

**The Beatback System:
Exploring Interactive Percussion for
Promoting Rhythmic Practice**

by

Andrew Hawryshkewich
B.A., Carleton University, 2008

THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE

in the
School of Interactive Arts & Technology
Faculty of Communication, Arts and Technology

© Andrew Hawryshkewich 2011

SIMON FRASER UNIVERSITY

Fall 2011

All rights reserved. However, in accordance with the *Copyright Act of Canada*, this work may be reproduced, without authorization, under the conditions for *Fair Dealing*. Therefore, limited reproduction of this work for the purposes of private study, research, criticism, review and news reporting is likely to be in accordance with the law, particularly if cited appropriately.

APPROVAL

Name: Andrew Hawryshkewich
Degree: Masters of Science
Title of Thesis: The Beatback System: Exploring Interactive Percussion for Promoting Rhythmic Practice

Examining Committee:

Chair:

Dr. Kate Hennessy
Assistant Professor,
School of Interactive Arts and Technology

Dr. Philippe Pasquier
Senior Supervisor
Assistant Professor,
School of Interactive Arts and Technology

Dr. Arne Eigenfeldt
Supervisor
Associate Professor,
School of Contemporary Arts, SFU

Steve DiPaola
External Examiner
Associate Professor,
School of Interactive Arts and Technology

Date Defended/Approved: 23 September 2011



SIMON FRASER UNIVERSITY
LIBRARY

Declaration of Partial Copyright Licence

The author, whose copyright is declared on the title page of this work, has granted to Simon Fraser University the right to lend this thesis, project or extended essay to users of the Simon Fraser University Library, and to make partial or single copies only for such users or in response to a request from the library of any other university, or other educational institution, on its own behalf or for one of its users.

The author has further granted permission to Simon Fraser University to keep or make a digital copy for use in its circulating collection (currently available to the public at the "Institutional Repository" link of the SFU Library website <www.lib.sfu.ca> at: <<http://ir.lib.sfu.ca/handle/1892/112>>) and, without changing the content, to translate the thesis/project or extended essays, if technically possible, to any medium or format for the purpose of preservation of the digital work.

The author has further agreed that permission for multiple copying of this work for scholarly purposes may be granted by either the author or the Dean of Graduate Studies.

It is understood that copying or publication of this work for financial gain shall not be allowed without the author's written permission.

Permission for public performance, or limited permission for private scholarly use, of any multimedia materials forming part of this work, may have been granted by the author. This information may be found on the separately catalogued multimedia material and in the signed Partial Copyright Licence.

While licensing SFU to permit the above uses, the author retains copyright in the thesis, project or extended essays, including the right to change the work for subsequent purposes, including editing and publishing the work in whole or in part, and licensing other parties, as the author may desire.

The original Partial Copyright Licence attesting to these terms, and signed by this author, may be found in the original bound copy of this work, retained in the Simon Fraser University Archive.

Simon Fraser University Library
Burnaby, BC, Canada

STATEMENT OF ETHICS APPROVAL

The author, whose name appears on the title page of this work, has obtained, for the research described in this work, either:

(a) Human research ethics approval from the Simon Fraser University Office of Research Ethics,

or

(b) Advance approval of the animal care protocol from the University Animal Care Committee of Simon Fraser University;

or has conducted the research

(c) as a co-investigator, collaborator or research assistant in a research project approved in advance,

or

(d) as a member of a course approved in advance for minimal risk human research, by the Office of Research Ethics.

A copy of the approval letter has been filed at the Theses Office of the University Library at the time of submission of this thesis or project.

The original application for approval and letter of approval are filed with the relevant offices. Inquiries may be directed to those authorities.

Simon Fraser University Library
Simon Fraser University
Burnaby, BC, Canada

ABSTRACT

This thesis details the development and research of Beatback, an interactive percussion system for promoting rhythmic practice. Beatback is a software based system which with MIDI-enabled hardware controllers - such as an electronic drum-kit - allows users to play with their own rhythmic material in two interaction modes: (1) Call-response allows users to reflect on their own playing with system generated responses learnt from the user's own performance. (2) Accompaniment enables users to build up complex rhythmic patterns by layering their own looped drum patterns. The first of two studies focused on drummers practising patterns with the system filtering out (or zoning) drums being played by the user, which found significant benefits to the zoned method. Research focusing on the Beatback system in its entirety demonstrated that both naive and experienced drummers feel more competent (in call-response) and enjoy interacting (in accompaniment) with the system significantly more than having open time to play the drum-kit. The results from both of these studies suggest the possibility of employing systems such as Beatback to benefit those practicing or learning how to play the drum-kit.

ACKNOWLEDGEMENTS

I would like to thank my senior supervisor, Philippe Pasquier, who has been exceedingly patient whilst I learnt the ins and outs of grad school. Also, to all my colleagues who were there to lend some support or direction when I needed it most, thank-you greatly.

TABLE OF CONTENTS

Approval	ii
Abstract	iii
Acknowledgements	iv
Table of Contents	v
List of Figures	vii
List of Tables	ix
1: Introduction	1
1.1 Motivations	1
1.2 The Beatback System	3
1.3 Beatback Validation & Evaluation	4
1.4 Contributions	5
2: Interactive Percussion Design & Research	7
2.1 Designing Percussion Interaction	7
2.2 A Review of Musical Interface Design	8
2.2.1 A Proliferation of the Specific	8
2.2.2 Developing Commonalities	10
2.2.3 Musical Interaction Design	11
2.2.4 Augmentation	13
2.2.5 Interactive & Generative Musical Systems	16
2.2.6 Interactive Percussion	21
2.2.7 Methods of Evaluation	24
2.3 Design Goals for Percussion Interaction	27
3: The Beatback System	29
3.1 Implementation	30
3.2 Data Parsing	32
3.2.1 Parsing MIDI Data	33
3.2.2 Pattern Learning	34
3.2.3 Pattern Storage	36
3.2.4 Pattern Generation	39
3.2.5 Drum-kit Zoning	41
3.2.6 User Interaction	44
3.2.7 Pattern Loading, Storage & Training	45
3.2.8 The GUI	46
4: Beatback Research	47
4.1 Interactive Percussion in Rhythmic Exploration	47
4.1.1 Design Validation as Research	47
4.1.2 Benefits of Rhythm Exploration	48

4.2	Research Design	49
4.3	Drum-kit Zoning Research	49
4.3.1	Methods	51
4.3.2	Equipment and Setup.....	51
4.3.3	Study Procedure	52
4.3.4	Data Collection.....	54
4.3.5	Results & Discussion	56
4.4	Beatback Interaction Research.....	61
4.4.1	Methods	61
4.4.2	Equipment and Setup.....	62
4.4.3	Study Procedure	62
4.4.4	Data Collection.....	64
4.4.5	Results & Discussion	64
4.5	Limitations & Future Work.....	70
5:	Conclusions	74
	Appendices.....	79
	Appendix A.....	79
	Appendix B.....	80
	Appendix C: Data Appendix CD-ROM	82
	Reference List	83

LIST OF FIGURES

Figure 1: The first pattern as read by Beatback and then parsed into in-out pairings.	36
Figure 2: The new pattern as read by Beatback is added to the table, updating the values.	39
Figure 3: The user has played some more, resulting in some additional transitions. Highlighted are possible transition possibilities for {S, H}	40
Figure 4: Example of the default drum-kit zoning used in Beatback.	43
Figure 5: The Beatback user interface allows for control of tempo, interaction modes and quantization. Additionally there are more in depth setup options and an ability to load and save data.	46
Figure 6: Eight-bar rhythmic pattern given to participants.....	53
Figure 7: Participant's perceived enjoyment (higher percentage, more enjoyable)	56
Figure 8: Participant's perceived tension (higher percentage, more tense)	56
Figure 9: Participant's perceived competence (higher percentage, felt more competent).....	56
Figure 10: Note density as percentage of original pattern (between tasks)	59
Figure 11: Percent of notes played accurately (between tasks)	59
Figure 12: Comparison of note accuracy and density within the zoned task	59
Figure 13: Comparison of note accuracy and density within the non-zoned task.....	60
Figure 14: Graph showing IMI results for participants in three conditions. A higher percentage indicates higher perceived experience. The error bars indicate the standard deviation.	65
Figure 15: Average task time (in minutes) compared to metronome use between the three conditions. Error bars indicate standard deviation.....	66
Figure 16: Balance between participant and system performance amount within conditions. Measured by comparing the note density as performed by the participant and the system.....	66
Figure 17: Note density comparison between the participant and system during the accompaniment task.....	67
Figure 18: Note density comparison between the participant and system during the call-response task.....	68
Figure 19: A demonstration and description of how timing crawl occurred during study	71

Figure 20: Figure demonstrating the timing gaps that occurred during the study. 72

LIST OF TABLES

Table 1: Musical note lengths as ticks	31
Table 2: Intrinsic Motivation Inventory questions with terminology used for the study, and what each question assesses.	55
Table 3: Intrinsic Motivation Inventory questions with terminology used for the study, and what each question assesses (questions altered from Table 2 to clarify intent)	62
Table 4: Prior musical experience questions for exploration research participants.	63

1: INTRODUCTION

This thesis details Beatback: an interactive percussion system geared towards rhythmic practise and exploration. Included in the introduction is an overview of what interactive musical systems are, the motivations, research and contributions of Beatback, and an overview of the structure and content of this thesis document.

1.1 Motivations

The field of musical interface design is constantly expanding. There is enough development in the field to warrant its own conference titled New Interfaces for Musical Expression (NIME). Every year NIME features new ways in which we can interact with or manipulate sound. This expansion of new interfaces stems from tools such as MaxMSP that make it possible for performers or musicians to become designers of their own musical interactions (Section 2.2.2). Performers acting as designers results in the advantage of them being able to control the details of the musical interaction or output of a given interface. Allowing their interface designs to be tailored to their intents (Section 2.2.3).

One disadvantage that results, is that should the interaction involved be novel – it requires additional knowledge outside of a performer’s existing training – it requires a performer to retrain should they want to use that interface (Section 2.2.4). Although some interfaces are not intended for widespread use, those that

are and work with a novel style of interaction require the performer to become familiar or train on the interface. It is for this reason that Beatback works on augmenting an existing interface, the digital drum-kit, as it already has a trained set of performers, and augmentation subsequently forms a design goal of the system (Section 2.3).

With this goal of augmentation in mind, the research looked at examples of augmented systems, and in particular, what constitutes a successful system (Section 2.3). One such system was the Continuator (Pachet, 2003a), which worked specifically with melodic content from pianos and guitars to generate musically appropriate responses to a performer's input. The work on the Continuator led Pachet to develop the concept of *interactive reflexive musical systems* (IRMS), which are designed to produce an impression of similarity to the user's input into the system (Pachet, 2003b). Achieving stylistically appropriate responses such as outlined with IRMS forms another goal behind the design of the Beatback system (Section 2.3).

The final motivation behind the development of the Beatback system stems from research performed with the Continuator, looking at how children interacted with the system (Addessi, 2007). With kids, when given open time to play with the system they discovered that they demonstrated high levels of focused attention and enjoyment in their interactions with the system. Both of these are demonstrated benefits in learning or practise scenarios (Section 2.2.4), a benefit which Beatback seeks to capitalize upon.

1.2 The Beatback System

The Beatback system is designed primarily to work as an augmented drum-kit. This means that it takes an existing interface – an electronic drum-kit – and augments its functionality through a system built in Max 5 (Cycling '74, 2010). The purpose of augmentation rather than the development of a new interface are twofold: (1) the system benefits by working with an interface (the drum-kit) that has a field of professionals (drummers) already familiar with it (Section 2.2.4). (2) The development of novel interaction can be the focus, in lieu of concerns relating to hardware or interface development.

Beatback uses variable-order Markov models to parse, store and generate responses to user input in a perceptual real-time interaction (Section 2.2.5). The use of Markov models enables the system to store rhythmic patterns as paired portions, with input and output patterns matched through a set of weighed transitions. The weighing and matching of patterns ensures the system is capable of providing varied output that is still a suitably expected response to the user's input. In other words, enabling it to emulate the user's style.

User input influences the system's response in one of two interaction models, either call and response or accompaniment, both of which are based on common models for musical interaction (Section 2.2.3). In call and response, the user can engage in a dialogue where the system learns from the user's performance while they play, and then responds based on the learnt material whenever the user stops playing. Within the accompaniment mode, the system

learns and responds to each region of drums, called zones, separately (Section 3.2.5).

1.3 Beatback Validation & Evaluation

In order to evaluate Beatback, two studies were conducted: the first study focused on (1) drum-kit zoning – the filtering of drums from a pattern being learnt – and its effects on practised percussionists. The second study looked at the Beatback system’s interaction models, having participants interact with a learning and responsive system. Both studies relating to Beatback use a mixed methods design, presenting quantitative data with qualitative feedback for analysis and discussion of the results.

Drum-kit zoning research (Section 4.3) looks at practiced percussionists in a pattern-learning task that compared participant performance in zoned and non-zoned conditions. Drums which are zoned are filtered out of the pattern being learnt, where no zoning meant the pattern played continuously. Percussionists with more than four years of informal experience and/or two years of formal training participate in two learning tasks; in which they are asked to play along with a complex rhythmic pattern as best as they can, given the two conditions listed above. MIDI data is recorded in order to analyse how the percussionists’ practice changes given the task, and the Intrinsic Motivation Index (IMI) (University of Rochester, 2010) was used to measure their self-reported enjoyment, tension and competence in each task. The analysis that follows the research (Section 4.3.5) discusses how percussionists engaged in the zoned task demonstrated a significantly higher motivation than without zoning.

