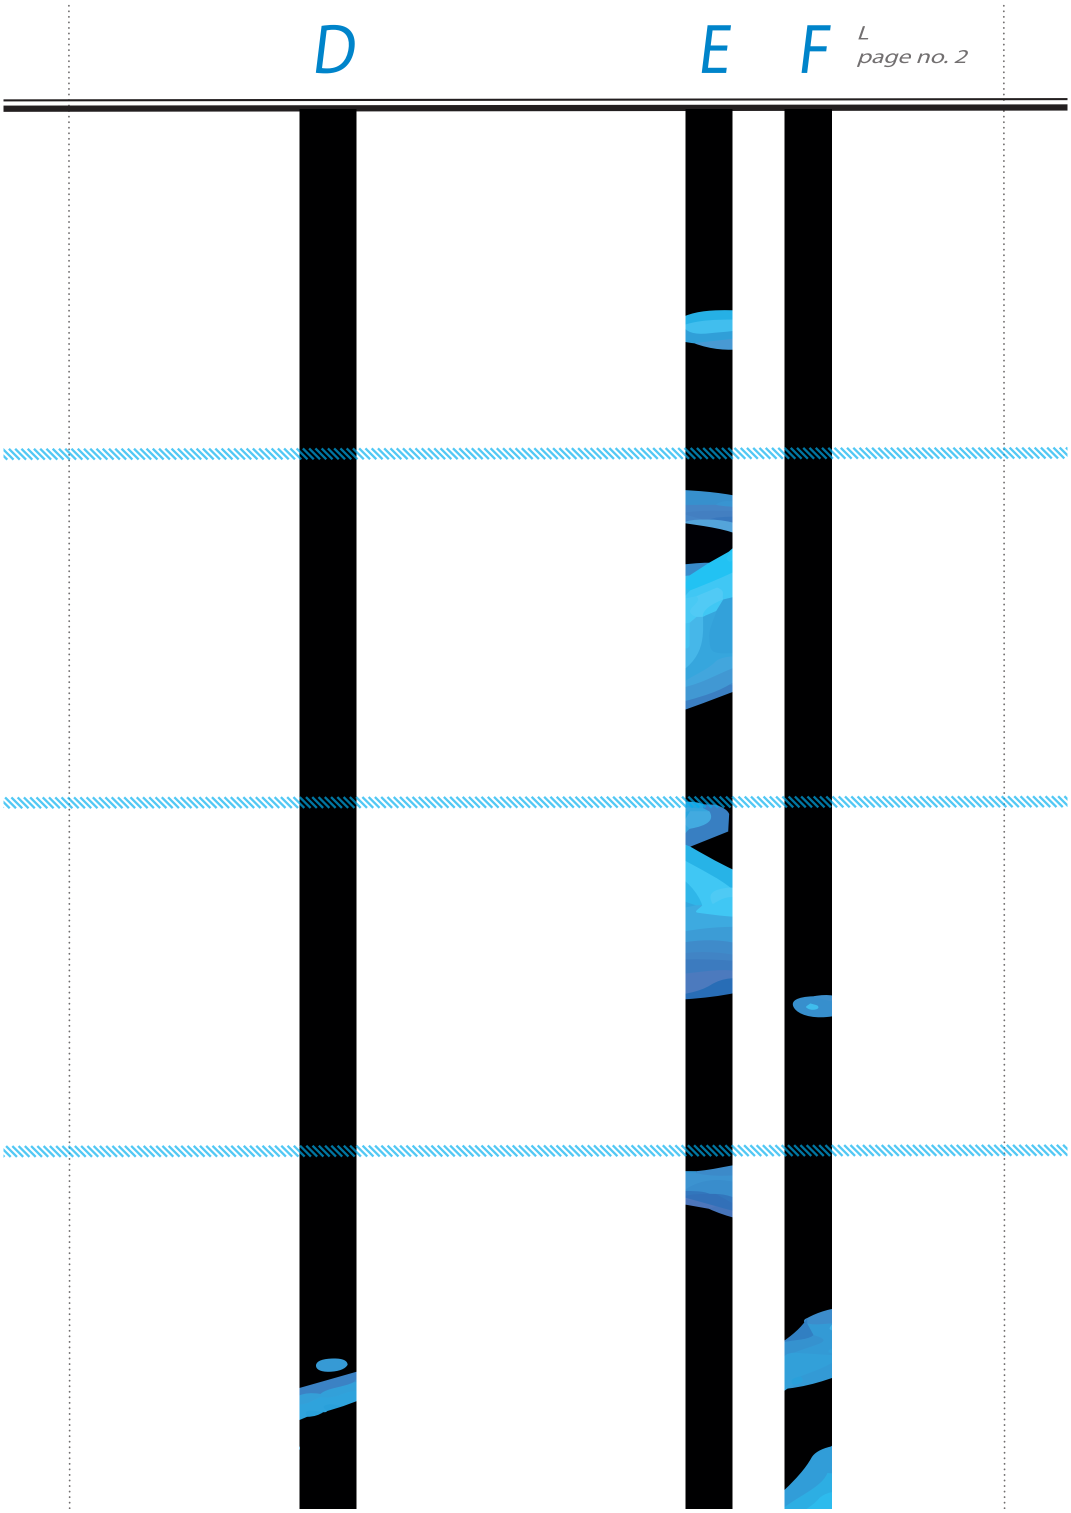
D
E
F
L
page no. 2

G E

L
page no. 3

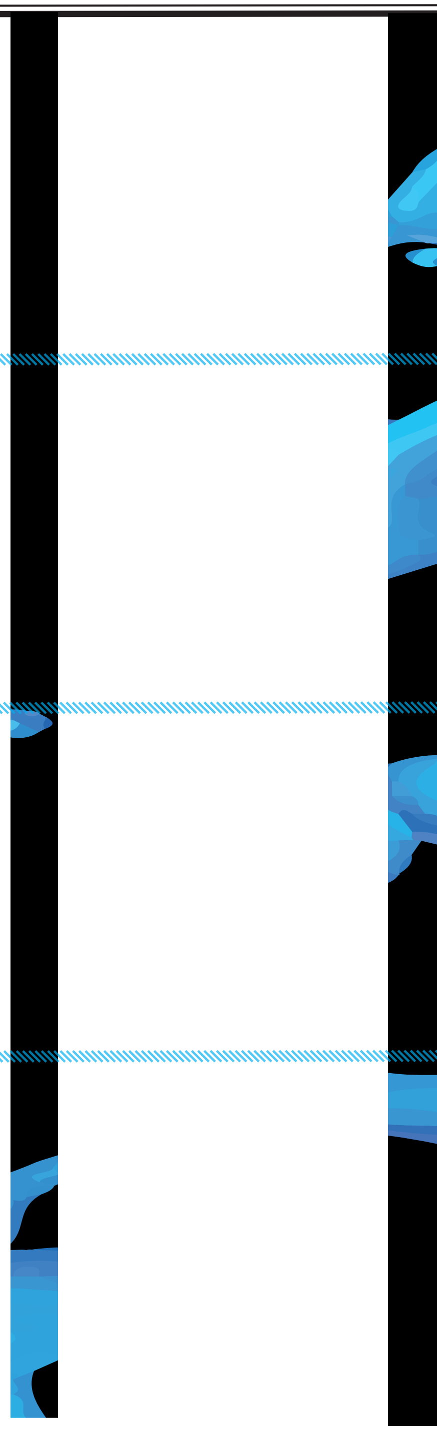

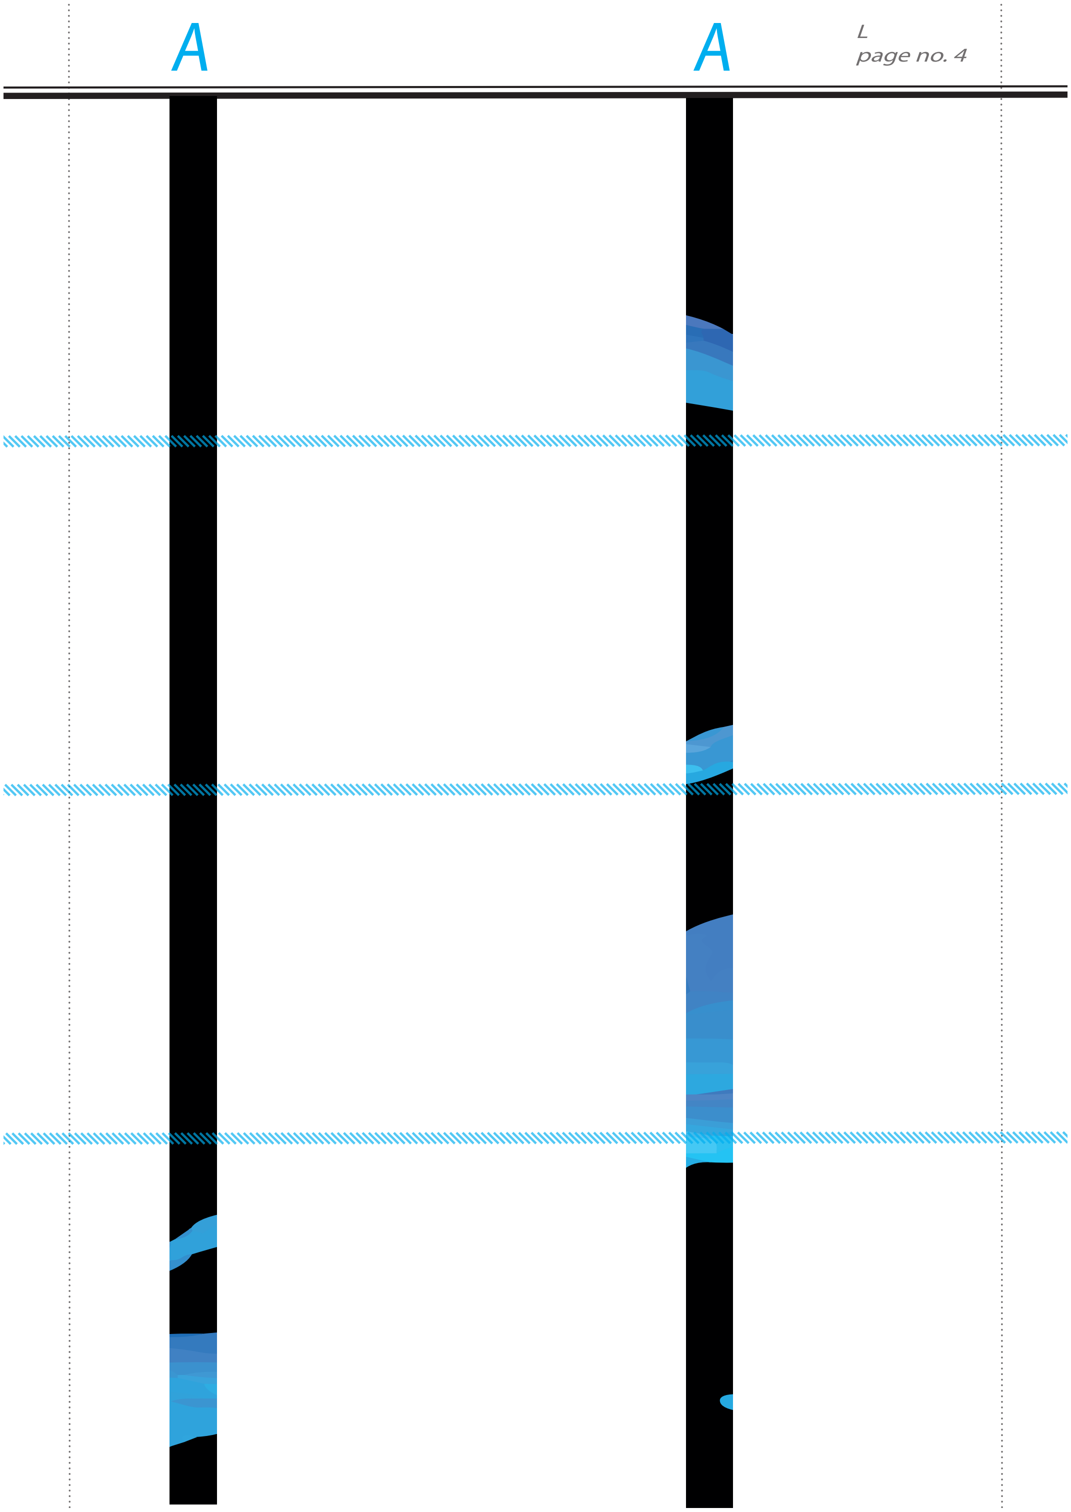
A
A
L
page no. 4

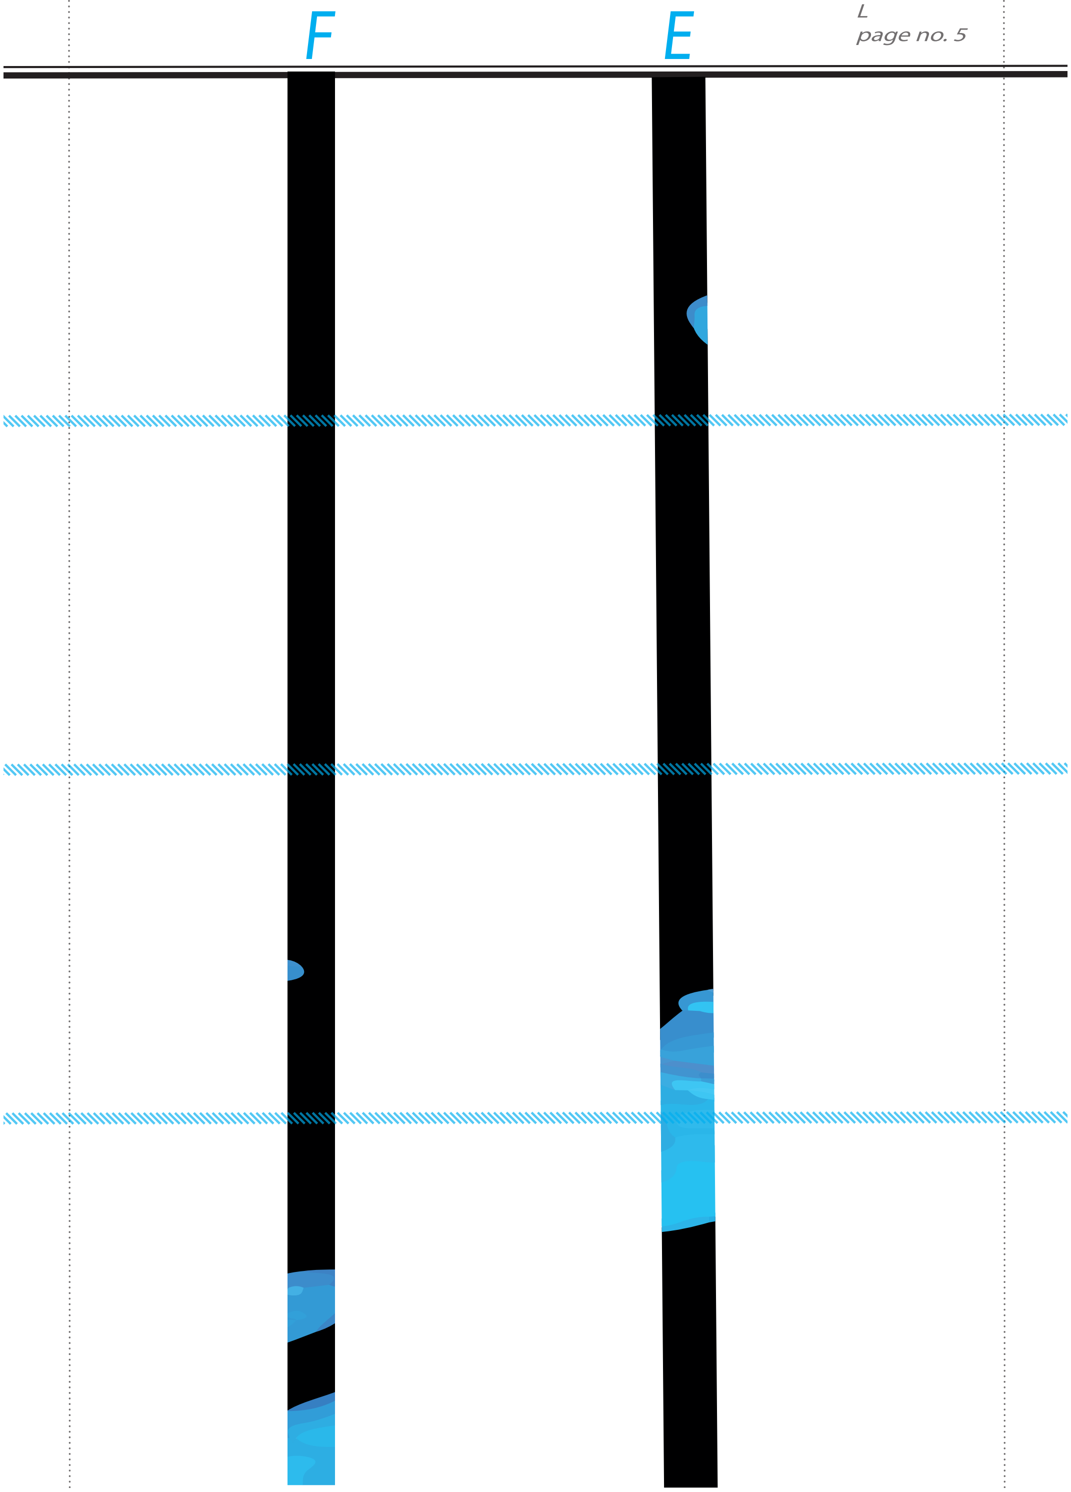
F
E
L
page no. 5

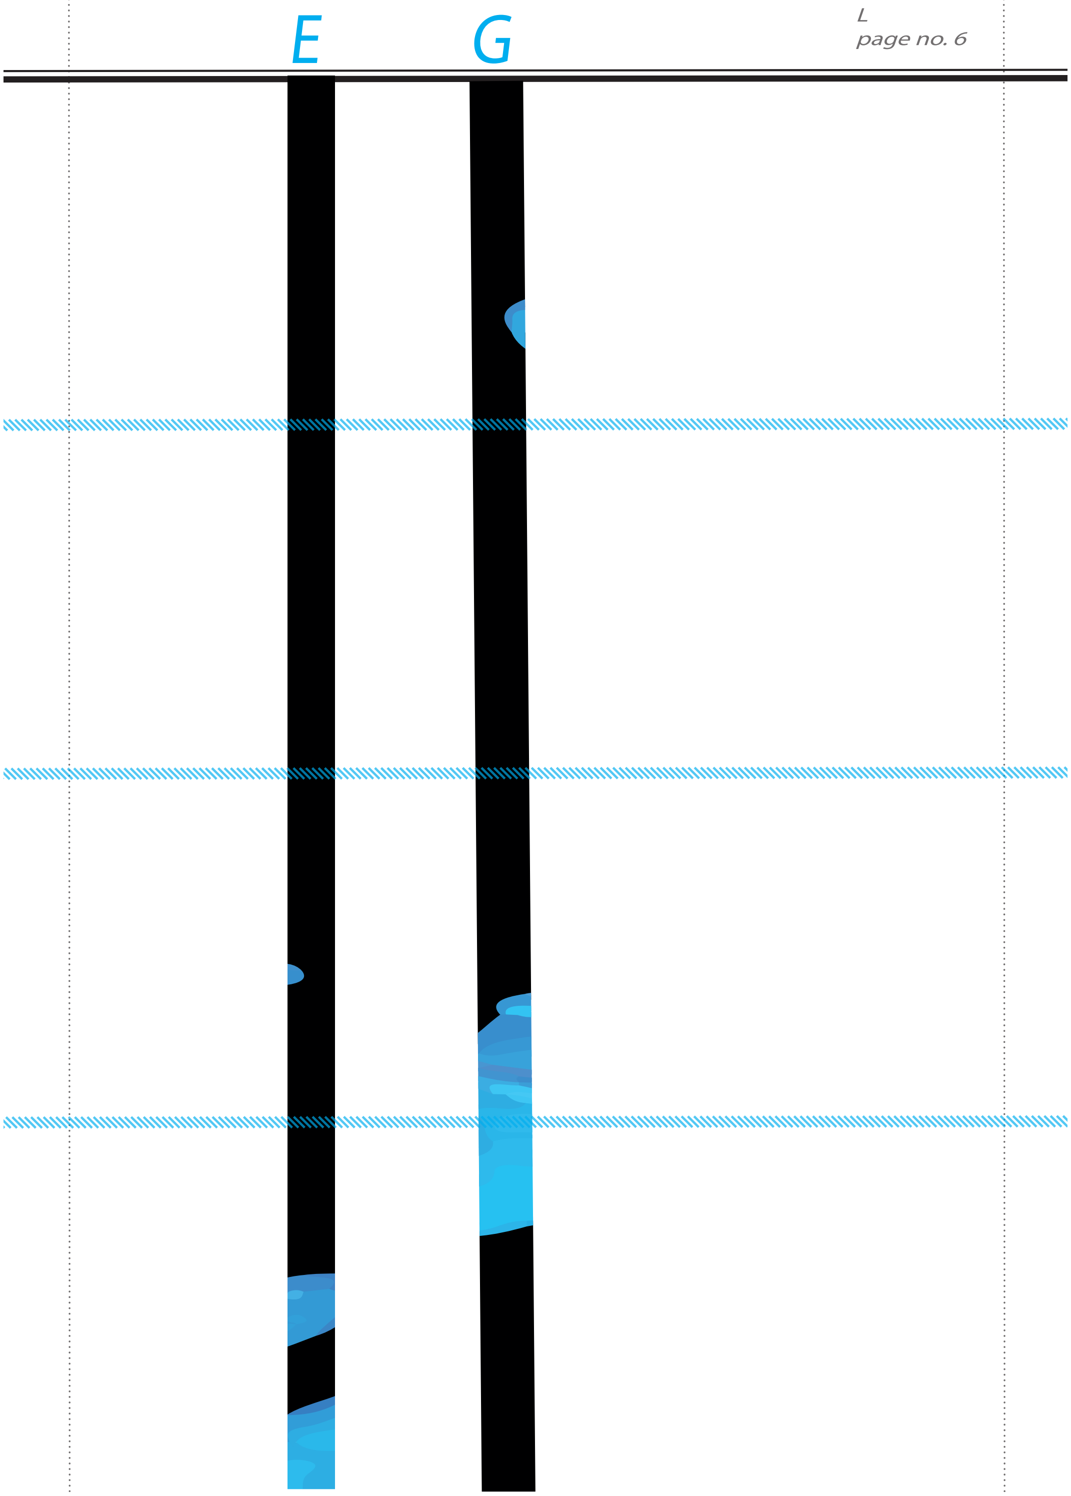
E
G
L
page no. 6

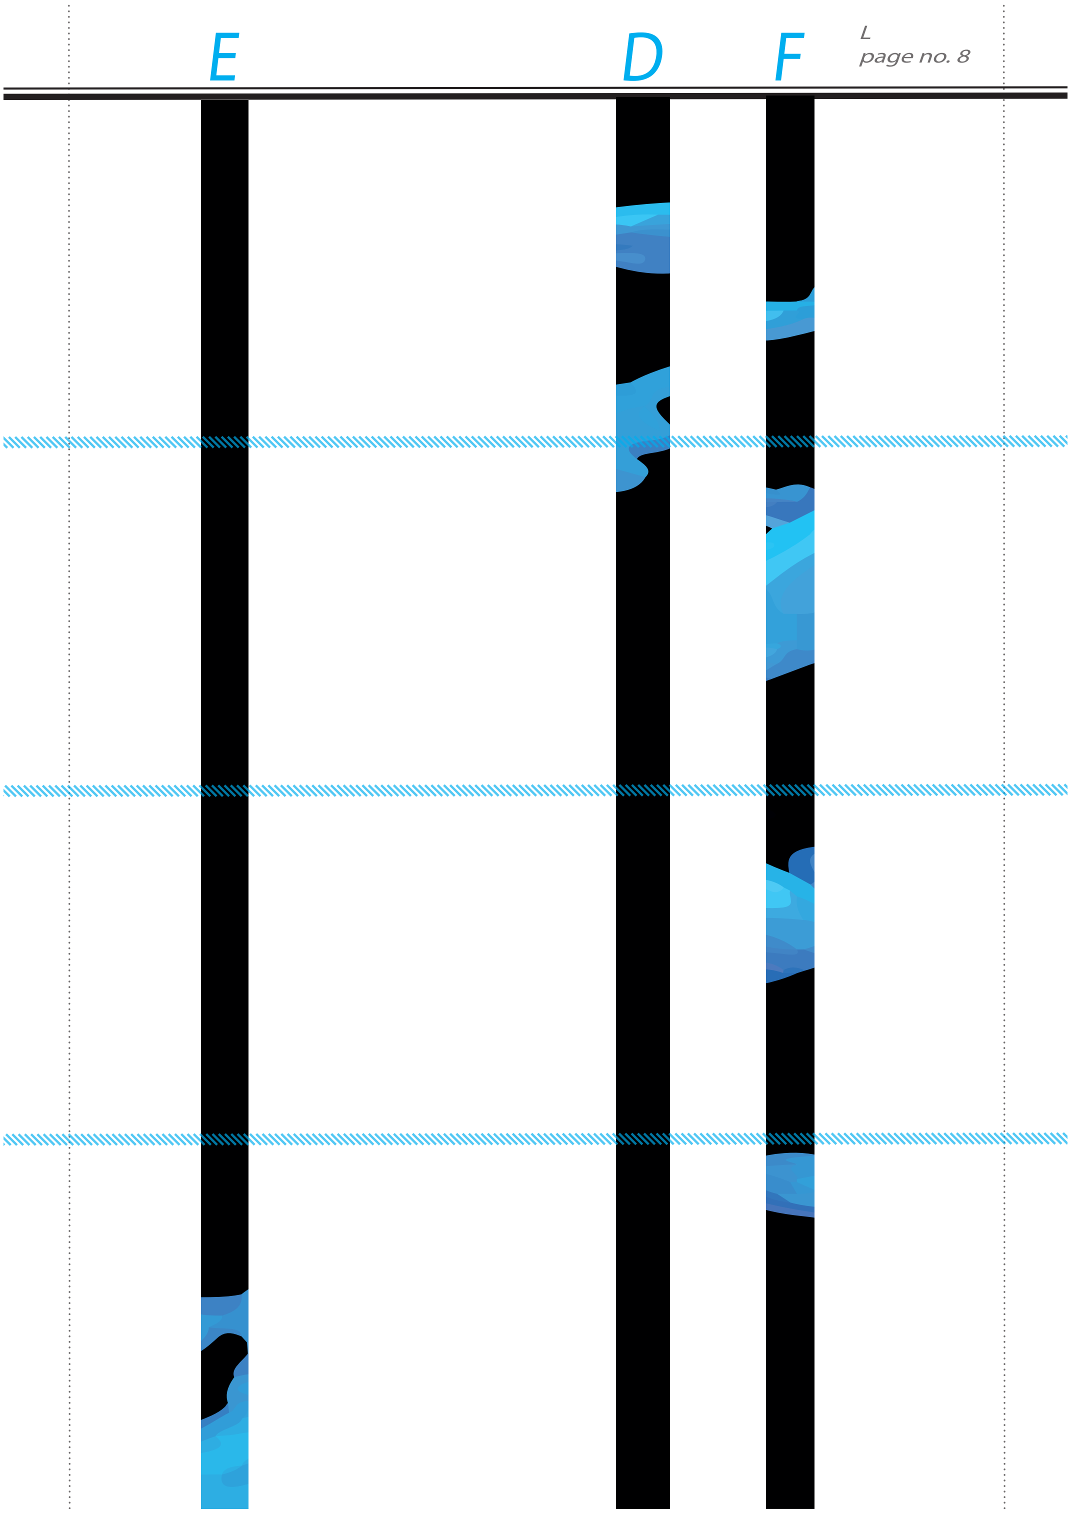
E
D
F
L
page no. 8

F

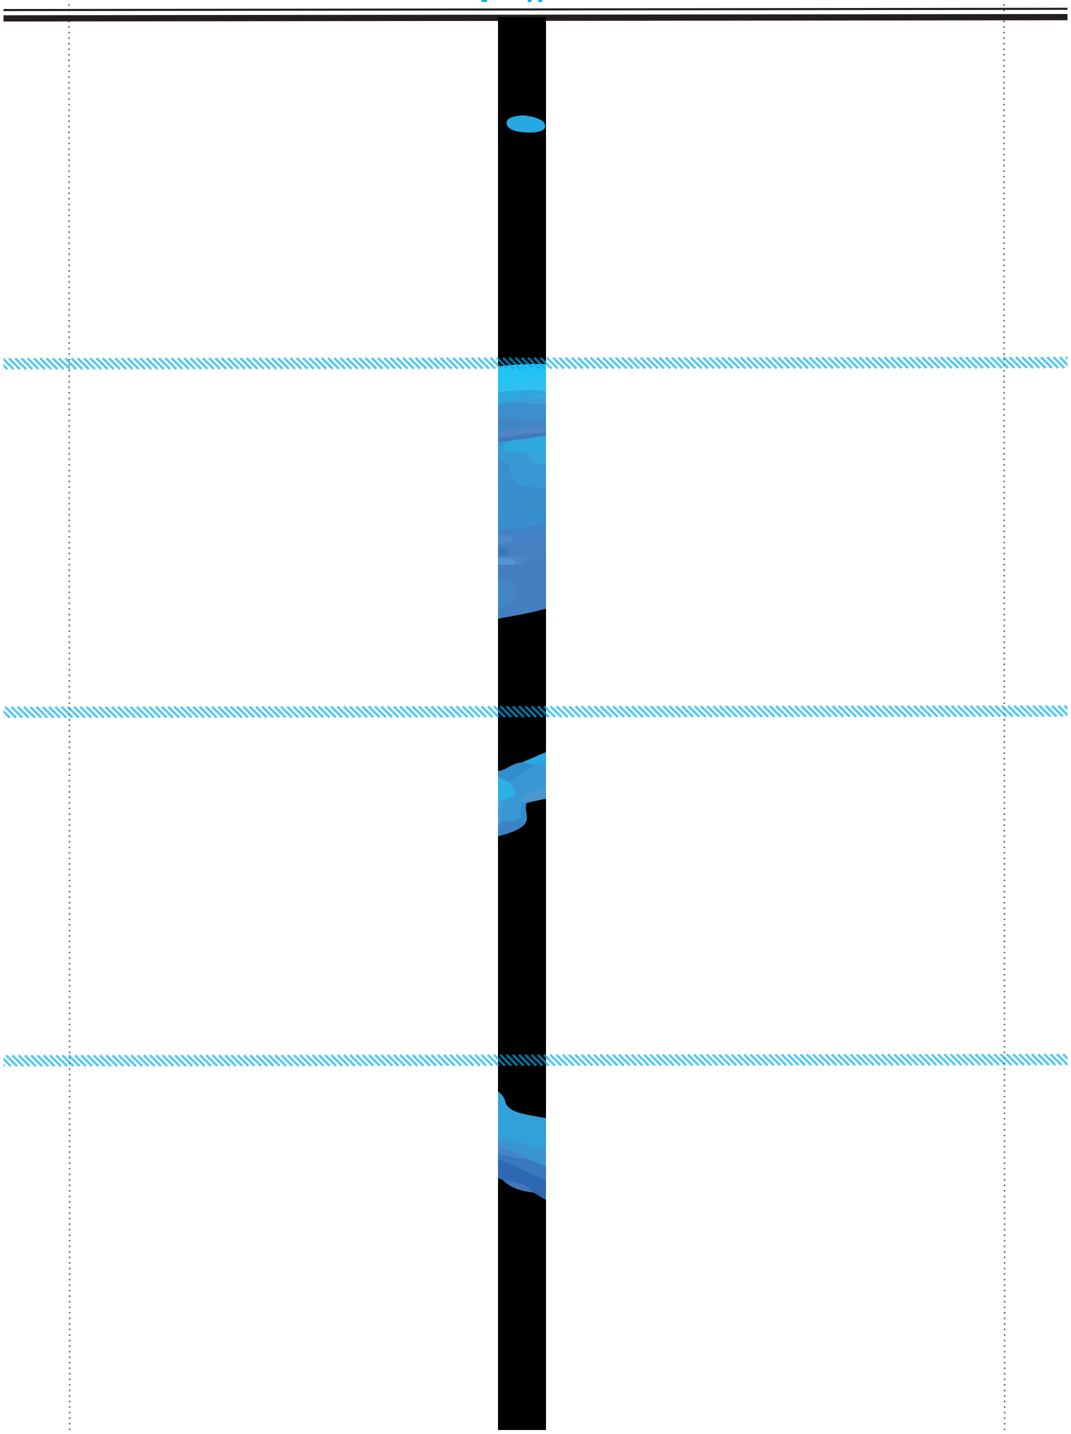

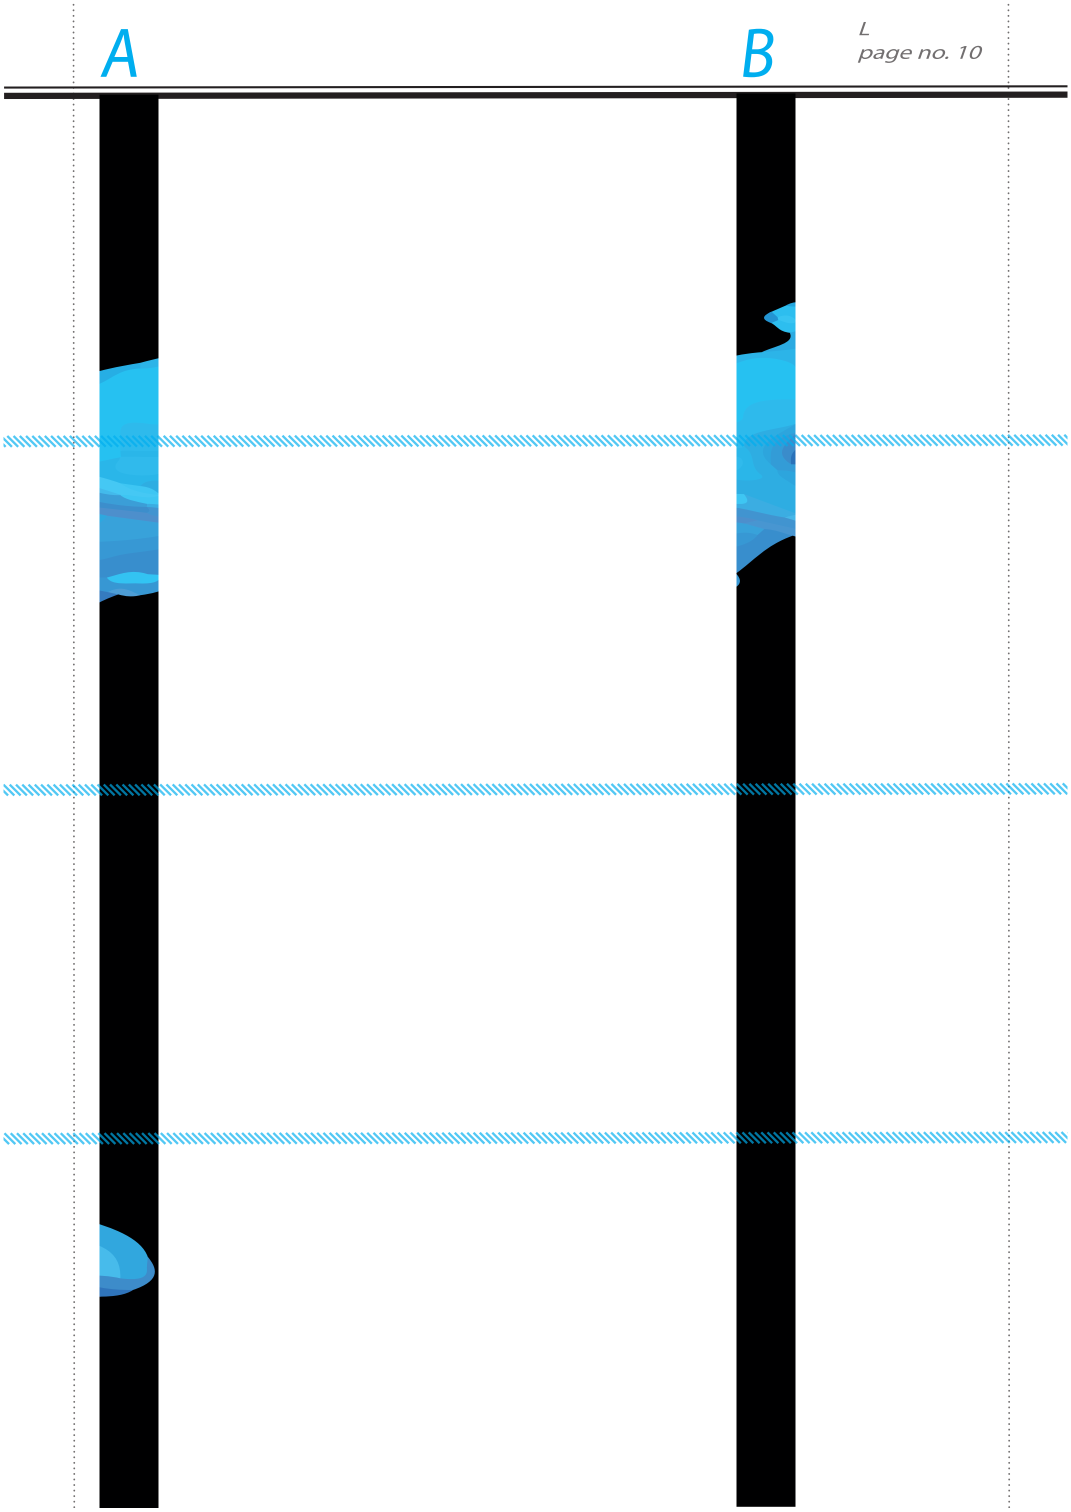
A
B
L
page no. 10

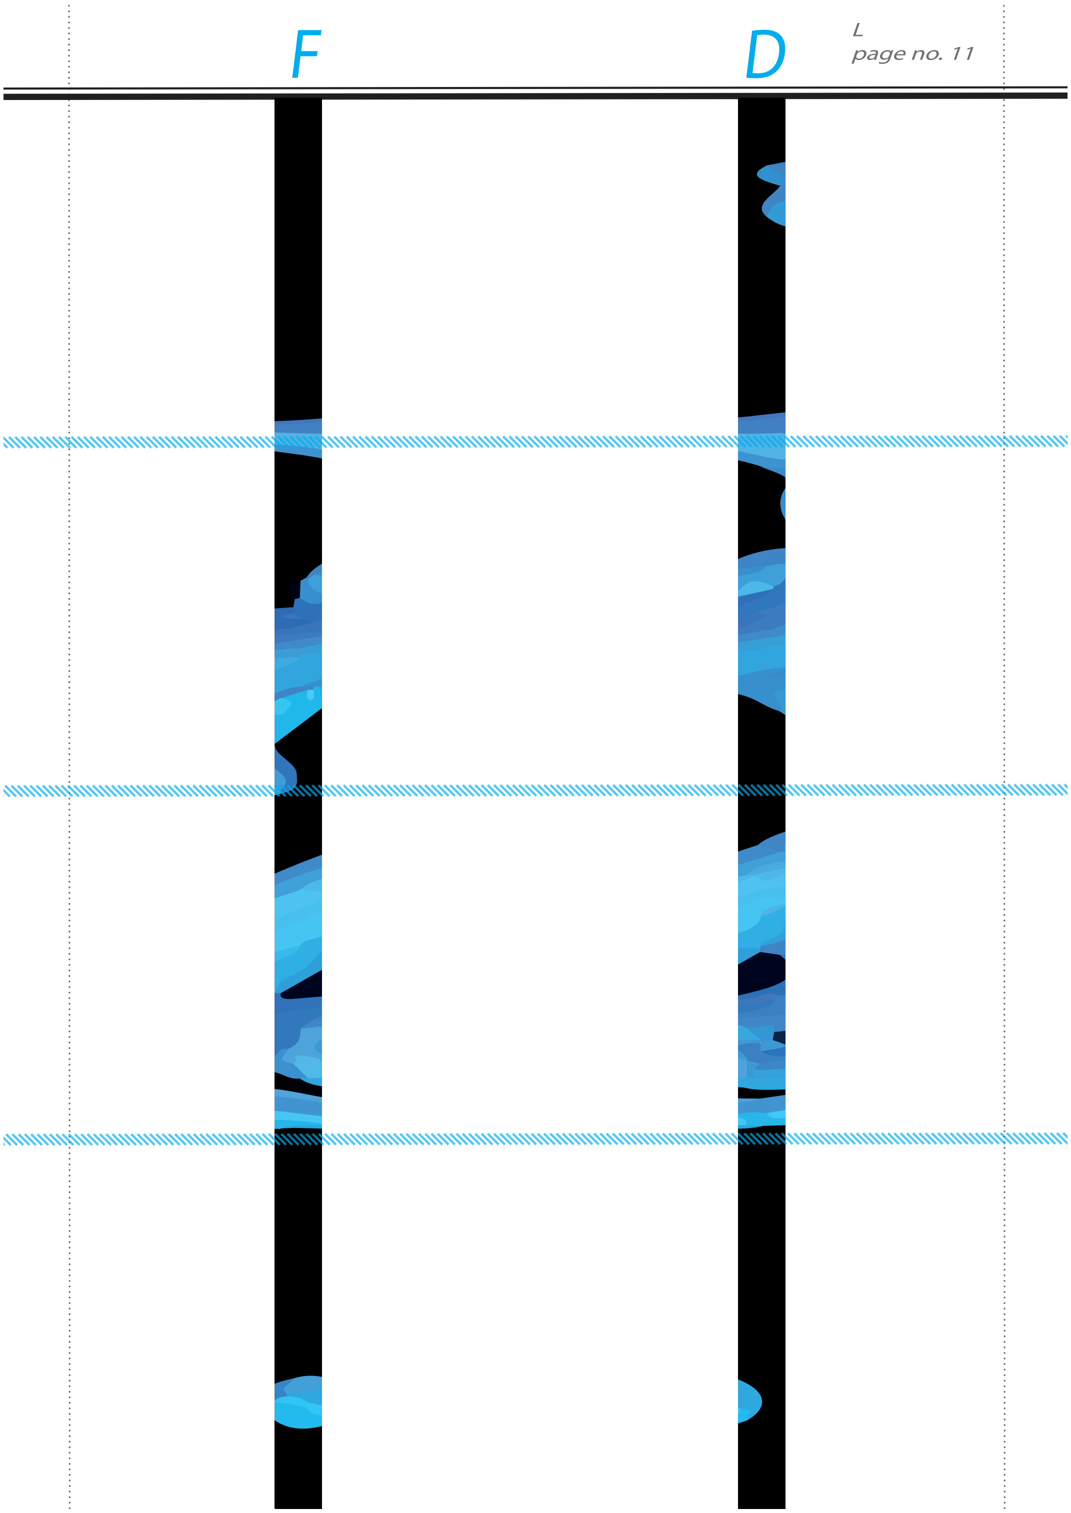
F
D
L
page no. 11

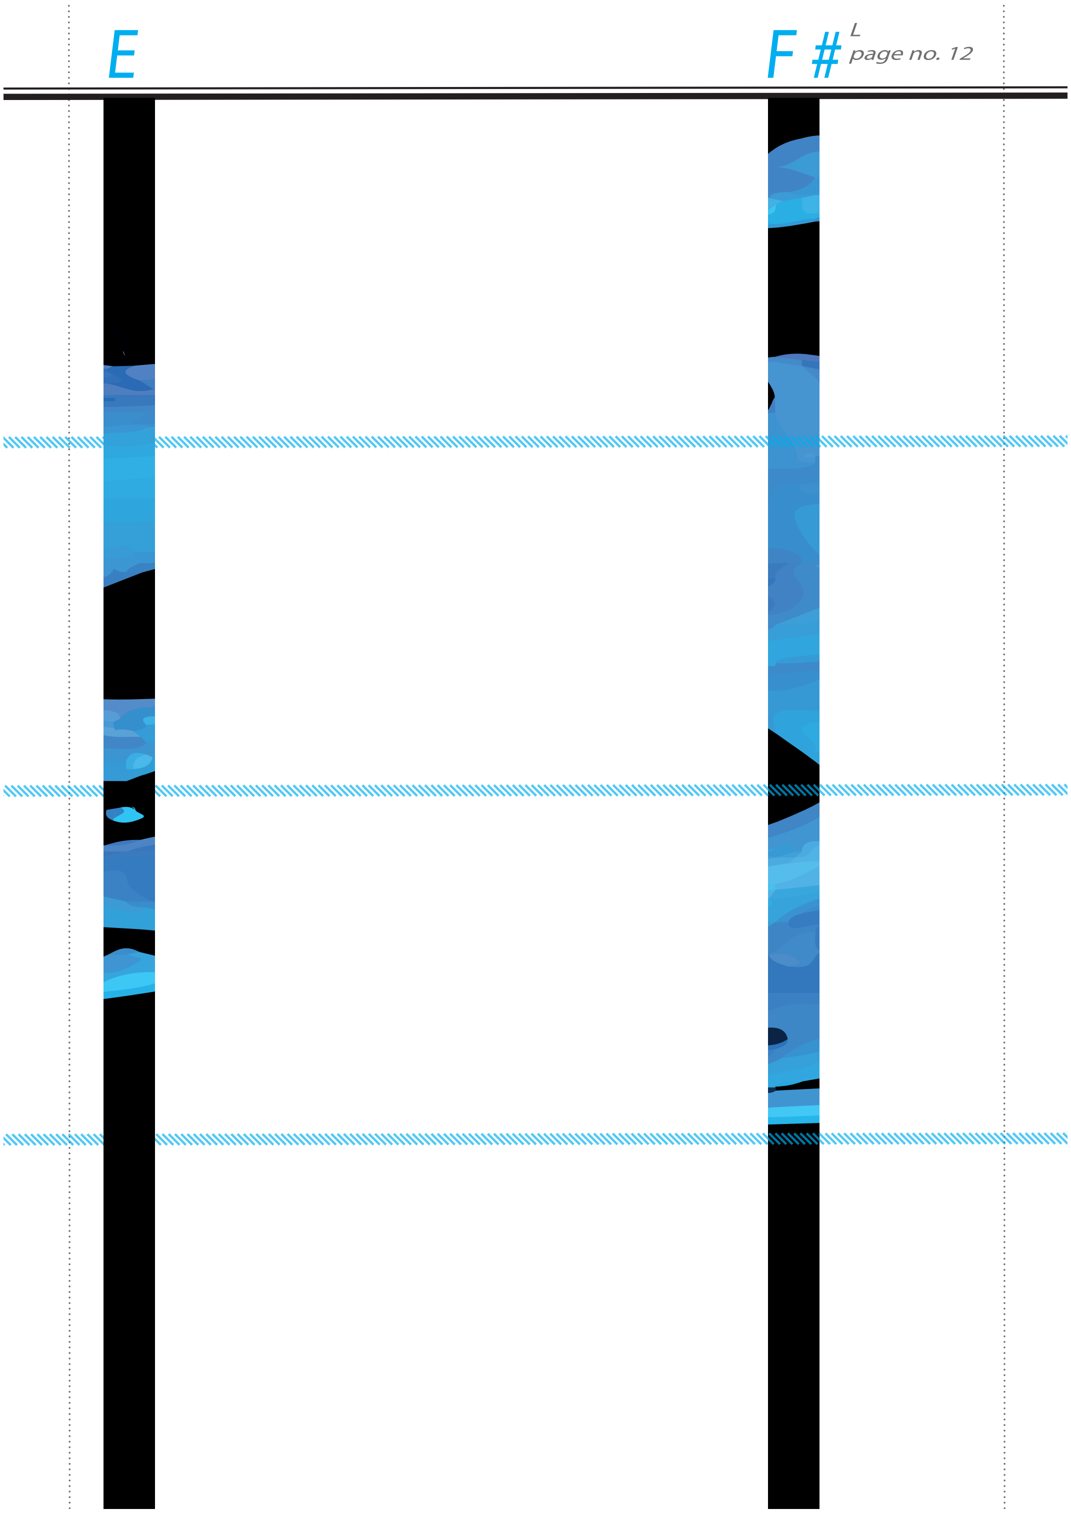
E
F #
L
page no. 12

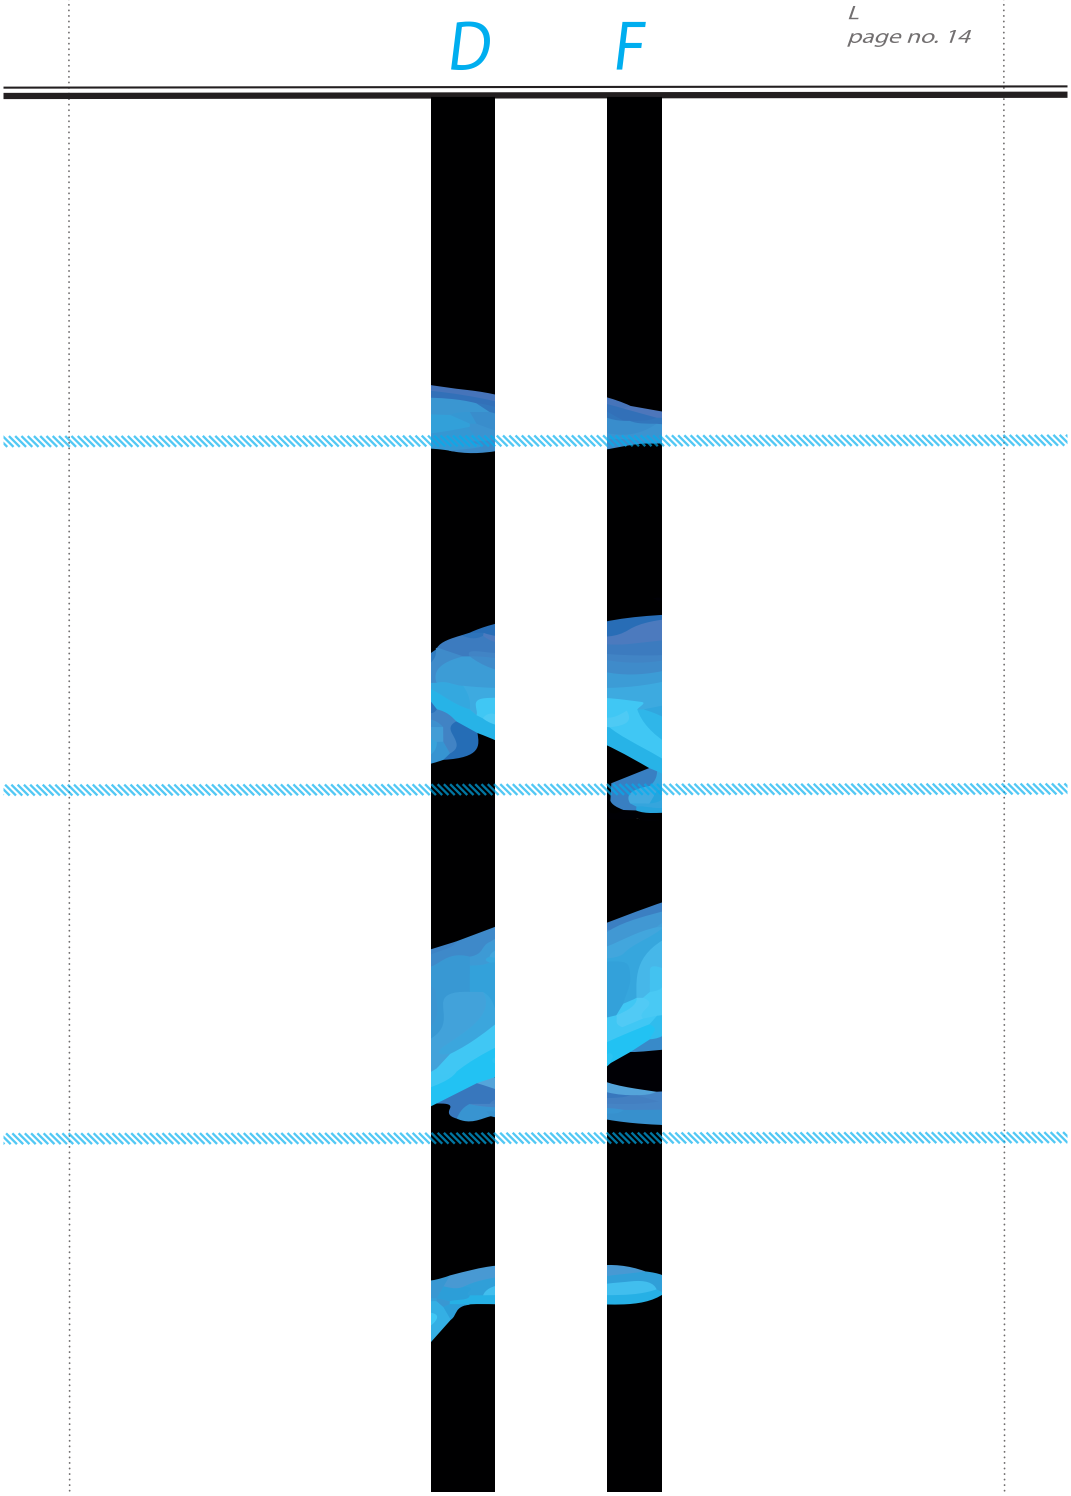
D
F

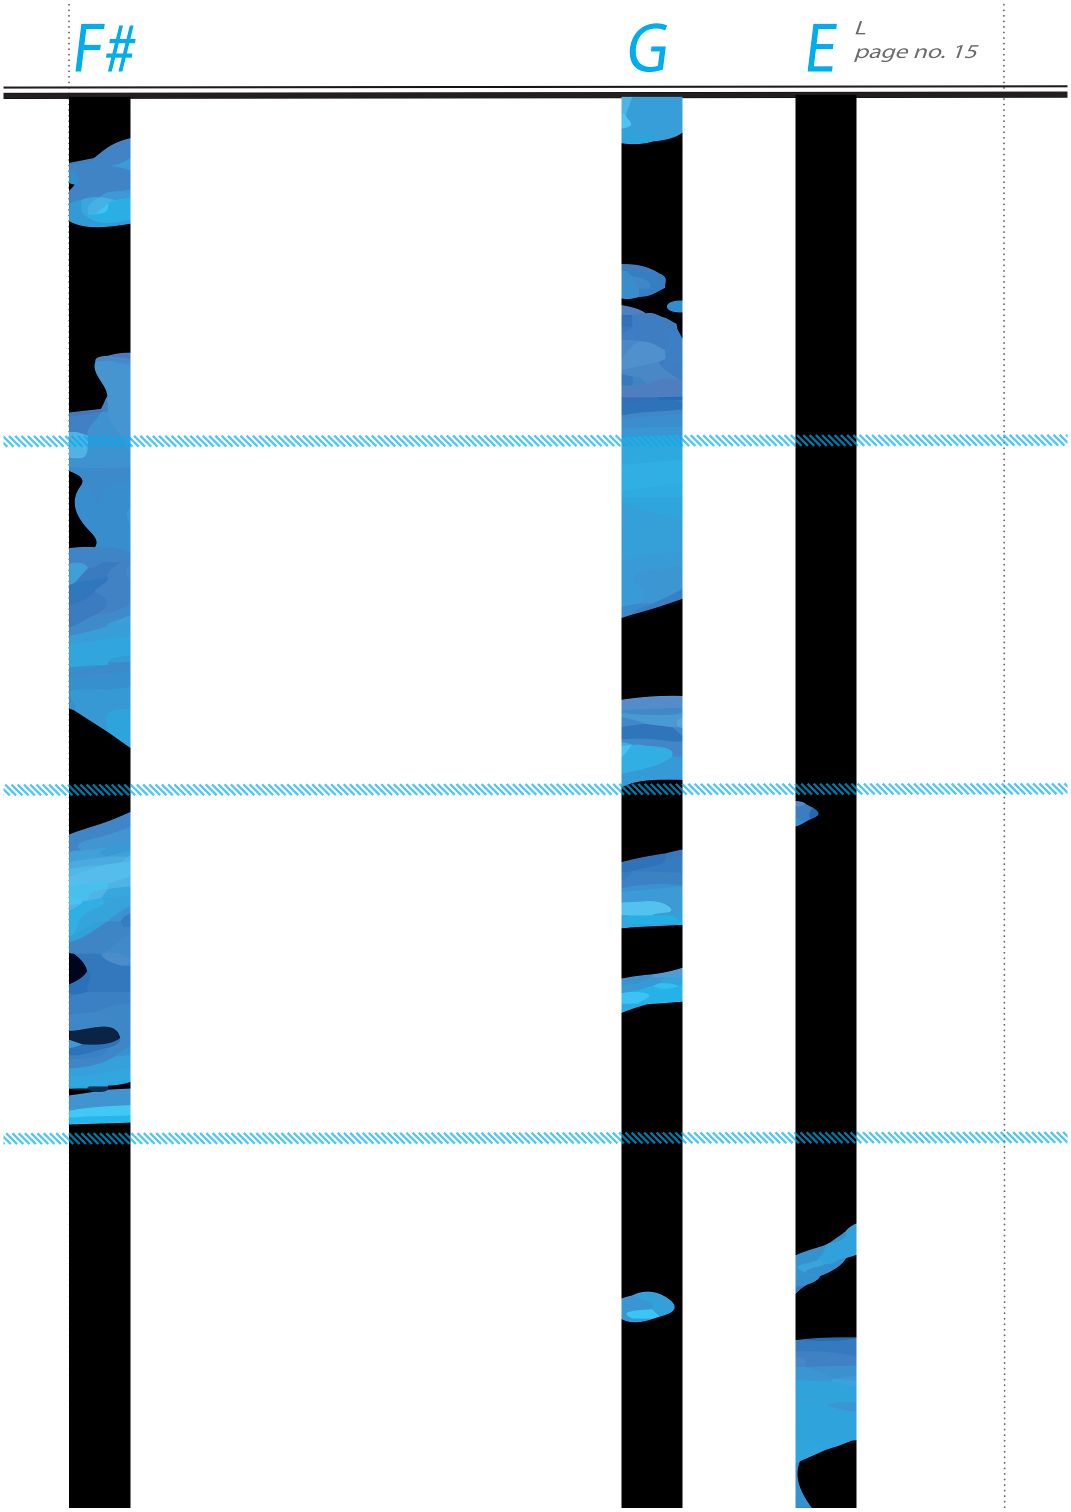
F#
G
E
L
page no. 15

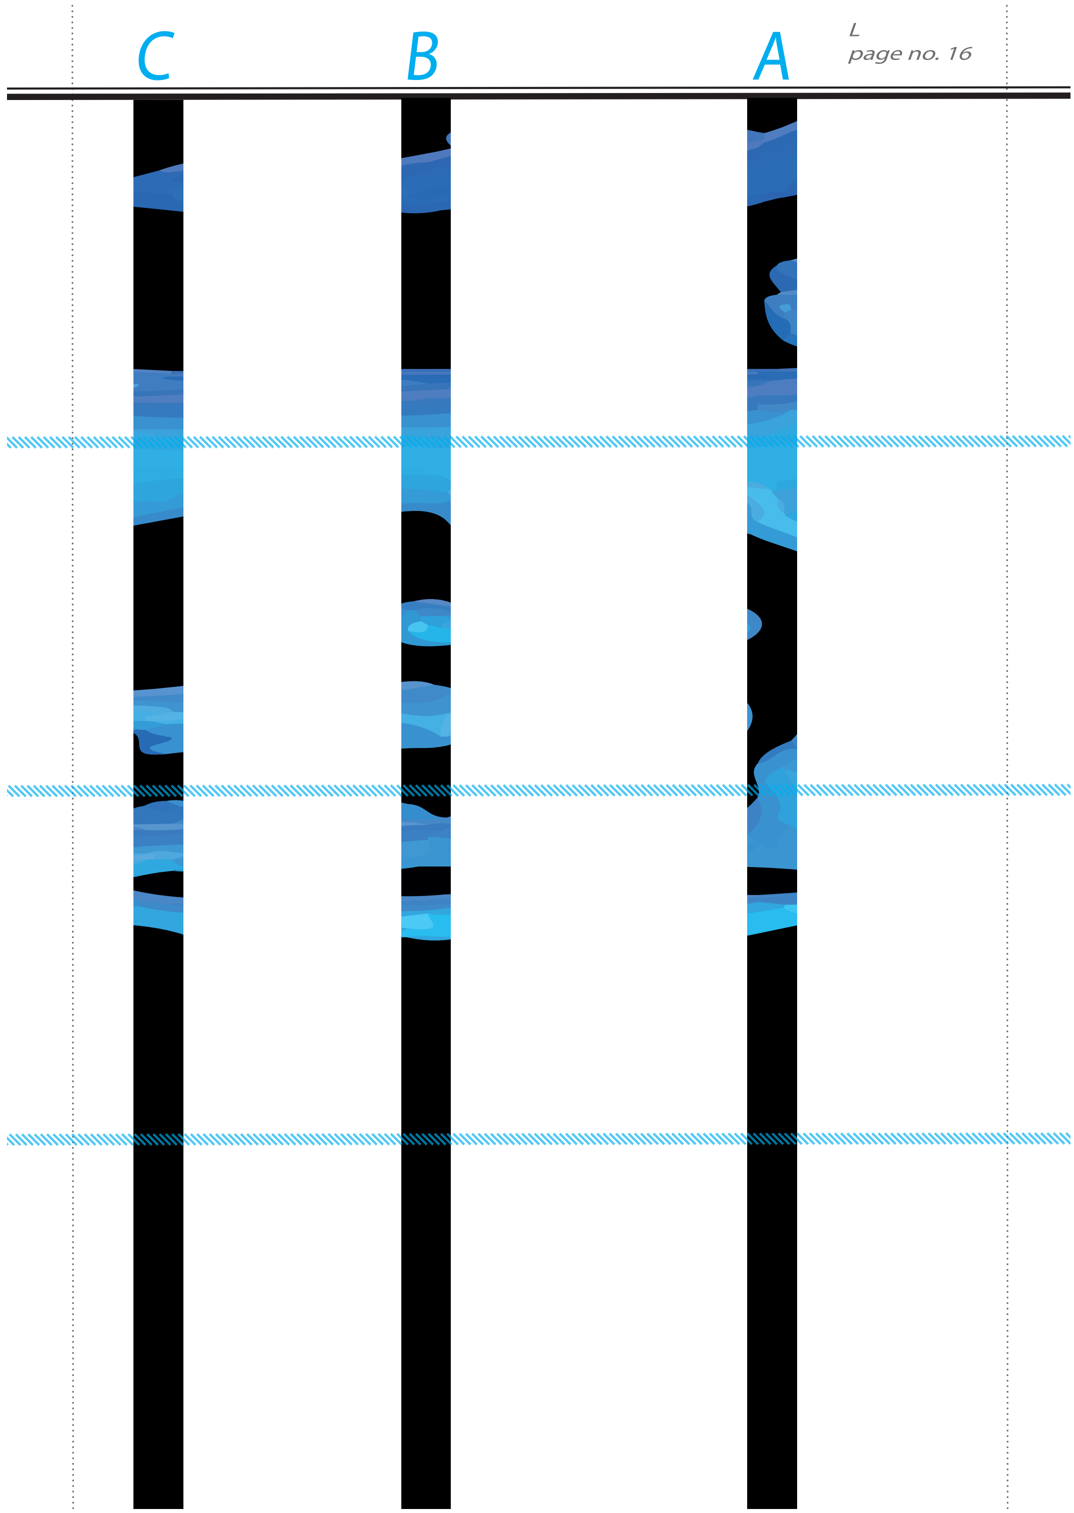
C
B
A
L
page no. 16

4 hands
1 piano + reach
+ .wav
richard stenton
2018