Research relating to the interaction modes within Beatback (Section 4.4) focuses instead on both novice and practiced percussionists playing with Beatback and their performance whilst doing so. Earlier work with children and an interactive piano (Pachet & Adnessi, 2004) suggests that interactive musical systems can offer benefits to musical exploration, which can translate to benefits in learning. In this study, participants interact with Beatback in both call-response and accompaniment interaction modes (Section 3.2.4) to assess whether Beatback offers benefits similar to those identified during research by Pachet & Adnessi (2004). Participants filled in the IMI and had their performance recorded as MIDI data for analysis. Results demonstrated that both call-response and accompaniment interaction models offer benefits to user's enjoyment and competence when playing the drums (Section 4.3 & 4.4).

1.4 Contributions

As far as we know, the Beatback system is the first interactive reflexive musical system (Section 2.2.3) that focuses on rhythm and percussion. Though, as identified, there are limitations to the current system (Section 4.5), studies run with participants using the system have demonstrated that it improves their intrinsic motivation (Section 4.3.5 & 4.4.5) by increasing self-perceived enjoyment and competence, while reducing perceived tension. Within the studies, both practised and naïve percussionists expressed these improvements to motivations, in learning a pattern and simply interacting with the Beatback system itself. To the best of our knowledge, the first implementation and study of drum zoning (Section 3.2.5) – or the filtering out of drums being played by a user

– is explored with Beatback. With relation to drum-kit zoning, that the IMI results demonstrate significant improvement to motivation when using zoning suggests that it could be a beneficial learning or practise tool for percussionists (Section 4.3). Stemming from the IRMS model set out by Pachet, the Beatback system works with rhythmic content and has demonstrated that IRMS can be used to benefit learning and practise in a rhythmic context.

2: INTERACTIVE PERCUSSION DESIGN & RESEARCH

This section provides insight into the development and history of musical system design¹. Starting from current commentaries on musical interface and system design, this literature review fleshes out themes prevalent in musical systems to help understand how the Beatback system contributes to the field. This will assist in establishing goals for what an effective musical system should constitute, how to approach effective percussive interaction, and appropriate evaluation strategies for such a system.

2.1 Designing Percussion Interaction

Within the field of musical system design, there are no written standards to the design process. The variety of approaches to design, discussion, and trying to draw out standards echoes the interdisciplinary and creatively unhindered nature of the field. This also makes effectively communicating new designs and ideas difficult when the audience themselves may have a varying level of understanding. The variety of methods by which designers² report their process, code, develop their interface, and validate it for its purpose speaks to the oddly broad number of disciplines and professions which contribute to the development of new musical systems. This makes any discussion in the design of new musical

¹ “Musical system design” is used in this paper to refer to the design of any musical interfaces and/or software

² Designers in the context of this review are simply the developers of new interfaces, regardless of their discipline or training.

interfaces require the consideration of any research domains they encompass to establish a more rigorous design. Articles like that of Serra (2004) (Section 2.2.2 and 2.2.7) further demonstrate how to ensure music technology research is effectively presented across all the fields it encompasses. To ensure an effective presentation, this review draws from (but is not exclusive to) musicology, human-computer interaction (HCI) and artificial intelligence to more clearly establish the motivations behind Beatback in relation to prior related works.

2.2 A Review of Musical Interface Design

This literature review is designed to provide background to the design goals for the Beatback system. These goals are organized and presented in Section 2.3.

2.2.1 A Proliferation of the Specific

Machover (2002) identified that the development of new musical systems revolved around their tailoring to a specific performance or function, rather than aiming for applications that are more generic. He prescribed that musical system design should move past being predominantly applied to very specific contexts, and move towards enticing further interest and interaction from general users. This identifies an argument that comes to the fore often in discussions on musical systems which is that of their prolific yet specific nature.

The reason for a proliferation of the specific is two-fold: firstly, the availability of low-cost software (i.e. Max/MSP, Puredata) and hardware (i.e. Arduino) to develop new systems enables a performer to become the designer.

In this paradigm, the performer prescribes what they need a musical system to be able to do, be it with regards to control setup, sound generation or otherwise. Given a specific purpose and easy access to tools (as mentioned), the performer can become the designer of their own system.

Jordà (2004) expressed a similar concern over the proliferation of rarely used interfaces, which become obsolete after their purpose has been fulfilled, such as for a performance. Focusing on the physical interface, he identified that virtuosity³ with new interfaces was not common for performers because of standardized mouse and keyboard, or MIDI-based systems being available. This preference for more common interface devices is suggestive of something that is also later confirmed; that performers beginning with a new system are frequently frustrated with having to learn the complexities of its interface. Furthermore, while citing that virtuosity on traditional interfaces – such as the piano or guitar – can require upwards of ten years of experience (Lehman, 1997), Jordà questioned whether it was even possible to design an interface with a low starting complexity but with no cap to the level of virtuosity attainable. This concept of virtuosity or specialization in an instrument leads to the second reason for why there is such a proliferation of musical interfaces, one of access.

This difficulty of access is not concerning code or specifications, but instead cognitive access. Repeatedly identified (Magnusson and Mendieta 2007, Orio and Wanderley 2002, Dannenberg 2007, Serra 2005) is that the design and research of new interfaces for musical expression sits mostly in the academic

³ Jordà describes virtuosity as the point of being able to control the performance of a given interface or instrument.

domain, which means that specialists in fields such as music, computer sciences, or psychology are all contributing to the design of new systems in various ways. While this diversity is excellent for broadening the field, it complicates discussions, where the detailing of high-level theory or precise functionality of a system's musical material or interaction may be difficult to understand for an individual not specialized in the field that is describing the system. As a result, it is more difficult to establish common ground with which to assess (and discuss) musical interface.

2.2.2 Developing Commonalities

Orio and Wanderley (2002) approach the development of musical systems from a human-computer interaction (HCI) methodology. Their argument is that without guidelines, the creation of new devices is highly unrestricted, resulting in an open creativity, a creativity from which the field has benefited, exploring the development of new interfaces openly. Yet, the lack of a common system of discussion also makes it difficult to generate new interfaces that build effectively on prior works. Orio and Wanderley (2002) therefore propose the four metrics of; learnability, explorability, feature controllability and timing controllability by which to define and categorize musical systems. As suggested by the names (learnability, control-ability) these metrics demonstrate a focus on the user's experience and abilities with a given system. The approach of looking at the user experience is one that is drawn upon in the development of Beatback. Poepel (2005) uses similar considerations though directed more towards performance with an instrument by using a model of encoding expressive cues within the

performance. Elements like the accuracy of timing, pitch or dynamics are suggested as possible user cues to encode, ultimately garnering insight into the expressiveness of a system. Most recently, Magnusson and Mendieta (2007) offer a similar though simpler model, suggesting that interface design is interested primarily in either ergonomic (physical setup) or timbral (audible output) considerations. These three different approaches to grouping interfaces based on high-level interaction affordances – user experience, expressive cues and ergonomic or timbral considerations – demonstrate how attempting to group systems based on their interaction tends to miss other important elements.

Perhaps the most successful generalization of musical interfaces is found in the work of Miranda and Wanderley (2006), who instead focus on hardware affordances. In lieu of trying to draw out commonalities in interaction, the focus is on sensorial input – for example, gesture or bio-signal based sensors – through which they develop a more common ground upon which to discuss musical system design. Miranda and Wanderley's approach to generalization demonstrates a reasonable solution for using language appropriate across disciplines, an issue commonly found when dealing with interdisciplinary writing (Salter & Hearn, 1996).

2.2.3 Musical Interaction Design

As was just discussed, there is no one universal approach to designing or analysing musical interaction, though a language and process appropriate to the field is important. With this in mind, some consideration should be given to musical interaction design, to help draw out directions from prior work in the field.

One approach to interaction design is covered by Hunt et al (2000), who discuss the mapping of controls to a sound and quite reasonably posit that this mapping is important to understanding musical interaction with a given system. If the mapping of an interface changes, the user will need to develop their skills to encompass the change. Continuing with the focus on the importance of user and interface interaction, Dobrian and Koppelman (2006) discuss how musical expression is the performers' doing. This not a new concept, but Dobrian and Koppelman also present that control does not equal expression. For when a performer engages a system with a high level of control, they may not be able to comfortably control all the elements to generate an expression they may deem reasonable. A system that is too limited in control of input may not offer the expressive ability a performer would like, and similarly a system with too many control variables may be too complex for a performer to consider. Interaction with a musical system must strike a balance between control and the performer's capabilities, achieving what Moore describes as *control intimacy*, or a match between desirable sounds and the capabilities of the performer (1988).

Concerning striking this balance, a good example is Pachet's Continuator⁴, a software based machine learning system. When coupled with a piece of hardware such as a MIDI keyboard, the Continuator is capable of modelling a user's performance, and then interpreting input for generating musical continuations. The system has been tested with both children and professional pianists (Pachet, 2003a), and in both cases achieved the balance

⁴ Technical specifics of the system are detailed further in section 2.2.5

between desirable output and the abilities of the performer. As a result, Pachet (2006) describes developing systems to achieve a similar musical interaction style, entitled *interactive musical reflective systems* (IMRS). Pachet discusses that the focus of IMRS is the interaction with the performer, where the system provides the performer with a musical response reflecting their own style and abilities (call and response), enticing them to explore interaction with it further (Pachet, 2003b). The *control intimacy* (Moore, 1988) described earlier provides an apt description of the Continuator as users suggested a high level of comfort in the interaction, in addition to desirable output by the system (Pachet, 2006). It also achieves a form of *interaction intimacy*: an unexpected though appropriate response, as perceived by the performer. This is further compounded by the use of a common musical interface – the piano – that offers the benefit of not requiring the user to learn a new style of physical interaction. As is discussed in the next section, this is one benefit offered by the augmentation of existing interfaces.

2.2.4 Augmentation

By moving towards the augmentation of an existing instrument rather than the development of a novel interface, more time can be spent focusing on the development of meaningful interaction. Augmentation of a musical instrument or interface is the addition of functionality through software or hardware additions that capitalize the existing base of professional users of the instrument. The creation of augmented interfaces or environments has been demonstrated (Price and Rogers 2004, Pachet and Adressi 2004, Ferguson 2006) to benefit user

experience and interaction, which is of relevance to Beatback (see 4.1.2, Benefits of Rhythm Exploration). This portion of the review fleshes out a rationale for using augmented interfaces in promoting musical exploration.

Looking at interaction in augmented spaces (not music specific) Price and Rogers discuss how physicality in interaction is important to the development and promotion of self-directed learning (2004). Though somewhat separate in context from musical interfaces, the study is important for providing possible benefits of augmentation. This includes (but is not limited to), enabling creative exploration through interaction, creating an authenticity to the interaction by using physical objects (as opposed to virtual ones), and an anticipation (for activity) which comes out of the expected physical and unexpected augmented interactions (Price & Rogers, 2004). All of these benefits were reflected in the research of Pachet and Adressi (2004), which looked at children's interactions with the Continuator (augmented piano). Their results demonstrated that children spent more time focused on interacting with the system, were more attentive to it, and were more explorative in their playing when compared to time at a standard piano. Part of the interest could be associated with the simple fact that a musical cue followed their performance, yet the research demonstrated that the interest was prolonged more than it would be in a simple action-reaction scenario.

Developing from the informal learning experience of Pachet and Adressi, others such as Ferguson (2006), have taken the concept of augmentation to a more formalized learning situation. For Ferguson, augmented instruments or interfaces geared towards formal musical learning – such as in a music school –

should offer interactive sonification to their students or learners (Ferguson, 2006). Interactive sonification being an activity in which the user engages with a common musical interface, and while repeating a similar action receives live audio or visual feedback on the qualities of their performance. Visual feedback in interactive sonification is a more commonly discussed application, offering the user an abstracted view of their auditory input for them to interpret (Hunt and Hermann 2004, Pauletto and Hunt 2004). The one challenge with using visual feedback is that the user must interpret visual cues while moderating their auditory performance. While with auditory cues, such interpretation can be blended with their own sound formation. Specifically, Ferguson introduces auditory icons; an example of which would be the metronome. It helps musicians develop an innate understanding of rhythm or tempo but the ticking can also become lost to one's hearing through grouping or masking effects; where a weaker sound is covered up by a louder one. Ferguson (2006) even discusses how novice musicians can tend to play after the metronome, to ensure that they are matching the beating of the metronome.

Considering the uses of augmentation in both informal and formal learning presented above, it is important to note a couple of themes. (1) That the interaction with the system is prolonged (and iterative), (2) that the system's feedback (or performance) is distinguishable from that of the user, and (3) that the system itself offers live interaction, a concept discussed further below.

2.2.5 Interactive & Generative Musical Systems

Starting with a look at live music performance, Dannenberg (2007) started by addressing the field of popular music performance, and discusses how systems designed by academia rarely became popularized or used commonly. To create more accessible interfaces, Dannenberg suggested parameters to promote interface use in live musical performance. These parameters suggested that the system should be autonomous, not requiring constant user attention, and should use an established tempo but be capable of moving through musical motifs with the performer (Dannenberg, 2007). The concept of having an autonomous system is an idea which Pachet (2006) further addressed with IMRS: by having no interface available to the performer apart from the piano itself, the user was forced to focus on the system as if it was the piano (see Section 2.3).

Joel Chadabe suggested similar concepts to that of Pachet, particularly that the quasi-predictable nature of the interaction should add information to the user's performance, and be able to cue further performance by the user (Chadabe, 1984). Addressing interactivity more recently, Croft (2007) suggested that recent interactive musical systems rarely offer interactivity to a user. With this in mind, he also establishes a series of parameters that can be generalized to two points: that interaction should offer (1) musically appropriate responses by the system, and those responses should be (2) sufficiently expected by the performer (Croft, 2007). So for Croft, a 'live' interaction requires the computer to almost act like another performer might, listening and responding to the

performer's input, and maintaining an expected relationship (through response) between the performer and the system. This expectation of Croft reiterates points addressed by Chadabe, Pachet, and the concept of *interaction intimacy*, where the relationship between the user and the system must not be entirely predictable. Johnston et al (2009) echo a need for quasi-predictability and term it "conversational interaction" or that interaction with a musical system to be similar to the kind of interaction offered by real users, it should not be entirely predictable. Although the musical content should be similar to the user's input, integrating a level of unexpected responses ensures that a system closer approximates the way musicians would interact.

To generate musical responses, either expected or unexpected, a distinct set of computational approaches is frequently employed in interactive musical systems. The majority of these approaches are concerned with effectively representing musical information as a series, and being capable of interpreting and reworking that series. Since there are a couple of approaches to generative music, they have been grouped algorithmically.

A number of rhythm specific generative musical systems use an agent-based design to generate musical material (Eigenfeldt and Kapur, 2008; Murray-Rust, 2003; Levisohn and Pasquier, 2008). An agent-based system is one in which a series of software agents work in conjunction with one another to attempt to achieve a programmatic goal (Alonso, 2002). With regards to musically generative system, these agents are considered much like a musical ensemble, with each of the agents performing and communicating with one another

(Eigenfeldt & Kapur, 2008). The one concern with agent-based works such as the Kinetic Engine (Eigenfeldt & Kapur, 2008), BeatBender (Levisohn & Pasquier, 2008) or VirtuaLatin (Murray-Rust, 2003) is that frequently the control of these systems is limited once the agents start performing. In the Kinetic Engine, the generated material is interesting due (in part) to the agents that interact with minimal human interference to try to coalesce to a common rhythm. Often the intent of systems such as the Kinetic Engine is to have a composer define a space in which the system should perform, and then let it run within the defined parameters (Eigenfeldt, 2006). This design tends to result in these systems catering more towards autonomous generation of material, offering limited human interaction. Although a valid source of interesting generative material, the concern with being able to take user performance material into generation is important when considering *interaction intimacy* (Section 2.2.4).

Employing genetic algorithms in the generation of musical material offers a great variety in (autonomously) generated material, allowing input material to be bred into unique hybrids of the original (Weinberg, Godfrey, Rae, & Rhoads, 2007). Genetic algorithms work by taking a population and using biology inspired operations – such as crossovers (breeding) or mutations – to create a new generation with a set of fitness values. The intent of these operations is to find population children who are closest to the fitness function. As a result of the large-scale generation of populations, these algorithms are relatively popular for use in the generation of musical material. An early example of a generative system employing genetic algorithms is GenJam (Biles, 1994), which used the

ability of these algorithms to explore variation in the generation of jazz solos. Similarly exploring the generation of material, though within the rhythmic domain was the work of Horowitz (1994), who developed an “interactive” genetic algorithm by having the user select favourable rhythms with which to select populations that should continue. The challenge of interacting with genetic algorithms while they work is best demonstrated by the work of Weinberg and Haile – a robotic percussionist. Working within the perceptual real-time generation of musical material, Haile presents exploration into real-time use of genetic algorithms in music generation. The algorithm employed in the generation of material listens to user input and then generates output based on a pre-established repertoire of “contextually relevant” material (Weinberg, Godfrey, Rae, & Rhoads, 2007). Weinberg et al also discuss the limitations of genetic algorithms in responding to real-time user input. Within the Haile system, the real-time interaction offered is mimicking user input, until the system has enough time to generate a new population or select a rhythmic pattern from existing repertoire. One system that does provide a possible solution to balancing generation and interaction using genetic algorithms is the most recent iteration of the Kinetic Engine (Eigenfeldt, 2009). The solution to interaction in generation offered by Kinetic Engine is through parsing given musical content to form the base from which it generates content. Although the Kinetic Engine begins to break the limitations of genetic algorithms, those same limits lead next to the discussion of Markov models.

Markov model based systems offers the ability of the system to generate material based on a corpus, and their application can enable a system to become reactive through listening to user input. In the case of rhythm, looking at the last three input notes by a performer, a Markov chain (in this case 3rd order – three notes) would be generated by stochastically choosing one of the probable continuations for the input pattern (Kochanski, 2005). Perhaps the most accomplished example is that of the Continuator by Pachet (2003), which in evaluation was able to generate live musical continuations based on the live performance of pianists (Pachet F. , 2004). Similarly, the use of Markov model based systems seems more common in interactive or live settings, as the user's input can quickly and readily influence the output of the system. Work such as that of Kelley (2003) or Eigenfeldt and Pasquier (2010), and discussions such as those of Rae (2009) further reinforce this fact; Markov models provide a means of encoding user input for generating musical material. Finalizing this, Weinberg's shift to using a mix of genetic algorithms and Markov models for the more recent Shimon (robotic percussionist), away from the purely genetic algorithms used with Haile is indicative of the limitation of genetic algorithms when used in a live, interactive, setting (Weinberg, Raman, & Mallikarjuna, 2009).

Systems that mix different methods of generating and listening to user input move much more towards interacting with a user. As has been mentioned, Shimon, the robotic musician (Weinberg, Raman, & Mallikarjuna, 2009) used a combination of Markov models and genetic algorithms to interpret and generate musical material. Perhaps more successful, is the work of Assayag et al on the

OMax system (Assayag et al, 2006), which uses a combination of Variable-Order Markov models, probabilistic suffix trees, and factor oracles. This combination of pattern parsing methods allows the OMax system to learn, capture and work with patterns performed by a user allowing for the ability to generate sophisticated musical responses in an interactive setting.

2.2.6 Interactive Percussion

Within the field of interactive percussion, there are a few common approaches to designing the interaction. As with the prior section, there is a variety of methods by which to digitize, interpret and generate a percussive performance or material. One approach used in capturing user input is through looking at the gesture – the up and down movement of the drumstick or the placement of hands on a table – and there are a variety of systems that employ this approach when providing feedback (Borchers, 2007; Burgess and Mynatt, 1994; Young and Fujinaga, 2004; Chun, Hawryshkewich, Jung, & Pasquier, 2010).

Perhaps one of the earliest examples of using a gesture-based approach is WorldBeat (Borchers, 1997). While also capable of mimicking other musical gestures, it offered a simple drumstick playing gesture by holding and moving two batons. It was through this that the user could elicit percussive sounds from the system. Along similar lines, both the Virtual Drum Kit (Burgess & Mynatt, 1994) and Aobachi (Young & Fujinaga, 2004) exemplify a similar approach to mimicking percussion interaction: both systems use baton or drumstick like objects for the user to move and generate sound with. Recently, the Wii-mote

has become quite popular for use in general hardware hacking or interface design, mostly due to the ease with which it can be refactored into a multi-axis sensory device. Further exemplifying hardware hacking, internet searches tend to garner the results of hobbyists tinkering with converting Wii-mote or iPhone gestural input into percussive output (Wii Drum Machine, 2006; More Realistic Wii Drum Kit, 2007; Tanaka & Parkinson, 2010)

The Wii-mote (as a gaming interface) draws attention next to the recent proliferation of video games that focus on interactive percussion. With the ubiquity of mobile devices growing ever larger, it is becoming more common to see percussion-based games or interactions. Games such as ZOOZBeat for the iPhone, or Drummer for the Nintendo DS, simply exemplify how much more accessible percussion interaction is becoming, though the translation of common percussive actions is often simplified. ZOOZBeat functions by having the user shake the iPhone, which is closer to a gestural interaction (Weinberg, Beck, & Godfrey, 2009), while Drummer has the user tap or slide a stylus across a screen to interact with the game (Bianchi & Yeo, 2009). Perhaps in response to the number of mobile or portable percussion games available, larger entertainment systems employ more common physical models of drum-kits as their interface. Game franchises such as Rock Band (2010) and Guitar Hero (2010) both offer players the ability to use actual drumsticks on a model kit to play along with popular music songs. Rock Band offers players the ability to play the entirety of drum parts – e.g. kick, snare, cymbals, toms, hi-hat – instead of simply providing

four pads and a pedal (2010). This moves the gaming interfaces to the point that the player is performing the actions of playing an actual, physical, drum-kit.

While playing video games allows for feedback on one's performance based on in-game cues and fosters the social feelings of belonging in an ensemble (Bhaskaran et al, 2009), there are other interactive percussion systems that further focus on interaction with other users as well as the system itself. The Jam-O-Drum was one such collaborative drumming interface, combining a six-person percussion ensemble with generative visual feedback to guide or suggest different types of percussion interaction (Blaine & Perkis, 2000). Similarly, Beatbug developed a collaborative percussion system, albeit simpler than Jam-O-Drum, with users creating and sharing rhythmic expressions through small handheld percussion interfaces (Weinberg, Aimi, & Jennings, 2002). Still within the realm of interacting with another physical performer, robotic systems – though not as focused on the interaction between users – emulate human users by inserting a robotic performer. Systems such as Shimon (Weinberg, Raman, & Mallikarjuna, 2009), Haile (Weinberg & Driscoll, 2007) or MahaDeviBot (Kapur, Singer, Benning, & Tzanetakis, 2007) provide users with a physical being that provides a physical presence to the generated musical responses, as opposed to only speaker-based output. Collaborative systems give the user the sense of belonging in a group or interacting with another percussionist, whereas speaker based setups often seem like output is generated rather than learned or performed.

Using an actual drum or drum-kit for interaction introduces the idea of augmentation, where the focus is more on the interaction, rather than the novelty of the interface (discussed in Section 2.2.4). One can consider the Augmented Djembe and Augmented Snare, both which use real drums with camera tracking systems to augment the user's performance. The Augmented Djembe (as the name suggests) uses a djembe with a camera mounted inside to allow the system to learn and interact with the user's performance (Maki-Patola, Hämäläinen, & Kanerva, 2006). Similarly, the Augmented Snare uses a webcam for reading input data from a snare drum to control audio effects. Although it also offers further augmentation through flex sensors and piezo-electric transducers – embedded in the brushes and drum head respectively – to provide more data for the system to interpret user input with (Gray, Lindesll, Minster, Symonds, & Ng, 2009). Identified by the developers of both the Augmented Djembe and Snare are the issues inherent with digitizing real-time physical input of a user. Using real percussion instruments offers the challenge of having to develop a robust system of evaluating drum strikes, where physical interfaces already exist for digitizing percussion data. The simplest example of such an interface being more recent video game hardware (Rock Band 3), and for more advanced digitization, professional electronic drum-kits which offer MIDI input and output.

2.2.7 Methods of Evaluation

Even with authors offering frameworks for assessing musical interfaces (Section 2.2.2), there is a lack of actual assessment of interface designs. In a 2008 paper for the New Interfaces for Musical Expression (NIME) conference,

Kiefer et al discussed that 37% of the 2007 conference proceedings included formal usability testing in their discussion. This is a valid concern when trying to assess new systems broadly, and is a sentiment echoed by other authors - (The Cambridge Companion to Electronic Music, 2007; Hsu and Sosnick, 2009; Orio and Wanderley, 2002; Jordanous, 2011). Yet, when considering the interdisciplinary nature of the field, a formal usability approach may not be at all appropriate, or valid.

In his “Roadmap for Research”, Serra (2005) suggested a more discipline-agnostic approach to assessment. Recognizing the interdisciplinarity of the field, Serra suggested three criteria: that a work should be understandable in the disciplines that it broaches, that it should affect a balance between them, and it should still be effective in advancing understanding within those disciplines (Serra, 2005). With these in mind, the lack of formal assessment in the field of musical interface design might seem reasonable. The evaluation of a customized system built by a performer (for his or her own use) is less likely to partake in a formal HCI study, as in this case, the performer is the designer. Also given the specializations of those designing the interfaces, they may not feel qualified to perform more formal user testing.

Returning to Serra’s three criteria, his suggestion would be that those who broach the interdisciplinary nature of interface design should be able to advance understanding and present reasonable work in whichever fields their work encompasses, which requires communication on a common ground. Given the predominance of requests for more HCI-grounded usability studies, this

approach warrants being addressed. Most recently, Hsu and Sosnick (2009) brought up the HCI approach to evaluation and identified two high level goals for interactive systems: to provide a usable and musically interesting interaction for the user. This view reflects a larger change in HCI research that is shifting from the efficiency or efficacy of an interaction, to that of the interactions' affect (Waterworth & Fallman, 2005). This is an opinion further echoed by Kaye et al (2007), who posit that the user experience is important in the process of evaluation. Returning to Hsu and Sosnick, introduced in their goals are also those of Collins and d'Escrivan (2007), which reflect a similar – though more specific – approach. Collins proposed that evaluation of a system requires technical criteria, relating to the code, operation or cognitive models employed in a system, the audience reaction, and the resulting interaction. Collins' goals present a more specific, though similar approach to that of Hsu and Sosnick, featuring a requirement for the assessment of the audience or user experience, and of the actual interaction that is occurring.