A.7 Goodbye Pork Pie Hat - Charles Mingus

A.8 Musica Ricercata n.7 - György Ligeti

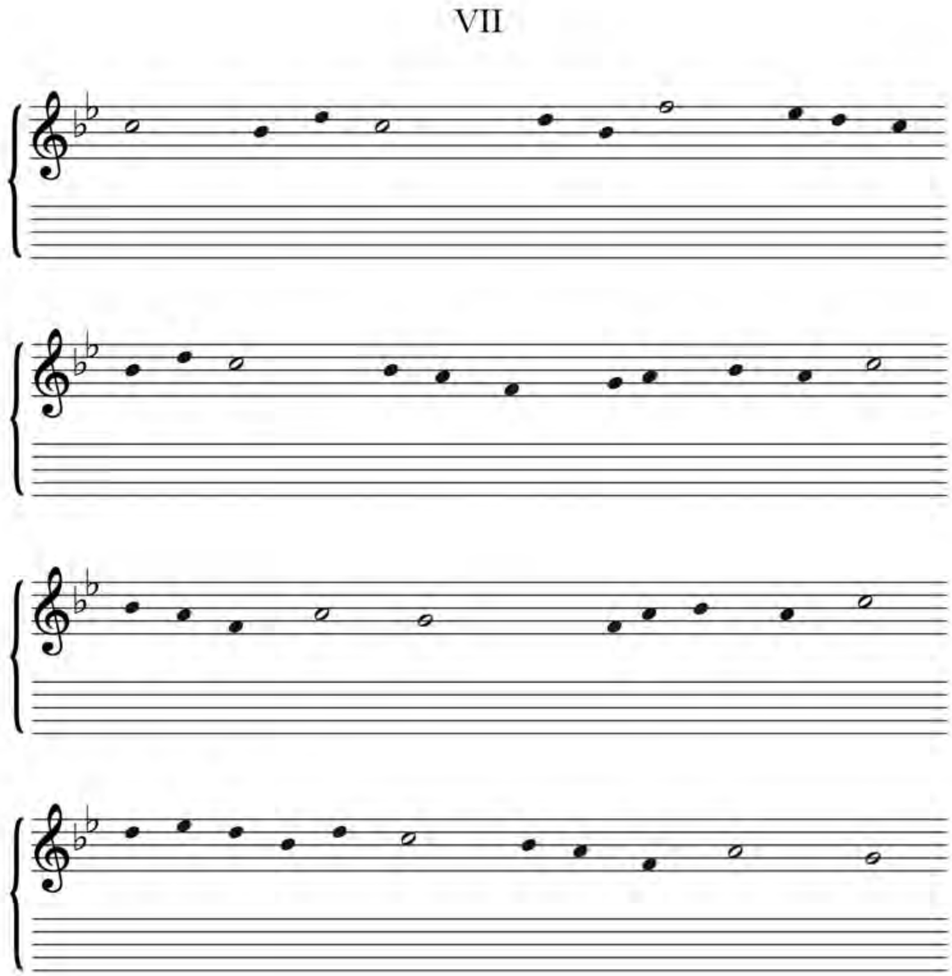

A.9 Music

Corea, C. (1998) 'Spain'. In: *Light As A Feather* [audio stream] Available through: https://www.spotify.com/uk/ [Accessed July 2019]

A.10 Websites

120 Years of Electronic Music (2019) *120 Years of Electronic Music - The history of electronic musical instruments from 1800 to 2019* Available at: http://120years.net [Accessed February 2019]

FCC greenlights Soli, Google's radar-based gesture tech (2019) *TechCrunch* Available at: https://techcrunch.com/2019/01/02/us-fcc-approves-google-soli-project/ [Accessed March 2019]

Pure Data (2019) *Pure Data*. Available at: http://puredata.info [Accessed February 2019]

Reach v0.1 (2017) *Reach v0.1*. Available at: https://github.com/NiccoloGranieri/Reach-v0.1

UEQ - User Experience Questionnaire (2018) *UEQ*. Available at: https://www.ueq-online.org [Accessed January 2019].

A.11 Videos

//Aria alla Francese for Piano and Reach System (2018) *//Aria alla Francese*. [video] Available at: https://www.youtube.com/watch?v=gzLTwMbGjxM [Accessed July 2019].

12.jpg, 4 hands 1 piano + reach + .wav (2018) *12.jpg*. [video] Available at: https://youtu.be/X2FMMex2LKI [Accessed September 2019].

Amon Tobin Interview (2014) *Amon Tobin Interview*. [video] Avialable at: https://vimeo.com/91024101 [Accessed 23rd of October 2017].

Building virtual instruments: case studies of gestural innovation in works for piano and electronics (2015) *Building Virtual Instruments*. [video] Available at: https://medias.ircam.fr/x07703d [Accessed June 2015].

Listen (2018) *Listen*. [video] Available at: https://youtu.be/pC8Fze2sJVs [Accessed September 2019].

Loop – Music in Motion: Performing with mi.mu gloves. (2017) *Loop – Music in Motion: Performing with mi.mu gloves*. [video] Available at: https://www.youtube.com/watch?v=7fczzuaXxBw [Accesed June 2017].

Machine Learning for Musicians and Artists (2016) *Machine Learning for Musicians and Artists*. [video] Available at: https://www.kadenze.com/courses/machine-learning-for-musicians-and-artists/info [Accessed June 2016].

Wax & Wane (live feat. Reach) (2018) *Wax & Wane*. [video] Available at: https://youtu.be/a0wCGydfons [Accessed September 2019].

Bibliography

Aguilera, F., Alarcón, R. A., Guerrero, L. A. and Collazos, C. A. (2006) *A Cognitive Model of User Interaction as a Guideline for Designing Novel Interfaces*. In: Ochoa, S. F. and Roman, G.