Continuing with these goals in mind, Dobrian and Koppelman (2006) stated that the best test of a system is having multiple performers try it. A larger group of users offers a broader set of experience, viewpoints, and leads to a higher level of statistical validity. Hsu and Sosnick (2009) posit shorter times in testing may garner interesting interaction, but longer testing can demonstrate the efficacy of an interface to engage users. Both of these suggestions come back to the movement away from solely testing technical capabilities to also looking at

user interactions and experiences as valid means to provide meaningful design insight and direction (Kaye, Boehner, Laaksolahti, & Staahl, 2007).

2.3 Design Goals for Percussion Interaction

Beatback develops on prior design knowledge to explore *intimate* percussion interaction with novice users, while still enabling practiced performers to engage with the system in a rewarding fashion. Drawing from the literature review (1) *augmented functionality*, (2) *stylistic appropriateness*, and (3) *engaged interaction* form the three design goals of the Beatback system.

Augmented functionality refers both to the graphical user interface (GUI) and the operations of the system itself. Minimizing the amount of user interaction with the GUI ensures that their focus is on the interaction with the musical output of the system itself (Section 2.2.4). Similarly, ensuring the system can operate without requiring constant adjustment by the user further promotes this ideal. Pursuing *augmentation* (Section 2.2.4) instead of designing a novel interface also offers the benefit that there are already individuals who are experienced at performing with a drum-kit. The combination of augmentation and a minimal GUI focuses the user's attention on interacting with the musical material instead of the control devices.

Stylistically appropriate (Section 2.2.3) responses are rhythmically consistent with the user's input into the system, and are important in promoting *conversational* and *live interaction* (Section 2.2.5). In other words, while output is stylistically appropriate it should not be always predictable and it should operate

within the user's perceived real-time (Section 2.2.5). Ensuring that the output of the system is stylistically consistent with the user's input helps entice the user into interacting with the system.

The final goal continues this focus on the users' interactions with the system and promoting *engaged interaction*; or, *focused* and *iterative* interaction with the system. This type of interaction is important in users who are developing or exploring their musical abilities (Section 2.2.4), and is therefore of relevance to Beatback's design. Though research with the system does not address this directly, though through testing of various attributes which contribute to engagement (enjoyment, task time), and by enabling iteration (saving and loading sessions) this form of interaction is promoted and assessed.

3: THE BEATBACK SYSTEM

Detailed below is the Beatback system itself, including how it parses and interprets incoming information, generates output, interaction models offered to the user, and how these elements help achieve the goals set out for interactive percussion systems (Section 2.3).

The Beatback system is designed primarily to work with digital drum-kits, but it can work with other MIDI drum interfaces – such as trigger pads – to learn and generate patterns from a user’s performance. The two interaction modes of call-response and accompaniment provide modes through which the user can interact with the system (Section 2.2.3). In the call-response mode, Beatback simply performs when the user is inactive, enabling a back-and-forth conversation with the system. The accompaniment mode has the system fill in the drums that the user is not playing, using a process called zoning: once a drum is struck in a set region of the kit – such as the toms – none of those drums sound until the user stops striking them (Section 3.2.5). This allows for the performance of more complex patterns by having the user layer their performance over the system’s. Research involving Beatback looks at self-directed practice and exploration using both interaction modes (Section 4.4).

3.1 Implementation

Beatback is built in Max 5 (Cycling '74), and works with standardized MIDI interfaces for input, though it has been primarily designed and tested using a Roland V-Drum system. Max allows performance data to be parsed and interpreted by the system, which in turn selects and outputs rhythmic patterns based on the user's input. The following is an introduction to dataflow in Beatback, including references to the Max patches and sub-patches which are available in Appendix C.

Initial input into Beatback is either MIDI data (dealt with by *dataInput*) or keystrokes (dealt with by *input*) from a computer keyboard. The stream of note onsets (when the drum or key is struck) is converted in Beatback into an attribute group of the drum number, inter-onset delay and velocity, which are considered as a note. The drum number is a user-assigned value that ties a MIDI note number or keystroke to a drum within the system, and is used to ensure drums can be reassigned after input.⁵ The inter-onset delay is the temporal distance between two note onsets measured in ticks⁶, and is reliant on the universal (same tempo for all sub-patches) metronome used within Beatback. The metronome tempo is set by the user in *timingControl* and outputs an audible count (click) for them to synchronize their performance to. The count is four quarter notes to a full bar, with the click providing a different tone to allow the user to distinguish when the system is reading a pattern (a half-bar portion). For

⁵ This allows the user to change computer keyboard or drum-kit assignments whenever desired.

⁶ Ticks are a measurement used within Max to denote length within a given meter. Table 1 provides conversion examples.

example, a snare hit which occurred one-quarter note after the prior note would be encoded as Snare-480 (ticks). This stream of attribute groups is fed into the *dataParse* sub-patch which organizes and parses patterns.

Table 1: Musical note lengths as ticks

<i>Musical Value</i>	<i>Value in Ticks</i>
Half note	960 ticks
Quarter note	480 ticks
Eighth note	240 ticks
Sixteenth note	120 ticks
Thirty-second note	60 ticks
Sixty-fourth note	30 ticks

The streams of attributes (each note) are put together into patterns by *dataParse*, and every half-bar they are sent for storage. Each pair of consecutive half-bar patterns are considered as in-out combinations. Taking an example in which six half-bar patterns were played, or A B C D E F, the (1st order) in and out pattern combinations would be A-B, B-C, C-D, D-E, E-F. Each pairing is stored as soon as it is prepared, with the last half-bar completed also being forwarded to *dataQuery* to search for possible transitions.

Pattern storage in *dataStorage* first looks for matches in the existing table, and given no match, creates 1st, 2nd, and 4th order Markov entries for the input-output pairings (further detailed in Section 3.2.2). The 1st order pairings are as listed above, A-B, B-C, etc.; which consider one input half-bar portion to use for a

continuation. 2nd order pairings consider two half-bar portions (one whole bar), or AB-C, BC-D, and so on.

The last in-pattern played by the user is also used for queries in *dataQuery*. So starting with the 4th order (and the prior example), ABCD would be the first search. Through a process further detailed in Section 3.2.4, the system searches through each of the orders from highest to lowest for in-pattern matches, which are then used to select output patterns out to *dataPlay* for playback.

The *dataPlay* sub-patch ties together *drumZoning* (Section 3.2.5) and *dataInput* to listen for when the user is playing, and given the interaction mode – accompaniment or call-response – output the patterns. In accompaniment, the setup of *drumZoning* filters the system’s drum output accordingly, while in call-response, *dataInput* informs *dataPlay* when the user is inactive. Given this general overview of the implementation of Beatback, the next section focuses on the specific ways that Beatback parses, stores and queries the user’s performance data.

3.2 Data Parsing

Beatback enables different models of self-directed percussion interaction, and requires the ability to generate stylistically appropriate musical material for this interaction. Taking the goals outlined for designing percussion interaction (Section 2.3), the system should be able to generate material stylistically consistent with the input, and offer enticing perceptibly real-time (Section 2.2.5)

interaction for the user. The combination at Beatback's core of MIDI and Max 5 enables the use of a variety of samplers or synthesizers in performance and generation along with approaching perceptual real-time latency without having the additional step of having to process live audio. As discussed in the prior section, tempo in Beatback is a fixed value set by the user before engaging with the system. This simply means that before or while engaging with the system, the tempo can be set by the user, but the value remains static for the duration of their interaction. Parsing and output of patterns is based on the set tempo that allows material learnt by the system at one tempo to be easily shifted to another.

3.2.1 Parsing MIDI Data

As the user plays, the three attributes of the drum type, inter-onset delay and velocity are collected as a stream of sequential input data and each note played by the user is a combination of these three attributes when stored (Section 3.1). Inter-onset delays are assigned in ticks (see Table 1 for musical equivalents) which Beatback is able to measure based on the user-set tempo. As a result, the first value in a recorded pattern receives a value of zero, as no note has come before it. All proceeding note values in the recorded pattern are given a length attribute that reflects their temporal distance from the prior note. Given a quarter and a half note are played, the first note in the pattern would have a duration of zero (no note before it), and the half note would receive a quarter-note inter-onset delay of 480 ticks as it is the temporal distance from the prior note. Similarly, when two notes that sound together, one is assigned the inter-onset delay (from the prior) and the other a value of zero. Given a pattern of

quarter-note snare then crash and kick at the same time, the stored pattern would show a 480 ticks crash with an accompanying 0 tick kick. There is no quantization of the pattern at this stage, as the tick value recorded by the system will be exactly as performed by the user. Quantization is applied during the pattern storage stage (Section 3.2.3).

3.2.2 Pattern Learning

Upon the completion of every half-bar after the user starts, the system will take the patterns input by the user and parse them for storage. In the input and learning of patterns, Beatback uses Variable-Order Markov Models (VOMM) that are probabilistic models in which the state of the pattern is described by a single discrete variable whose possible continuations describe all the learnt states (Russell & Norvig, 2003). As an example, a common system that functions similarly to how Markov models do is the T9 predictive text system available on most cell phones: as you start to input letters, the system begins narrowing down the possibilities and provides you with the most probable complete word for your set of letters (Dunlop & Crossan, 2000). In the case of Beatback, the complete set being described is the body of patterns being learnt, with each pattern being a set of notes with the aforementioned attributes (drum type, note distance, velocity). Patterns are stored in paired input and output tables, with a transition table that lists the number of times each input pattern has transitioned into an output pattern. The process of storage is explained in detail in Section 3.2.3, but as a quick example, the pattern Snare-1/2 note, Kick-1/2 would become an input pattern of Snare, an output pattern of Kick, with a single transition listed between

the two. This method of storage enables the query stage (Section 3.2.4) to quickly lookup possible output given an input pattern.

An important distinction between different types of Markov models is the order of the model, which defines how many previous states are considered when computing the probability of future states. VOMMs are capable of storing information on patterns of varying lengths, which is important musically, as the model considers and searches for longer patterns as well as shorter ones (when no match exists for the longer one). Fixed-order Markov models can only consider a fixed length of input pattern (Ching & Ng, 2006). For example, if given a transition of A to B to C to D, or for this paper's purposes A,B,C,D, a first order Markov model will consider on one prior order: given A, B would result, or given B, C is the only existing continuation, and so forth. A second order Markov model, given A,B would provide C, or two orders (notes) leading to the next value. Beatback draws on the benefit of variable-order Markov models and works with 4th order (four half-bars; two bars), 2nd order (two half-bars; one bar), 1st order (one half-bar) searches. Allowing it to work from larger two bar patterns down to half-bar searches as required (detailed in Section 3.2.4).

To help clarify the benefits of higher orders, consider that the pattern A,B,C,B,A has occurred. In a first-order search, a search for B can equiprobably result in the continuation of C or A. Given one step up – a second-order – search for A,B, the only result is C. Enabling higher orders of Markov models allows for more musically appropriate continuations of learnt patterns based on user input.

3.2.3 Pattern Storage

To exemplify how pattern storage works, presume that the user has input some notes, for which the attributes (drum-distance-velocity) are collected. For the sake of example, only first-order Markov models are presented.

Example 1. The user has played Kick-0-60, Snare-1/2-80, Snare-1/8-60, Hihat-1/8-50, Kick-1/4-68, Snare-0-86.

As Beatback continuously portions patterns into half-bar sections (Section 3.1), the half-bars read by Beatback would be {K-0-60} (A), {S-0-80, S-240-60, H-240-50} (B), {K-0-68, S-0-86} (C). Recall that Beatback breaks the pattern into half-bar portions, which is why {K-0} is alone, and {S-1/2} from the original pattern becomes {S-0} as it is the first note in the half-bar portion. Each in-out pairing of start to end of the pattern would be placed in a temporary table along with the number of times each pairing has transitioned; B given A and C given B (Figure 1).




Figure 1: The first pattern as read by Beatback and then parsed into in-out pairings.

Once done encoding the input, the patterns are compared against the existing corpus for updating. The broken down input pattern's pairings – in and

out pattern combinations – are checked against existing transitions for matches. Looking for existing patterns compares each of the pattern's three attributes (drum, distance, velocity) to ensure a match. Each attribute has a different method for matching, so the process goes as follows: first checked is the length of the pattern (number of notes). For checking drum types, the system cycles through both patterns to ensure that the drum type for each note in the patterns are an exact match. If matched, the note lengths of both patterns are checked to ensure that they are sufficiently close given the set quantization value. The quantization determines the variance in note length that is acceptable for finding a match, and is an assignable length value – half, quarter, eighth, sixteenth note. So returning to the example, a sufficiently close pair of notes would have a difference in lengths less than the quantization value. With an eighth note quantization (240 ticks or +/- 120 ticks), any difference between two notes which is less than a sixteenth – e.g. a note measuring 235 ticks and one at 142 ticks, a 93 tick difference (and less than 120 ticks) would be considered sufficiently close for a match. The system checks all notes within a pattern to determine a match. Unless a match is found, this process repeats until all the stored patterns are exhausted, in which case the system adds the new pattern. Given a matching pattern, an updating process occurs to the stored pattern, in which note lengths and velocities for notes are averaged with the existing pattern. All averaging of note lengths and velocities only averages the existing value (an average) with the new input's note lengths and velocities.