-C., eds., *Advanced Software Engineering: Expanding the Frontiers of Software Technology*, pp. 62–76. Springer US.

Antle, A. N., Corness, G. and Droumeva, M. (2009) *Human-Computer-Intuition? Exploring the Cognitive Basis for Intuition in Embodied Interaction*. *International Journal of Arts and Technology*, vol. 2(3):pp. 235–254. Available at: https://doi.org/10.1504/IJART.2009.028927.

Barrett, E. and Bolt, B. (2014) *Practice as Research: Approaches to Creative Arts Enquiry*. I.B.Tauris.

Bernardo, F., Arner, N. and Batchelor, P. (2017) *O Soli Mio: Exploring Millimeter Wave Radar for Musical Interaction*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'17, pp. 283–286. Copenhagen, Denmark: Aalborg University Copenhagen.

Bolt, B. (2007) *The Magic Is in the Handling*. In: *Practice as Research: Approaches to Creative Arts Enquiry*. I.B.Tauris.

Bolter, J. D. and Gromala, D. (2003) *Windows and Mirrors: Interaction Design, Digital Art, and the Myth of Transparency*. Cambridge, MA, USA: MIT Press.

Bowker, G. and Star, S. (2000) *Sorting Things Out: Classification and Its Consequences*. Inside Technology. MIT Press.

Bowman, D. A., Gabbard, J. L. and Hix, D. (2002) *A Survey of Usability Evaluation in Virtual Environments: Classification and Comparison of Methods*. *Presence: Teleoperators and Virtual Environments*, vol. 11(4):pp. 404–424. Available at: https://doi.org/10.1162/105474602760204309.

Breiman, L. (2001) *Random Forests*. *Machine Learning*, vol. 45(1):pp. 5–32. Available at: https://doi.org/10.1023/A:1010933404324.