One final important element to the storage process is with regards to notes that sound together, and how they are represented within the pattern. When parsing input, Beatback ensures that when two or more drums sound within the quantization value of each other, any drum assigned a zero length come after the note that retains the inter-onset distance from the prior. This occurs to ensure the initial structure of the input pattern is maintained.

Example 2. The user plays Kick-0-70, Snare-1/2-48, Hihat-1/4-35, Kick-1/2-85.

The pattern is read by the system to be {K-0-70}, {S-0-48, H-480-35} and {K-480-85}. Starting with {K-0-70}, the system would find a match both in length and drum-type results of the first 'in' value in the table being updated (Figure 1). As the velocity of the original table value (60) and the new value (70) do not match, the new value would be 65, an average of the two (Figure 2). All values for the velocity are not stored and instead each becomes part of the new average. Moving on to the second section {S,H}, there is one output match in length {K,S}, but since it does not match in drum-type, a new output pattern would be added, and transitions updated accordingly for {K} to {S,H}. Since no patterns matching in length to {S,H} exist for input the new input pattern would be added, and the system moves on to {K-480-85}, which is also added to the pattern table. Note that even though drum-type patterns may match, note lengths not within the quantization value require Beatback to add new patterns to the table. Note lengths within quantized values are averaged to create the new stored value. Having established a pattern table to work with, albeit small,

Beatback can now continue and use the parsed and stored patterns to begin offering generating patterns back to the user.




Figure 2: The new pattern as read by Beatback is added to the table, updating the values.

3.2.4 Pattern Generation

While storing patterns, Beatback is also constantly querying for possible transitions based on the last half-bar portion. The only point at which Beatback is not generating patterns is when the user first starts the system, as it will not yet have learnt any material. Once there is material though, Beatback will generate patterns in two stages: query and build.

In the query stage, the last half-bar is used to query for possible patterns to transition into. Every half-bar, the system will take the given pattern and query the input patterns for a match, which ensures that the system is constantly prepared to begin generating output. For the variable orders of the model being

used, the system will start the query with four half-bar portions (2 bars – 4th order), and failing a match move onto two half-bar portions (1 bar – 2nd order), and then a single half-bar portion (1st order). For example, given only two half-bar portions of query input, the system will only be able to start with 2nd order searches.




Figure 3: The user has played some more, resulting in some additional transitions. Highlighted are possible transition possibilities for {S, H}

Example 3. The user plays Snare-0-57, Snare-1/2-84, Hihat-1/4-30. Which is queried for continuations.

The half-bar portions read by the system are {S-0-57} and {S-0-84, H-480-30}. The system would start with a second-order search, looking for {S-0} - {S-0, H-480}. Since {S-0} - {S-0, H-480} has not occurred, doesn't exist (Figure 3), the system moves down to a first order search of {S-0, H-480}. As a match exists for {S-0, H-480} (see Figure 3) the number of transitions are used to weigh the

possible options, {S-0, H-480} is 33% (one transition) and {K-480} is 66% (two transitions) probable as the continuation for the pattern. Both options and their percentage weights are then taken into consideration during the build stage.

As described above, the build stage similarly uses the weights of possible transitions identified in the query stage and begins to build a pattern. Continuing from the earlier example of {S, H}, the system would stochastically select from the weighed options, and outputs the resulting pattern. Let us assume that given {S, H} the less probable transition of {S-0, H-480} was chosen, which is output (audibly) by the system, and is used as a new query for continuing. Beatback then returns this output to the query stage only to discover that the new query pattern – {S-0, H-480} – is available in the input patterns and continues building the pattern with. Higher orders work similarly, though the system searches for chains of half-bar patterns to match.

3.2.5 Drum-kit Zoning

Perhaps the most significant yet simplest addition to Beatback is drum zoning. Zoning allows user-assigned groupings of drums (As described above, the build stage similarly uses the weights of possible transitions identified in the query stage and begins to build a pattern. Continuing from the earlier example of {S, H}, the system would stochastically select from the weighed options, and outputs the resulting pattern. Let us assume that given {S, H} the less probable transition of {S-0, H-480} was chosen, which is output (audibly) by the system, and is used as a new query for continuing. Beatback then returns this output to the query stage only to discover that the new query pattern – {S-0, H-480} – is

available in the input patterns and continues building the pattern with. Higher orders work similarly, though the system searches for chains of half-bar patterns to match.) to be filtered out of the system's playback when the user strikes a drum within that zone: when the user is playing any of the toms, the system will not play any of the toms. While Beatback by default uses the drum-zoning model shown in As described above, the build stage similarly uses the weights of possible transitions identified in the query stage and begins to build a pattern. Continuing from the earlier example of {S, H}, the system would stochastically select from the weighed options, and outputs the resulting pattern. Let us assume that given {S, H} the less probable transition of {S-0, H-480} was chosen, which is output (audibly) by the system, and is used as a new query for continuing. Beatback then returns this output to the query stage only to discover that the new query pattern – {S-0, H-480} – is available in the input patterns and continues building the pattern with. Higher orders work similarly, though the system searches for chains of half-bar patterns to match. (described below), the user is also able to assign their own zoning.

Beatback is set up to work with a digital drum-kit, so the default model used for zoning it is as follows: the set of cymbals (ride and crash), the toms, the snare, the kick and the hi-hat are all assigned separate zones. As long as the user does not play within a zone, any of the drums within that zone can be played by the system when in either the practice or accompaniment modes. When the user strikes a drum within a zone – for example the low tom – none of

the drums in that zone – the high, mid and low toms – will sound until the user has not played within that zone for a period of two beats (half a bar).

The one important contingent for playing an accompaniment in a given zone is that the user has to have played patterns within that zone before. To maintain the rhythmic composition of the user's originally inputted material, Beatback does not take patterns learned with one drum and associate them with another. Therefore to have Beatback play the hi-hat while the user plays the snare, the user first has to play the hi-hat. Then, once the system has learnt how to play the hi-hat, the user can play along with the snare.

Drum zoning supports the user when they are playing or learning new patterns. When using drum-kit zoning in a practice context, users can load a pattern (Section 3.2.7), and then learn the pattern drum by drum – having the system play the missing parts – until they are comfortable with the entire pattern. Zoning also provides auditory cues in interaction: Beatback can be continuously cueing users with their own musical material on the non-engaged drum zones.




Figure 4: Example of the default drum-kit zoning used in Beatback.

Interaction where the user hears their own musical material can provide audible suggestions for patterns to play or explore.

3.2.6 User Interaction

There are two main interaction modes in Beatback: call-response and accompaniment, which as their names suggest, respond to or accompany inputted patterns. In either mode, the system begins output after two beats (or half a bar) of user inactivity, a length referred to as the output delay. Within the call-response mode, once the user has stopped playing and the output delay has transpired, the system can output patterns continuously until the user begins playing the drums again. Any number of drums struck by the user (the call), silences the system immediately until the user stops playing and the system begins generating output again (the response). Within call-response, all data is stored into one transition table for all the drums; no zoning (Section 3.2.5) is applied.

The output of the system in the accompaniment mode is the similar to that of the call-response mode, but the output – coupled with the drum-zoning feature (Section 3.2.5) – filters out the drums that the user is playing. Different drums of the kit are grouped together (Figure 4, pg.44), so that should one drum be played by the user in that zone, the entire set is silenced. For example, in the accompaniment mode, should the user play a pattern on the snare, the system will continue that pattern for the user until they play a new pattern on the snare. This also marks the difference in how the two modes store data, with accompaniment learning and storing a transition table for each zone separately.

Unlike call-response mode which stores all patterns in one table, accompaniment will have a table for the toms, another for the snare, and so on. As a result, the snare part the system learnt will play while the user tries another pattern on another zone of the drums, such as the toms. The accompaniment mode results in a form of intelligent live loop machine for the drums; allowing the user to build up patterns they may not be able to otherwise due to limitations in their physical skill level or musical competency.

One final distinction between the two modes is the ability of the two to vary patterns. Within the call-response mode, the system is learning patterns for the entire drum-kit at once, whereas with the accompaniment mode, each drum zone is learnt separately from one another. This means that within the accompaniment mode, there can be a higher variability in the patterns produced by the system, as each drum will generate patterns agnostic of what the other drums might be playing.

3.2.7 Pattern Loading, Storage & Training

In addition to the ability to learn patterns from user input, Beatback also has the ability to play along with prior performances with the system through saving and loading prior sessions. To enable this functionality, Beatback encodes and saves the table data (stored patterns) into a set of three files – the in patterns, out patterns, and transitions – which can then be later reloaded, or shared with others. Additionally, the system allows users to disable learning to play with already learnt material.

3.2.8 The GUI

The graphic user interface for Beatback allows for the user to tailor the system to their needs, though it is not intended for adjustment in a live setting. The GUI control has been kept to a minimum to ensure the user is focused on the interaction instead of adjusting settings as identified in Section 2.3. Rather than expending their attention dealing with adjustment of the system, they will focus on the physical action and auditory feedback of the system.

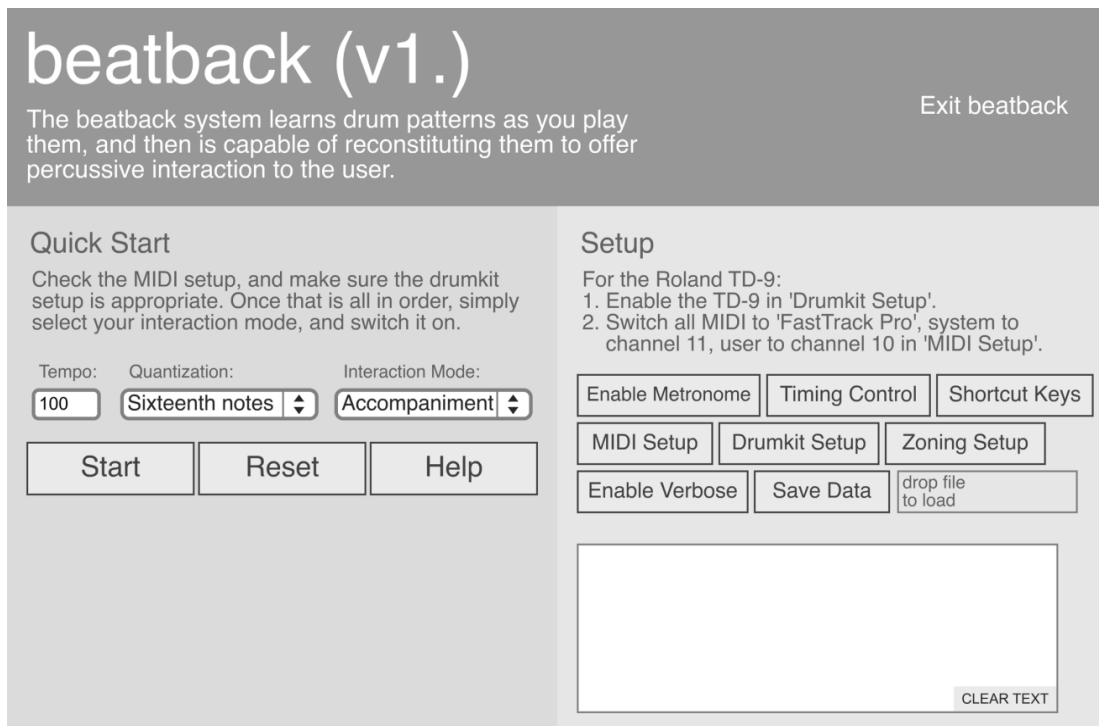


Figure 5: The Beatback user interface allows for control of tempo, interaction modes and quantization. Additionally there are more in depth setup options and an ability to load and save data.

4: BEATBACK RESEARCH

Research involving the Beatback system follows two directions: (1) testing how user's experience learning rhythmic patterns differs between a zoned and non-zoned drum-kit, and (2) whether Beatback can promote rhythmic exploration. This chapter details the design of these studies, their procedure, results and analysis.

4.1 Interactive Percussion in Rhythmic Exploration

As explored by Pachet and Addessi (2004), interactive musical systems that are capable of a conversational interaction style have shown benefit to users' focus, length of interaction, and exploration of different musical motifs. These benefits are the focus of the empirical study of Beatback; exploring whether an interactive drum-kit can assist self-directed practice and rhythmic exploration. This involves the validation of the Beatback system as an interactive percussion system and looking at user interactions with the system. Two questions come out of validation, which are (1) how to validate the system's design while researching its abilities, and (2) what are the possible benefits offered by rhythmic exploration.

4.1.1 Design Validation as Research

Within the existing literature on musical interface design, approach to research has been progressing towards HCI models, such as in the work of Orio

et al (2002), where they looked at borrowing tools from HCI research to enable a standardized means of assessment. This theme of attempting to develop a common grounding for testing musical systems or interfaces repeats frequently in literature, with authors offering varying generalized goals for design of musical interaction (Sections 2.2.2 & 2.2.7). Given all their input on how to assess these systems, two themes are prevalent: (1) interaction and the (2) user's experiences are important in the development of an interactive musical system.