Brent, W. (2012) *The Gesturally Extended Piano*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'12. Ann Arbor, Michigan: University of Michigan.

Brown, A. R. and Sorensen, A. (2009) *Integrating Creative Practice and Research in the Digital Media Arts*. In: Smith, H. and Dean, R. T., eds., *Practice-Led Research, Research-Led Practice in the Creative Arts*, Research Methods for the Arts and Humanities, pp. 153–165. Edinburgh University Press.

Brown, D., Nash, C. and Mitchell, T. (2018) *Simple Mappings, Expressive Movement: A Qualitative Investigation into the End-User Mapping Design of Experienced Mid-Air Musicians*. *Digital Creativity*, vol. 29(2-3):pp. 129–148. Available at: https://doi.org/10.1080/14626268.2018.1510841.

Bullock, J., Beattie, D. and Turner, J. (2011) *Integra Live : A New Graphical User Interface for Live Electronic Music*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'11, pp. 387–392. Oslo, Norway.

Bullock, J., Michailidis, T. and Poyade, M. (2016) *Towards a Live Interface for Direct Manipulation of Spatial Audio*. In: *Proceedings of the 2016 International Conference on Live Interfaces*. University of Sussex: Experimental Music Technologies (EMuTe) Lab with REFRAME Books.

Cadoz, C. and Wanderley, M. M. (2000) *Gesture - Music*. In: Battier, M., ed., *Trends in Gestural Control of Music*, pp. 71–94. Paris, France: IRCAM - Centre Pompidou.

Caramiaux, B., Bevilacqua, F., Palmer, C. and Wanderley, M. (2017) *Individuality in Piano Performance Depends on Skill Learning*. In: *Proceedings of the 4th International Conference on Movement Computing*, MOCO '17, pp. 14:1–14:7. London, United Kingdom: ACM, Available at: https://doi.org/10.1145/3077981.3078046.

Chadefaux, D., Le Carrou, J.-L., Fabre, B., Daudet, L. and Quartier, L. (2010) *Experimental Study of the Plucking of the Concert Harp*. In: *Proceedings of the 20th International Symposium on Music Acoustics*.

Christopoulos, C. (2014) *From Gesture to Sound: A Study for a Musical Interface Using Gesture Following Techniques*. Master's Thesis, Universitat Pompeu Fabra, Barcelona, Spain.

Cronin, D. (2014) *Usability of Micro-vs. Macro-Gestures in Camera-Based Gesture Interaction*. San Luis Obispo, California, USA: California Polytechnic State University.

Csikszentmihalyi, M. (2014a) *Attention and the Holistic Approach to Behavior*. In: *Flow and the Foundations of Positive Psychology*, pp. 1–20. Springer.

Csikszentmihalyi, M. (2014b) *Play and Intrinsic Rewards*. In: *Flow and the Foundations of Positive Psychology*, pp. 135–153. Springer.

Csikszentmihalyi, M. and Csikszentmihalyi, I. (1975) *Beyond Boredom and Anxiety. Experiencing Flow in Work and Play*, vol. 721. San Francisco, USA: Jossey-Bass.

Dahl, L. (2016) *Designing New Musical Interfaces as Research: What's the Problem? Leonardo*, vol. 49(1):pp. 76–77. Available at: https://doi.org/10.1162/LEON_a_01118.

Dahlstedt, P. (2017) *Physical Interactions with Digital Strings - A Hybrid Approach to a Digital Keyboard Instrument*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'17, pp. 115–120. Copenhagen, Denmark: Aalborg University Copenhagen.

Delalande, F. (1988) *La Gestique de Gould: Élements Pour Une Sémiologie Du Geste Musical*. In: *Glenn Gould Pluriel*, Text Collected and Presented by Ghyslaine Guertin. L. Courteau.

Di Donato, B. (2018) *MyoSpat: A Gesture Controlled System for Embodied Audiovisual Interaction in Live Performance*. PhD Thesis, Royal Birmingham Conservatoire, Birmingham City University, Birmingham, United Kingdom.

Dourish, P. (2004) *Where the Action Is: The Foundations of Embodied Interaction*. Bradford Bks. MIT Press.

Erol, A., Bebis, G., Nicolescu, M., Boyle, R. D. and Twombly, X. (2007) *Vision-Based Hand Pose Estimation: A Review. Computer Vision and Image Understanding Special Issue on Vision for Human-Computer Interaction*, vol. 108(1):pp. 52–73. Available at: https://doi.org/https://doi.org/10.1016/j.cviu.2006.10.012.

Fels, S., Gadd, A. and Mulder, A. (2002) *Mapping Transparency through Metaphor: Towards More Expressive Musical Instruments*. *Organised Sound*, vol. 7(2):pp. 109–126. Available at: https://doi.org/10.1017/S1355771802002042.

Fischer, G. (1998) *Beyond "Couch Potatoes": From Consumers to Designers*. In: *Proceedings. 3rd Asia Pacific Computer Human Interaction*, pp. 2–9. Available at: https://doi.org/10.1109/APCHI.1998.704130.

Gaver, W. (2012) *What Should We Expect from Research Through Design?* In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, pp. 937–946. Austin, Texas, USA: ACM, Available at: https://doi.org/10.1145/2207676.2208538.

Gay, G. and Hembrooke, H. (2004) *Activity-Centered Design: An Ecological Approach to Designing Smart Tools and Usable Systems*. Acting with Technology. MIT Press.

Gillian, N. and Paradiso, J. A. (2014) *The Gesture Recognition Tookit*. *The Journal of Machine Learning Research*, vol. 15(1):pp. 3483–3487.

Godøy, R. and Leman, M. (2010) *Musical Gestures: Sound, Movement, and Meaning*. Routledge.

Grandhi, S. A., Joue, G. and Mittelberg, I. (2011) *Understanding Naturalness and Intuitiveness in Gesture Production: Insights for Touchless Gestural Interfaces*. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pp. 821–824. Vancouver, BC, Canada: ACM, Available at: https://doi.org/10.1145/1978942.1979061.

Green, O. (2014) *NIME, Musicality and Practice-Led Methods*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'14, pp. 1–6. London, United Kingdom: Goldsmiths, University of London.

Green, O. (2016) *The Situation of Practice-Led Research Around NIME, and Two Methodological Suggestions for Improved Communication*. *Leonardo*, vol. 49(1):pp. 78–79. Available at: https://doi.org/10.1162/LEON_a_01119.

Guna, J., Jakus, G., Pogačnik, M., Tomažič, S. and Sodnik, J. (2014) *An Analysis of the Precision and Reliability of the Leap Motion Sensor and Its Suitability for Static and Dynamic Tracking*. *Sensors*, vol. 14(2):pp. 3702–3720.

Hadjakos, A., Aitenbichler, E. and Mühlhäuser, M. (2008) *The Elbow Piano : Sonification of Piano Playing Movements*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'08, pp. 285–288. Genoa, Italy.

Hadjakos, A. and Waloschek, S. (2014) *VIP: Controlling the Sound of a Piano with Wrist-Worn Inertial Sensors*. In: *International Conference on New Interfaces for Musical Expression, Keyboard Salon Workshop*, vol. 30.

Haken, L., Abdullah, R. and Smart, M. (1992) *The Continuum: A Continuous Music Keyboard*. In: *Proceedings of the 1992 International Computer Music Conference*, pp. 81–84.

Han, J., Shao, L., Xu, D. and Shotton, J. (2013) *Enhanced Computer Vision With Microsoft Kinect Sensor: A Review*. *IEEE Transactions on Cybernetics*, vol. 43(5):pp. 1318–1334. Available at: https://doi.org/10.1109/TCYB.2013.2265378.

Hayes, G. R. (2011) *The Relationship of Action Research to Human-Computer Interaction*. *ACM Transactions on Computer-Humam Interaction*, vol. 18(3):pp. 15:1–15:20. Available at: https://doi.org/10.1145/1993060.1993065.

Heidegger, M. (1962) *Being and Time*. Oxford: Blackwell publishers.

Hemery, E., Manitsaris, S., Moutarde, F., Volioti, C. and Manitsaris, A. (2015) *Towards the Design of a Natural User Interface for Performing and Learning Musical Gestures*. *Procedia Manufacturing*, vol. 3:pp. 6329–6336. Available at: https://doi.org/https://doi.org/10.1016/j.promfg.2015.07.952. 6th International Conference on Applied Human Factors and Ergonomics (AHFE 2015) and the Affiliated Conferences, AHFE 2015.

Hutchinson, H., Mackay, W., Westerlund, B., Bederson, B. B., Druin, A., Plaisant, C., Beaudouin-Lafon, M., Conversy, S., Evans, H., Hansen, H. et al. (2003) *Technology Probes: Inspiring Design for and with Families*. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '03, pp. 17–24. Ft. Lauderdale, Florida, USA: ACM, Available at: https://doi.org/10.1145/642611.642616.

Ihde, D. (1979) *Technics and Praxis*. Boston Studies in the Philosophy of Science. D. Reidel Pub. Co.

Jack, R. H., Harrison, J., Morreale, F. and McPherson, A. P. (2018) *Democratising DMIs: The Relationship of Expertise and Control Intimacy*. In: Luke Dahl, T. M., Douglas Bowman, ed., *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'18, pp. 184–189. Blacksburg, Virginia, USA: Virginia Tech.

Jack, R. H., Stockman, T. and McPherson, A. (2017) *Rich Gesture, Reduced Control: The Influence of Constrained Mappings on Performance Technique*. In: *Proceedings of the 4th International Conference on Movement Computing*, MOCO '17, pp. 15:1–15:8. London, United Kingdom: ACM, Available at: https://doi.org/10.1145/3077981.3078039. ACMID: 3078039.

Johnston, A. (2011) *Beyond Evaluation : Linking Practice and Theory in New Musical Interface Design*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'11, pp. 280–283. Oslo, Norway.

Johnstone, E. (1985) *The Rolky: A Poly-Touch Controller for Electronic Music*. In: *Proceedings of the 1985 International Computer Music Conference*. Burnaby, BC, Canada.

Jordà, S. (2005) *Digital Lutherie Crafting Musical Computers for New Musics Performance and Improvisation*. PhD Thesis, Universitat Pompeu Fabra, Barcelona, Spain.

King, E. (2006) *Supporting Gestures: Breathing in Piano Performance*. In: Gritten, A. and King, E., eds., *Music and Gesture*, pp. 142–164. Ashgate.

Kursa, M., Rudnicki, W., Wieczorkowska, A., Kubera, E. and Kubik-Komar, A. (2009) *Musical Instruments in Random Forest*. In: Rauch, J., Raś, Z. W., Berka, P. and Elomaa, T., eds., *Foundations of Intelligent Systems*, pp. 281–290. Berlin, Heidelberg: Springer Berlin Heidelberg.

Lähdeoja, O., Wanderley, M. M. and Malloch, J. (2009) *Instrument Augmentation Using Ancillary Gestures for Subtle Sonic Effects*. In: *Proceedings of the 6th Sound and Music Computing Conference*, pp. 327–330. Porto, Portugal.

Lamb, R. and Robertson, A. (2011) *Seaboard : A New Piano Keyboard-Related Interface Combining Discrete and Continuous Control*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'11, pp. 503–506. Oslo, Norway.

Latour, B. (1987) *Science in Action: How to Follow Scientists and Engineers Through Society*. Harvard University Press.

Laugwitz, B., Held, T. and Schrepp, M. (2008) *Construction and Evaluation of a User Experience Questionnaire*. In: Holzinger, A., ed., *HCI and Usability for Education and Work*, pp. 63–76. Berlin, Heidelberg: Springer Berlin Heidelberg.

Leman, M. (2008) *Embodied Music Cognition and Mediation Technology*. Mit Press. MIT Press.

Lewis, A. and Pestova, X. (2012) *The Audible and the Physical: A Gestural Typology for 'Mixed' Electronic Music*. In: *Proceedings of the Electroacoustic Music Studies Network Conference*, pp. 1–13. Stockholm.