Working with these themes, there is no reason that the validation of a system through user testing cannot also offer results that provide insight into how to improve or further musical interaction design. Although, as demonstrated (Section 2.2.7), it is important that any validating research be presented in a fashion appropriate to the fields which they broach. Research presented appropriately within its context ensures it will be of benefit to its respective field.

4.1.2 Benefits of Rhythm Exploration

Augmentation of standardized interfaces offers a practiced performer the ability to translate their existing abilities on an existing interface to one with heightened functionality. Similarly, it can offer novice users an interactive outlet with which to start exploring an instrument. As was discussed (Section 2.2.4), employing an augmented rather than novel interface benefits the performer as there is no requirement to relearn basic motor or muscle-memory skills. Instead, the focus of playing with such a system becomes the nature of the interaction itself, and of specific interest to Beatback – rhythmic exploration.

In a paper discussing the Western music learning model, Davidson and Jordan (2007) discuss how free exploration (or play) with a musical instrument is integral to the development of musical abilities. Similarly, Davidson et al (1997) discuss how children who do not engage in informal learning activities, frequently give up playing altogether. For these reasons, assessing rhythmic exploration with Beatback provides insight into a method for promoting exploration of musical expression and the development of musical interest.

4.2 Research Design

The drum-kit zoning research focuses on looking at how the concept of drum-kit zoning alters practiced percussionists' – individuals with more than four years of experience and, or, two years of formal training – motivation or process in learning new rhythmic patterns. It is important to note that the drum-kit zoning research does not use the accompaniment mode of Beatback itself, but only explores the filtering that zoning offers (the experiment involves no learning by the system). The exploration research acts to validate the Beatback system, and to explore if interactive percussion can heighten users' exploration of rhythmic patterns. Each study is presented in detail on its own (Sections 4.3 and 4.4) with a larger discussion combining the results of the studies (Section 5).

4.3 Drum-kit Zoning Research

Drum-kit zoning – as detailed in Section 3.2.5 – is a simple process by which drums being played by the user are filtered out of a pattern being played back by a computer. It is a functionality included in Beatback, but which is looked

at separately here as it presents a model for interacting with a digital drum-kit not previously explored by research, and may be beneficial as a tool for interactive sonification (Section 2.2.4). Zoning can fill in the drums the user is not playing and thereby offers a model for interaction where the user can practice patterns part by part, with the parts of the pattern they are unable or unwilling to play, continuing in the background.

This research looks at how practiced percussionists experience learning a given drum pattern using zoning as opposed to an un-zoned drum-kit. During two tasks, one with zoning and one without, the participants' performance is captured as MIDI data and then analysed using the MIDI toolkit for Matlab (Eerola & Toiviainen, MIDI Toolbox: MATLAB Tools for Music Research., 2004). The toolkit enables an analysis of the recorded MIDI data. To provide further insight into these results, an Intrinsic Motivation Inventory (Section 4.3.4) and any commentary by participants post-task are also collected.

Prior to continuing an explanation of the study, self-directed learning needs definition as this study attempts to approximate such a learning environment. Self-directed learning is a model of learning where the student takes charge of their motivation in learning and developing a skill on their own (Knowles, 1975). Since practising the drum-kit is often a self-directed endeavour, the drum-kit zoning research environment seeks to replicate this solitary learning model to understand if zoning can benefit this learning method.

4.3.1 Methods

Looking at practiced percussionists, a balanced within-subject design is employed to ensure a random though even assignment of tasks to participants. The balanced-random design is used in addition to a complex rhythmic pattern (for practising in-task) to minimize learning effects in the study results. The study uses a questionnaire titled the Intrinsic Motivation Inventory (IMI – Section 4.3.4) which gives quantitative data on participants' self-reports of enjoyment, tension and competence. This data is compared with their actual MIDI performance while learning a given drum pattern in two tasks: (a) one task with the entire pattern playing continuously, and another (b) with the pattern being zoned, so that when the participant strikes a drum, that drum does not sound in the pattern. The balanced design ensures that both moving from un-zoned to zoned and vice versa are assessed by the study, along with the activities within task. All participation was voluntary, and participants were drawn from a pool of undergraduate and graduate students at Simon Fraser University, and drum instructors from Vancouver-based music schools.

4.3.2 Equipment and Setup

The equipment used for this study includes an electronic Roland TD-9KX drum-kit, which has skins on its drums to give as realistic a feel as possible to the participant, without using an acoustic drum-kit. Acoustic kits lack a built-in ability to provide MIDI data, which is what enables the capturing of each task. The drum-kit was patched via MIDI into a computer running Apple's Logic and Max 5. Logic was used to play back the drum pattern to Max 5, which returned the

zoned or un-zoned output to the drum-kit for the playback of drum sounds. During the task, MIDI data from the participant's performance was recorded by the computer for later analysis. To ensure that participants did not confuse their own performance with that of the original pattern being played – due to masking effects (Neuhoff, Kramer, & Wayand, 2002) – a stereo headphone setup was employed. The right earpiece output the sound of the participant's own performance, while the left earpiece output the pattern played by the system.

4.3.3 Study Procedure

At the onset of the study participants were greeted, and handed a consent form to fill out. They were also verbally informed that they could leave the study at any stage, and that all results would remain anonymous. This portion of the research had sixteen participants – fourteen male, two female – all students at Simon Fraser University or Vancouver-based percussion instructors.

Given consent, the study proceeded with the participant being given time to set themselves up at the drum-kit by adjusting the chair or drums to their liking. Once the participant was comfortable at the drum-kit, the eight bar pattern to be learnt (see Figure 6) in the tasks was played twice: the first time to ensure that volume levels while performing were appropriate, and the second time to ensure that all participants had a common knowledge of the pattern before beginning the first task. Once the participant was prepared to start, they were reminded to learn as much of the pattern as possible for the entire time the pattern was playing (ten repetitions). As it is not necessary to have them learn the entire pattern when trying to assess their process, the pattern played was deliberately complex.

Given a random (though balanced, see Section 4.3.1) first task assignment, the participant would either be learning the pattern using drum-kit zoning, or learning the pattern with it playing constantly in the background. In each task assignment, the pattern was repeated ten times, during which their performance was recorded using MIDI data. After completing the first task, participants received an Intrinsic Motivation Inventory (see Section 4.3.4) to fill out. Having completed the questionnaire for the first task, the participant was assigned the task they had not received the first time – zoned or not zoned – and repeated the process (including the IMI questionnaire). Once finished, the participant was thanked for their time, and they were asked if they had any comments or questions on their participation in the study. The researcher noted comments relevant to the participant’s experience during the study.




Figure 6: Eight-bar rhythmic pattern given to participants.

4.3.4 Data Collection

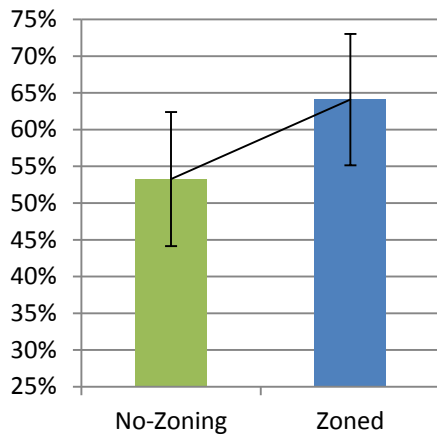
The first tool used to collect data in the study is the Intrinsic Motivation Inventory (IMI), a questionnaire for gaining quantitative data on participant's perceived intrinsic motivation (University of Rochester, 2010). The inventory is a statistically validated questionnaire (Tsigilis & Theodosiu, 2003) with four editions including nine, twenty-two, twenty-five or twenty-nine questions. Each version of the questionnaire assesses different motivations such as enjoyment, tension and competence. The IMI was chosen as a tool for research with Beatback as it has been used repeatedly in learning or classroom applications to explore student experience (Ryan, Connell, & Plant, 1990; Deci, Koestner, & Ryan, 2001; Ryan & Deci, 2009). As the questionnaire is designed for delivery on paper, and with participants responding themselves, each of the assessed motivations is a self-reflection by the participant. For example, their responses on the questionnaire that relate to enjoyment demonstrate their self-perceived enjoyment during the given task. For this study, the nine-question version of the IMI was used, as the longer questionnaires include details that were not the focus of Beatback's research. All nine questions were rated by participants on a Likert Scale of one (not at all experienced), to seven (totally experienced). Appendix A shows the questionnaire as presented to participants and Table 2 lists the questions and what perception each is meant to assess. The IMI is used after each task to assess the difference in motivations between the two tasks. Phrases such as "...compared to other students," force participants to self-assess their experience against what they perceive as a norm. The questionnaire provides responses relating to participants' perceived enjoyment, tension and

competence, which can affect their motivation in performing a task. Positive levels of enjoyment demonstrate the participant is being satisfied, low levels of tension show a minimal amount of stress and high levels of competence can indicate an understanding or knowledge in the given task. In addition to assessing these three motivations, the IMI is also employed to enable comparison of perceived against performed competence, which is inferred from the recorded MIDI data. Matlab is used for the MIDI data analysis, with each eight-bar phrase considered as one block in the learning of the pattern. With the participant repeating each phrase ten times during their task, the result is ten blocks in which the performance of the participant can be analysed.

Table 2: Intrinsic Motivation Inventory questions with terminology used for the study, and what each question assesses.

Question	Assesses
While I was practicing the drum-kit, I was thinking about how much I enjoyed it.	Enjoyment
I did not feel at all nervous while practicing.	Tension (Reversed)
This practice method did not hold my attention at all.	Enjoyment (Reversed)
I think I practiced the drum-kit pretty well.	Competence
I would describe this practice method as very interesting.	Enjoyment
I think I understood this practice method very well, compared to other students.	Competence
I enjoyed practicing the drum-kit very much.	Enjoyment
I felt very tense while practicing the drum-kit.	Tension
The drum-kit was fun to practice.	Enjoyment

4.3.5 Results & Discussion



Enjoyment:


$t(15) = -2.78, p < 0.05$

Standard deviation:

No zoning = 18%

Zoning = 18%

Figure 7: Participant's perceived enjoyment (higher percentage, more enjoyable)



Tension:


$t(15) = 2.21, p < 0.05$

Standard deviation:

No zoning = 15%

Zoning = 15%

Figure 8: Participant's perceived tension (higher percentage, more tense)



Competence:

$t(15) = -2.24, p < 0.05$

Standard deviation:

No zoning = 16%

Zoning = 17%

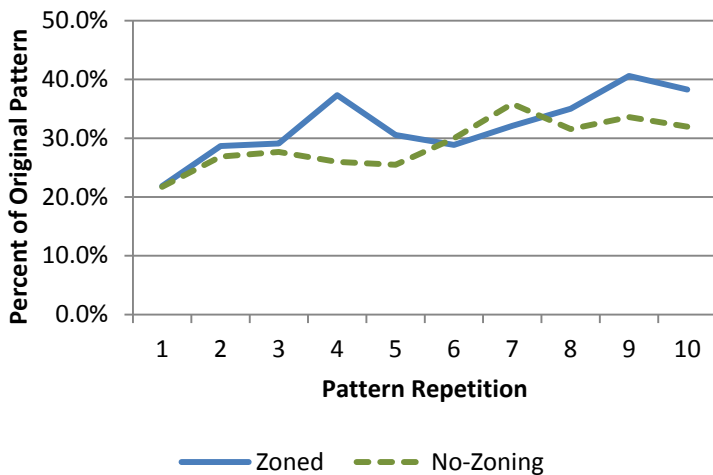
Figure 9: Participant's perceived competence (higher percentage, felt more competent)

Looking at the performance results (MIDI data), note-density and timing accuracy of participants were compared. Note-density was measured to provide insight into how much of the pattern the participants were practising, or a measure of their performance confidence (Figures 10-11). Note accuracy measured within a 32nd note is used to assess the timing accuracy of the participant. For both tasks, the results were analysed using a pair-wise repeated measures ANOVA, and the comparison of tasks for note-density shows a significant difference $F(1,16) = 5.30$, $p < 0.05$, with the zoned condition showing better participant performance ($M=0.31$, $SD=0.13$) over the non-zoned condition ($M=0.29$, $SD=0.15$). This significance indicates that when primed with the suggestion of building up patterns one drum at a time – as they become comfortable – the participants performance reflects a higher level of comfort with the zoned task as opposed to the non-zoned task. Timing accuracy demonstrates no significance towards either task $F(1,16) = 1.15$, $p=0.28$ which clarifies concerns with an earlier bias: as discussed, the masking effects (Section 2.2.4) were expected to provide the zoned condition with better timing accuracy results than in the non-zoned condition. The lack of significance suggests that the conditions of the experiment may have inherently reduced masking issues through its use of one earpiece for the pattern, and the other for the participant's performance (Section 4.2.2). Although, it is difficult to conclude on the concerns with masking considering the small subject size.

Note density is the percentage of the original pattern played by the participant. For example, if the pattern has two drums, and the participant plays

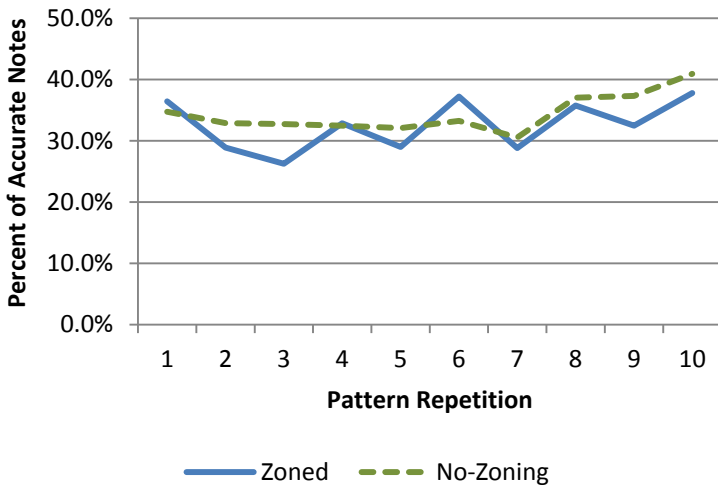
one, they would achieve a 50% density. Timing accuracy was calculated within a 1/32nd note of the original pattern and only for notes played by the participant. The graphs below show the average performance of participants between the two tasks over all ten repetitions of the pattern (Figure 7).

As presented above, zoning while practising a drum pattern helps participants perceive themselves as more confident, less tense and enjoying the task of learning a drum pattern significantly more than when not applied. One reason it may be preferred over trying to match a constantly playing pattern can be found in Figures 12-13; which shows how participants' timing accuracy and note density (comfort) changed over the duration of the tasks. In the zoned condition, there is a visibly stronger relationship between the note density and accuracy: frequently, when one value increases, the other decreases, such as in repetition one or nine. Conversely, in the non-zoned condition, both run closer to in parallel throughout the task. Perhaps what zoning offers is a hands-free method for focusing attention. As a participant focuses on practising a particular drum in the rhythm, they might lose track of the filtered out pattern (as it is not playing), but improve their performance by exploring the pattern as they need to. Then, when they need a reminder, they can listen to the pattern again and then repeat, to improve their performance. The relationship shown in Figures 12-13 is not perfectly consistent as repetition three and four in the zoned task see both the note density and accuracy increasing and decreasing at the same time. In either case, more research as proposed below, is required to further understand the benefits of zoning.



Two-way ANOVA,
repeated measures.
Between participants:
 $F(15,16) = 8.96, p < 0.05$
Between tasks:
 $F(1,16) = 5.30, p < 0.05$
Interaction:
 $F(15,16) = 3.82, p < 0.05$

Figure 10: Note density as percentage of original pattern (between tasks)



Between participants:
 $F(15,16) = 6.41, p < 0.05$
Between tasks:
 $F(1,16) = 1.15, p = 0.28$
Interaction:
 $F(15,16) = 0.55, p = 0.89$

Figure 11: Percent of notes played accurately (between tasks)

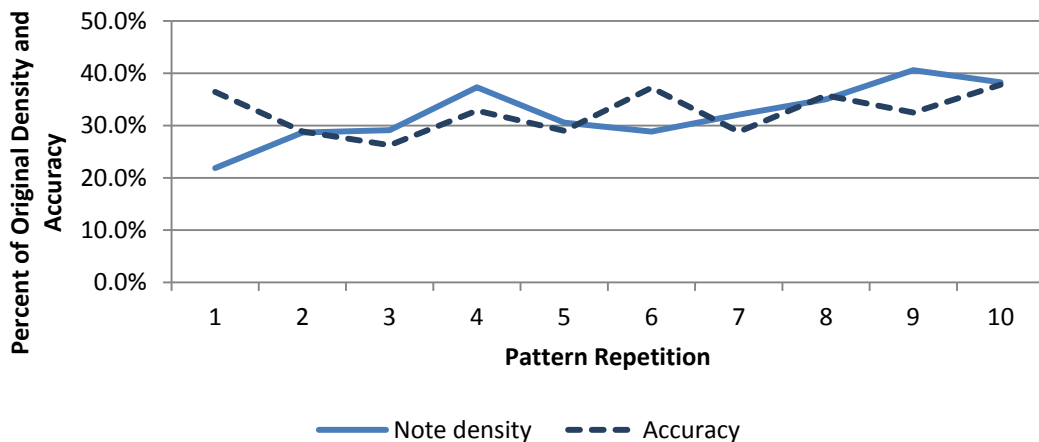


Figure 12: Comparison of note accuracy and density within the zoned task




Figure 13: Comparison of note accuracy and density within the non-zoned task

During the study the participants themselves did bring up the concern that the pattern used was far too complex to be learnt in the given task time. This is a legitimate concern, as should drum-kit zoning be considered for use as a practice or learning tool, more human-attainable patterns should be assessed. The intent with the study design was to ensure that participants did not master the pattern within one task time, and avoid any learning effects that may hinder the quality of results. Similarly, while the research looks at participants' timing accuracy and the density of the patterns they play, it doesn't assess the participants post-task. In particular, was the participant able to replicate the learnt pattern better post-task? More research into zoning focusing on post and pre-task quality of pattern performance would help determine under which precise conditions it is beneficial in practice or learning the drums. Additionally, a longitudinal look at participant's performance in learning a pattern using zoning for a prolonged period may help garner insight into reasons for user preference.

4.4 Beatback Interaction Research

While the drum-kit zoning research looks specifically at zoning, only one element of the Beatback system, the interaction research looks at the system as a whole and its application in fostering rhythmic exploration. The goal is to determine whether the system achieves its goals as an interactive reflexive percussion system, and if it enables participants to explore rhythmic expression. The interest in rhythmic expression specifically is to further research done on interactive reflexive music systems (IRMS, see Section 2.2.3) and their application to musical exploration. Should individuals demonstrate high levels of engagement and exploration with an interactive drum-kit, it may demonstrate that IRMS for percussion could garner similar benefits in self-directed learning, as initially explored with the Continuator (Addressi, 2007).

4.4.1 Methods

Employing a mixed methods design, the study focuses on participants' enjoyment, tension, and competence (IMI); their actual performance data (MIDI); and a short survey of their prior experiences with drums and music. There is one task in the study repeated with three different conditions; one condition without any system interaction, one using Beatback's call-and-response mode, and one using the accompaniment mode (Section 3.2.4). The tasks and conditions are detailed in Section 4.4.3.

Twelve participants – seven male and five female – averaging 30 years of age and approximately 18 years of their life playing or learning various musical instruments completed the three tasks. Within the twelve participants, nine were

concerned naïve percussionists, and three were considered practised (defined in Section 4.3.1). All participation was voluntary, and participants were invited to join the study via an email circulated through the Simon Fraser University, School of Interactive Arts and Technology graduate student community mailing list.

4.4.2 Equipment and Setup

The equipment and setup is the same as the zoning study (Section 4.2.2) exempting that Max 5 is used to record the MIDI data in lieu of Logic.

4.4.3 Study Procedure

Once arrived, the participant was greeted and it was requested that they fill out a consent form. While completing the form, they were notified of their ability to leave the study at any time and that all results will remain anonymous.

Table 3: Intrinsic Motivation Inventory questions with terminology used for the study, and what each question assesses (questions altered from Table 2 to clarify intent)

Question	Assesses
While I was playing the drum-kit, I was thinking about how much I enjoyed it.	Enjoyment
I did not feel at all nervous while playing.	Tension (Reversed)
This method of playing did not hold my attention at all.	Enjoyment (Reversed)
I think I played the drum-kit pretty well.	Competence
I would describe this method of playing as very interesting.	Enjoyment
I think I understood this method of playing very well, compared to other students.	Competence
I enjoyed playing the drum-kit very much.	Enjoyment
I felt very tense while playing the drum-kit.	Tension
The drum-kit was fun to play.	Enjoyment

Given consent, the study then began with an introduction to the drum-kit for naïve participants, or open time to set up the drum-kit for experienced participants. Participants were then introduced to the task with a randomly (though balanced; Section 4.3.1) assigned condition. Each participant completed the same task three times with each of the conditions, and received them in a different order to remove biases or learning effects that might have occurred from their order. Participants had from one to ten minutes for each task and that were free to end the task at any point if they felt a loss of interest. If the ten minutes simply expired, the task concluded. The task itself was to play or interact with the drum-kit in one of the three conditions: (1) no response by the system; meant as a control condition, as it allows for comparison against simply playing with a regular drum-kit. The (2) call-response condition worked with the Beatback system enabled, and the participants were offered the additional options of being able to reset the system (clear learnt patterns) and enable or disable the metronome. The (3) accompaniment mode had participants interact with the Beatback System, and also offered a reset and metronome toggle.

Table 4: Prior musical experience questions for exploration research participants.

Question	Context
How old are you?	Mean age
What is your first language?	Culture bias
Have you ever played any musical instruments? If so, please list them and your experience. (ie. Piano, 4 years with classes)	Musical motor-skills
Have you ever received formal musical training? If so, please list them and your experience. (ie. Vocal training, 2 years)	Musical learning
Have you ever played any musical video games? If so, please list them and your experience. (ie. Rockband, guitar, 2hours)	Musical games

Upon completion of each condition – either at a time decided by the participant or after the ten minutes had elapsed – the participants were given an Intrinsic Motivation Questionnaire to fill out (Table 1). At the end of the experiment, participants were given a short set of survey questions on their prior experience with musical interfaces – shown in Table 4 – to provide further context to the study’s results.

4.4.4 Data Collection

The Intrinsic Motivation Inventory (IMI), a questionnaire providing quantitative data on participant’s perceived intrinsic motivation (University of Rochester, 2010) was used in a nine-question form. Like that of the first study, the IMI measures the participants motivations, which is compared to their total task time, to help further understand and act as a check for their IMI results. The musical interface questionnaire (Table 4) gives more context to the statistical results offered by the IMI and the MIDI data. Questions meant to derive a mean overview of the age group, cultural background and musical experience of the participant pool, in addition to comments on their experience.

4.4.5 Results & Discussion

The IMI results are presented Figure 14 and they demonstrate that though there was an average preference for both call-response and accompaniment conditions in enjoyment and perceived competence, statistical significance was only found in participant’s enjoyment of the accompaniment condition

$t(11) = -2.2, p < 0.05$ and perceived competence during the call-response condition $t(11) = -2.38, p < 0.05$.

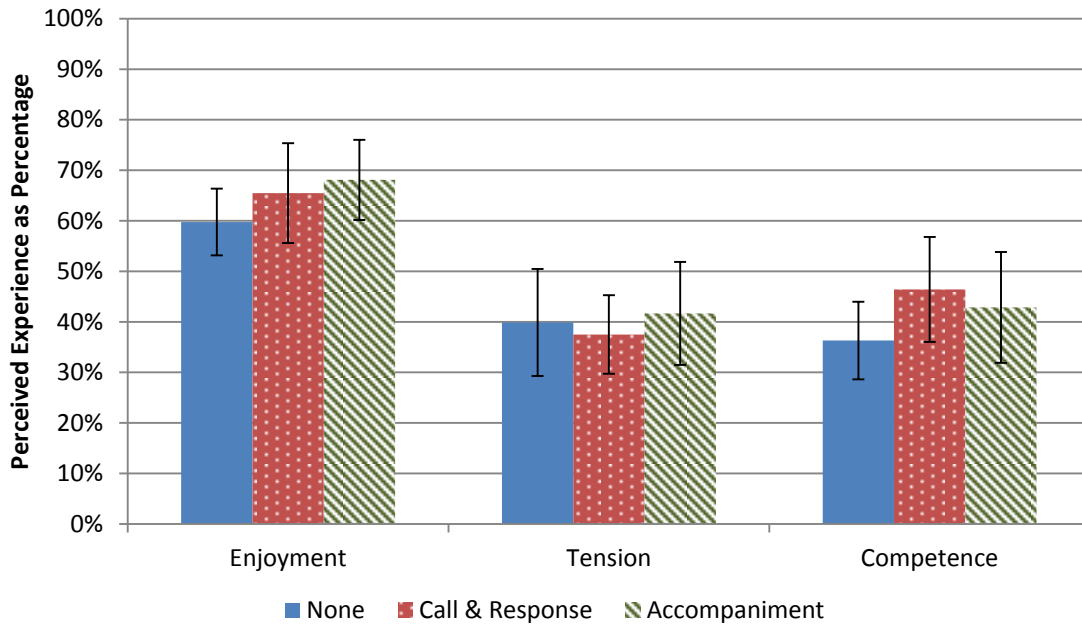


Figure 14: Graph showing IMI results for participants in three conditions. A higher percentage indicates higher perceived experience. The error bars indicate the standard deviation.

In addition to using the IMI, the participants' performance data was recording during each condition, including task time, metronome use, and any resetting of the system's learnt patterns. Figure 15 illustrates the average task time in comparison to metronome use between the three conditions, showing that metronome was on average lower in call-response and accompaniment conditions. Given the average task lengths listed below, participants also used the reset button once every 5.4 minutes in the call-response mode, and every 1.4 minutes in the accompaniment mode.