Lien, J., Gillian, N., Karagozler, M. E., Amihood, P., Schwesig, C., Olson, E., Raja, H. and Poupyrev, I. (2016) *Soli: Ubiquitous Gesture Sensing with Millimeter Wave Radar*. *ACM Trans. Graph.*, vol. 35(4):pp. 142:1–142:19. Available at: https://doi.org/10.1145/2897824.2925953.

Maes, P.-J., Leman, M., Lesaffre, M., Demey, M. and Moelants, D. (2010) *From Expressive Gesture to Sound*. *Journal on Multimodal User Interfaces*, vol. 3(1):pp. 67–78. Available at: https://doi.org/10.1007/s12193-009-0027-3.

Magnusson, T. (2009) *Of Epistemic Tools: Musical Instruments as Cognitive Extensions*. *Organised Sound*, vol. 14(2):pp. 168–176. Available at: https://doi.org/10.1017/S1355771809000272.

Magnusson, T. (2010) *Designing Constraints: Composing and Performing with Digital Musical Systems*. *Computer Music Journal*, vol. 34(4):pp. 62–73. Available at: https://doi.org/10.1162/COMJ_a_00026.

Mainsbridge, M. and Beilharz, K. (2014) *Body As Instrument: Performing with Gestural Interfaces*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'14, pp. 110–113. London, United Kingdom: Goldsmiths, University of London.

McNutt, E. (2003) *Performing Electroacoustic Music: A Wider View of Interactivity*. *Organised Sound*, vol. 8(3):pp. 297–304. Available at: https://doi.org/10.1017/S135577180300027X.

McPherson, A. (2010) *The Magnetic Resonator Piano: Electronic Augmentation of an Acoustic Grand Piano*. *Journal of New Music Research*, vol. 39(3):pp. 189–202. Available at: https://doi.org/10.1080/09298211003695587.

McPherson, A. (2012) *TouchKeys: Capacitive Multi-Touch Sensing on a Physical Keyboard*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'12. Ann Arbor, Michigan: University of Michigan.

McPherson, A. P. (2017) *2012: TouchKeys: Capacitive Multi-Touch Sensing on a Physical Keyboard*. In: Jensenius, A. R. and Lyons, M. J., eds., *A NIME Reader: Fifteen Years of New Interfaces for Musical Expression*, vol. 3, pp. 419–432. Cham: Springer International Publishing.

Michailidis, T. (2016) *On the Hunt for Feedback: Vibrotactile Feedback in Interactive Electronic Music Performances*. PhD Thesis, Royal Birmingham Conservatoire, Birmingham City University, Birmingham, United Kingdom.

Moore, F. R. (1988) *The Dysfunctions of MIDI*. *Computer Music Journal*, vol. 12(1):pp. 19–28.

Morreale, F., Guidi, A. and McPherson, A. (2019) *Magpick: An Augmented Guitar Pick for Nuanced Control*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'19, pp. 1–6. Porto Alegre, Brasil: Federal University of Rio Grande do Sul.

Morreale, F. and McPherson, A. (2017) *Design for Longevity: Ongoing Use of Instruments from NIME 2010-14*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'17, pp. 192–197. Copenhagen, Denmark: Aalborg University Copenhagen.

Morreale, F., McPherson, A. P. and Wanderley, M. (2018) *NIME Identity from the Performer's Perspective*. In: Dahl, L., Bowman, D. and Martin, T., eds., *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'18, pp. 168–173. Blacksburg, Virginia, USA: Virginia Tech.

Nardi, B. (1996) *Context and Consciousness: Activity Theory and Human-Computer Interaction*. Mit Press. MIT Press.

Nicolls, S. L. (2010) *Interacting with the Piano*. PhD Thesis, Brunel University, London, United Kingdom.

Nicolls, S. L. and Gillian, N. (2012) *A Gesturally Controlled Improvisation System for Piano.* In: *Live Interfaces: Performance Art Music Proceedings.*

Nijs, L. (2017) *The Merging of Musician and Musical Instrument.* In: Lesaffre, M., Maes, P. and Leman, M., eds., *The Routledge Companion to Embodied Music Interaction*, Routledge Music Companions, pp. 49–57. Taylor & Francis.

Nijs, L., Lesaffre, M. and Leman, M. (2009) *The Musical Instrument as a Natural Extension of the Musician.* In: *Proceedings of the 5th Conference of Interdisciplinary Musicology*, pp. 132–133. LAM - Institut Jean Le Rond d'Alembert.

Norman, D. (1990) *Why Interfaces Don't Work.* In: Laurel, B. and Mountford, S. J., eds., *The Art of Human-Computer Interface Design*, pp. 209–219. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc.

Nosulenko, V. N., Barabanshikov, V. A., Brushlinsky, A. V. and Rabardel, P. (2005) *Man–Technology Interaction: Some of the Russian Approaches. Theoretical Issues in Ergonomics Science*, vol. 6(5):pp. 359–383. Available at: https://doi.org/10.1080/14639220500070051.

O'Modhrain, S. (2011) *A Framework for the Evaluation of Digital Musical Instruments. Computer Music Journal*, vol. 35(1):pp. 28–42. Available at: https://doi.org/10.1162/COMJ_a_00038.

Paradiso, J. A. (2004) *Electronic Controllers for Musical Performance and Interaction.* p. 18.

Pestova, X. (2008) *Models of Interaction in Works for Piano and Live Electronics.* PhD Thesis, McGill University, Montreal, Canada.

Pestova, X. (2017) *Approaches to Notation in Music for Piano and Live Electronics: The Performer's Perspective.* In: Sallis, F., Bertolani, V., Burle, J. and Zattra, L., eds., *Live Electronic Music: Composition, Performance, Study*, Routledge Research in Music. Taylor & Francis.

Pierce, A. (2010) *Deepening Musical Performance Through Movement: The Theory and Practice of Embodied Interpretation.* Musical Meaning and Interpretation. Indiana University Press.

Pu, Q., Gupta, S., Gollakota, S. and Patel, S. (2013) *Whole-Home Gesture Recognition Using Wireless Signals.* In: *Proceedings of the 19th Annual International*

Conference on Mobile Computing, MobiCom '13, pp. 27–38. Miami, Florida, USA: ACM, Available at: https://doi.org/10.1145/2500423.2500436.

Puckette, M. S., Apel, T. and Zicarelli, D. (1998) *Real-Time Audio Analysis Tools for Pd and MSP*. In: *Proceedings of the 1998 International Computer Music Conference, ICMC 1998, Ann Arbor, Michigan, USA, October 1-6, 1998*.

Rabardel, P. (1995) *Les Hommes et Les Technologies; Approche Cognitive Des Instruments Contemporains*.

Raykov, Y. P., Ozer, E., Dasika, G., Boukouvalas, A. and Little, M. A. (2016) *Predicting Room Occupancy with a Single Passive Infrared (PIR) Sensor Through Behavior Extraction*. In: *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '16, pp. 1016–1027. ACM, Available at: https://doi.org/10.1145/2971648.2971746.

Rebelo, P. (2006) *Haptic Sensation and Instrumental Transgression*. *Contemporary Music Review*, vol. 25(1-2):pp. 27–35. Available at: https://doi.org/10.1080/07494460600647402.

Schrepp, M., Hinderks, A. and Thomaschewski, J. (2014) *Applying the User Experience Questionnaire (UEQ) in Different Evaluation Scenarios*. In: Marcus, A., ed., *Design, User Experience, and Usability. Theories, Methods, and Tools for Designing the User Experience*, pp. 383–392. Springer International Publishing.

Shapin, S. (1989) *The Invisible Technician*. *American Scientist*, vol. 77(6):pp. 554–563.

Smith, D. and Wood, C. (1981) *The 'USI', or Universal Synthesizer Interface*. In: *Audio Engineering Society Convention 70*.

Ssergejewitsch, T. L. (1928) *Method of and Apparatus for the Generation of Sounds*.

Stowell, D., Robertson, A., Bryan-Kinns, N. and Plumbley, M. (2009) *Evaluation of Live Human–Computer Music-Making: Quantitative and Qualitative Approaches*. *International Journal of Human-Computer Studies - Special issue on Sonic Interaction Design*, vol. 67(11):pp. 960–975. Available at: https://doi.org/https://doi.org/10.1016/j.ijhcs.2009.05.007.

Sturman, D. J. (1992) *Whole-Hand Input*. PhD Thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States of America.

Sung, K. (2011) *Recent Videogame Console Technologies*. *Computer*, vol. 44(02):pp. 91–93. Available at: https://doi.org/10.1109/MC.2011.55.

Taylor, R., Schofield, G., Hook, J., Ladha, K., Bowers, J. and Wright, P. (2013) *Crafting Interactive Systems: Learning from Digital Art Practice*. In: *CHI '13 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '13, pp. 3223–3226. Paris, France: ACM, Available at: https://doi.org/10.1145/2468356.2479652.

Thorn, S. D. (2018) *Alto.Glove: New Techniques for Augmented Violin*. In: Luke Dahl, T. M., Douglas Bowman, ed., *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'18, pp. 334–339. Blacksburg, Virginia, USA: Virginia Tech.

Tits, M., Tilmanne, J., D'Alessandro, N. and Wanderley, M. M. (2015) *Feature Extraction and Expertise Analysis of Pianists' Motion-Captured Finger Gestures*. In: *Proceedings of the International Computer Music Conference*, ICMC'15, pp. 102–105.

Waisvisz, M. (1985) *The Hands : A Set of Remote MIDI Controllers*. In: *Proceedings of the International Computer Music Conference*, ICMC'85.

Wanderley, M. M. and Battier, M. (2000) *Trends in Gestural Control of Music*. Paris, France: IRCAM - Centre Pompidou.

Wang, S., Song, J., Lien, J., Poupyrev, I. and Hilliges, O. (2016) *Interacting with Soli: Exploring Fine-Grained Dynamic Gesture Recognition in the Radio-Frequency Spectrum*. In: *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, UIST '16, pp. 851–860. Tokyo, Japan: ACM, Available at: https://doi.org/10.1145/2984511.2984565.

Wilson, A. D. (2010) *Using a Depth Camera As a Touch Sensor*. In: *ACM International Conference on Interactive Tabletops and Surfaces*, ITS '10, pp. 69–72. Saarbrücken, Germany: ACM, Available at: https://doi.org/10.1145/1936652.1936665.

Wilson, C. (2013) *Semi-Structured Interviews*. In: *Interview Techniques for UX Practitioners: A User-Centered Design Method*. Elsevier Science.

Winters, R. M. and Wanderley, M. M. (2012) *New Directions for Sonification of Expressive Movement in Music*. In: *Proceedings of the 18th International Conference on Auditory Display*. Atalanta, GA, USA.

Wixon, D. and Good, M. (1987) *Interface Style and Eclecticism: Moving beyond Categorical Approaches*. *Proceedings of the Human Factors Society Annual Meeting*, vol. 31:pp. 571–575. Available at: https://doi.org/10.1177/154193128703100521.

Xambo, A., Laney, R., Dobbyn, C. and Jorda, S. (2013) *Video Analysis for Evaluating Music Interaction: Musical Tabletops*. In: Holland, S., Wilkie, K., Mulholland, P. and Seago, A., eds., *Music and Human-Computer Interaction*, pp. 241–258. London: Springer London.

Yang, Q. and Essl, G. (2012) *Augmented Piano Performance Using a Depth Camera*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'12. Ann Arbor, Michigan: University of Michigan.

Zandt-Escobar, A. V., Caramiaux, B. and Tanaka, A. (2014) *PiaF: A Tool for Augmented Piano Performance Using Gesture Variation Following*. In: *Proceedings of the International Conference on New Interfaces for Musical Expression*, NIME'14, pp. 167–170. London, United Kingdom: Goldsmiths, University of London.

Zhao, C., Chen, K.-Y., Aumi, M. T. I., Patel, S. and Reynolds, M. S. (2014) *SideSwipe: Detecting In-Air Gestures Around Mobile Devices Using Actual GSM Signal*. In: *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, UIST '14, pp. 527–534. Honolulu, Hawaii, USA: ACM, Available at: https://doi.org/10.1145/2642918.2647380.