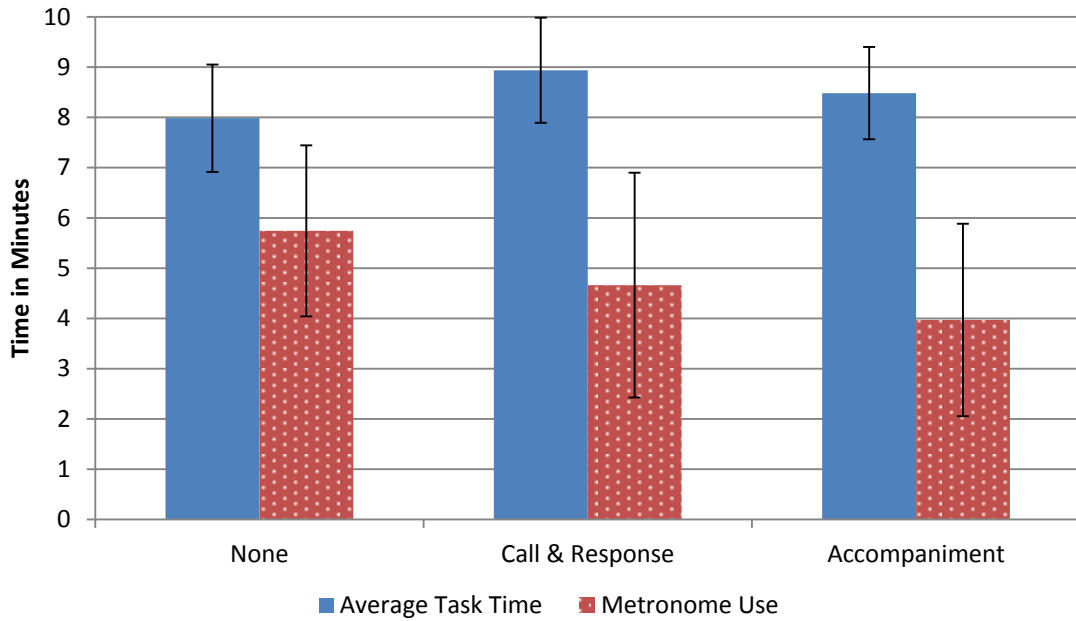


Figure 15: Average task time (in minutes) compared to metronome use between the three conditions. Error bars indicate standard deviation.

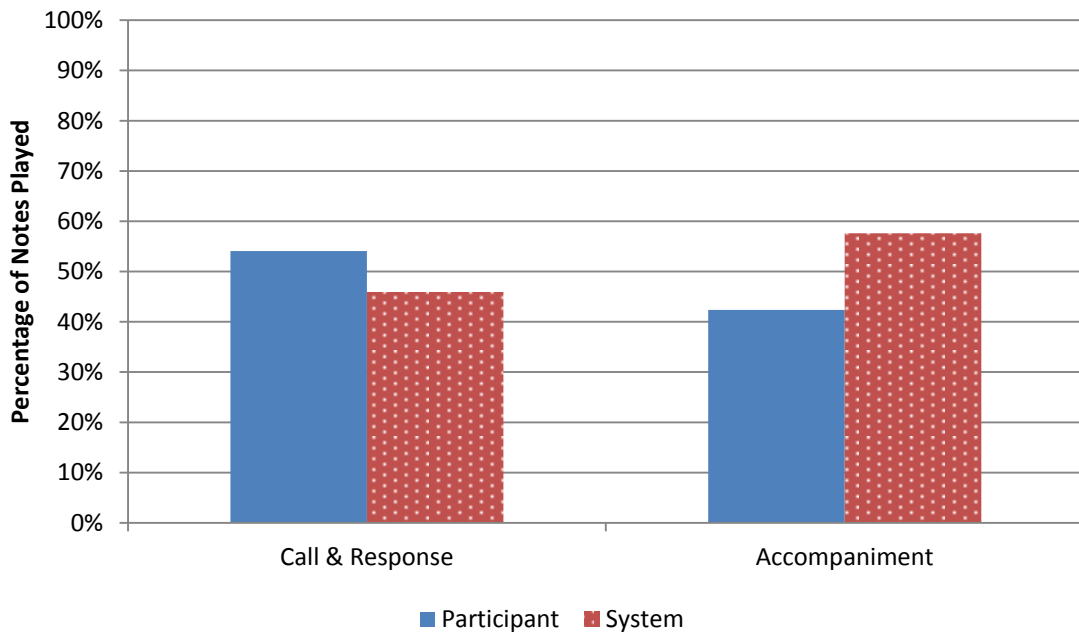


Figure 16: Balance between participant and system performance amount within conditions. Measured by comparing the note density as performed by the participant and the system.

The MIDI data recorded from participants was used to assess the responses by the system, and to look at the balance of the amount performed by participants as compared to the system. The intention in exploring this relationship is to ensure that there were no excessively strong biases towards the participant or system playing for the majority of the task time. Figure 16 shows the balance between participant and system in the call-response (54% to 46%) and accompaniment (42% to 58%) conditions, calculated based on average note-densities during performance. Figures 17-18 show the average note density (in notes per beat) within the call-response and accompaniment conditions (respectively), comparing the system and the participant.

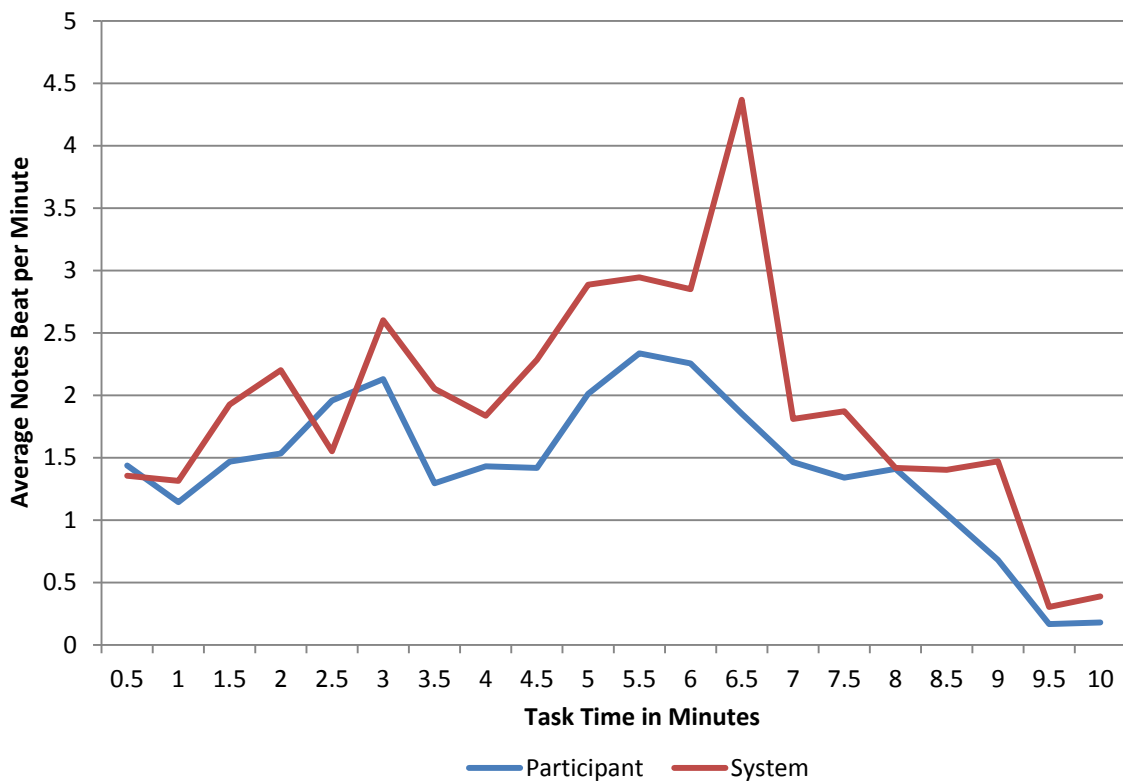


Figure 17: Note density comparison between the participant and system during the accompaniment task

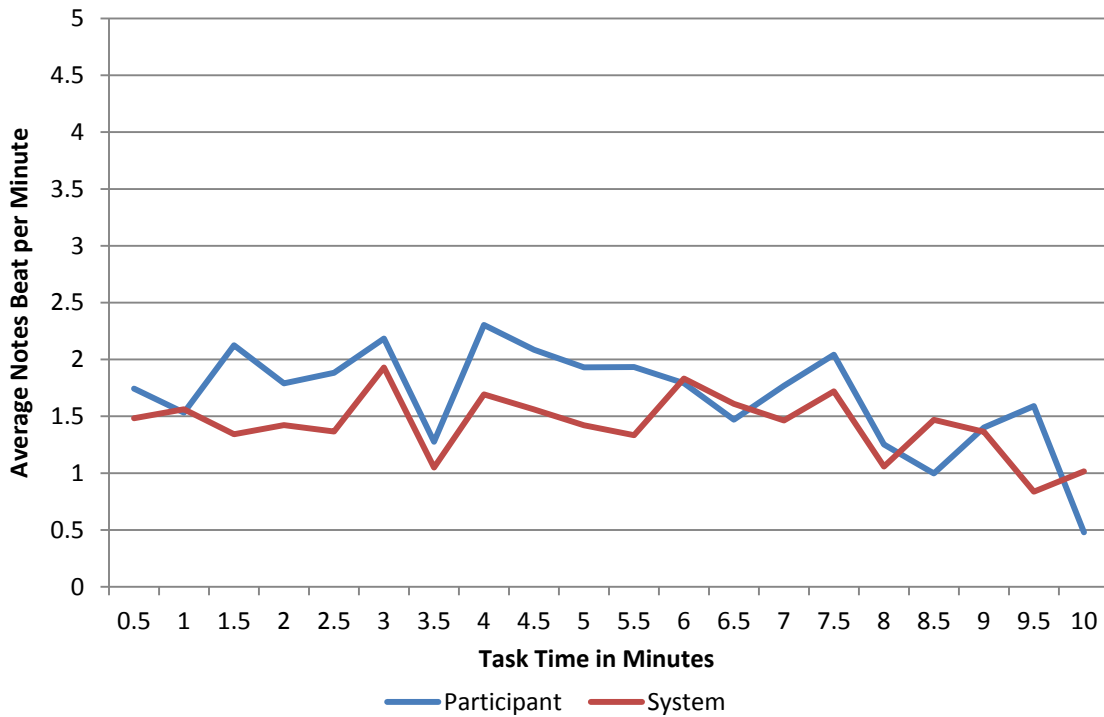


Figure 18: Note density comparison between the participant and system during the call-response task

The IMI results are promising; suggesting that both the call-response and accompaniment modes of interaction offer benefits in perceived enjoyment and competence. Tension in the call-response mode is less than in the control condition, though higher in the accompaniment; which is most likely a result of the quick build-up of patterns that occurs when participants used the accompaniment mode. Conversely the significantly higher enjoyment offered by the same mode is of interest, as it suggests that the somewhat playful (though ultimately more stressing – tension) ability of the participant to build up patterns offers more enjoyment than the simple call-response of patterns. Though, at the same time, the call-response condition makes individuals feel more competent in their performance.

The balance between participant and system performance with the call-response mode favoured the participant, and the accompaniment mode having more of the system. The heavier output of the system in accompaniment makes sense, considering that its output builds continuously until the user resets the system. That both these values are near even demonstrates that Beatback generates output which balances with user input; and that users are willing to attend to the generated output with a similar level of consideration to that of their own. During the study, some participants referred to “identifying with the system”, or trying to “game” or “beat” the system’s output. These comments echo results found by Pachet and Adressi in their research with the Continuator (2004), where kids were not only motivated to engage with the system but also felt that they were playing a game or toying with it. That Beatback resulted in similar comments is promising as it may be beneficial in rhythm practice or learning.

The final elements of interest are the graphs shown in Figures 17-18, which compare the note density of the participant and the system during both tasks. In the call-response condition, the note density of the system is frequently lower than that of the participant, as the system was only capable of outputting patterns when the participant was inactive. Whenever the participant plays, the system stops to listen, and as a result the only times when the system may achieve a higher note-density is when the participant allows it to play an extended response by listening for a longer period of time. Accompaniment demonstrates one of the main concerns with that mode of interaction; that of the user’s inability to easily disable drums that the system has been taught to play.

Particularly around the 6.5-minute mark the average system note-density is high, showing a point in the tasks where patterns were commonly built up to a number of drums, before the system being reset around the 7-minute mark. The inability of a user to reset drums individually (not the system as a whole) in accompaniment mode demonstrates one of the limitations of the Beatback system, which is discussed further in the next section.

4.5 Limitations & Future Work

The inability of a user to disable a drum once they have set that drum playing is one limitation of the Beatback system, but does require resolution none the less. Within the study, the accompaniment mode was expected to require user resetting, as there was no way to stop a drum from sounding once a participant had taught it a pattern. Participants averaged resetting the accompaniment mode once every 1.4 minutes (average 8.5 minute task), which is fairly frequent, and if the mode were to be used extensively, another solution to switching drums on and off would be required. A potential solution to this problem could be including a percussive 'shortcut' to stop the drum: mapping a drum rim-hit to act as an on-off switch for that zone. Interestingly, even while resetting the system the most frequently during accompaniment mode, participants also identified as significantly enjoying the mode the most of the three conditions (Figure 14). This suggests that even while quickly becoming sonically complex, perhaps being able to build up patterns in such a fashion is sufficiently entertaining to explore further.

Given Beatback’s reliance on the metronome to indicate to a user the breaking up of patterns, it was a poor choice to offer participants the option to toggle the metronome. However, understanding how participants interact with the metronome is just as important as the other functionalities of the system, and participant’s use of the toggle helped demonstrate a significant weakness in the design of Beatback. With use of the metronome averaging about 52% in the call-response mode, and 47% in the accompaniment mode many participants cited not liking the constant strict feedback of a metronome when trying to practice or explore different rhythmic patterns. Toggling the metronome resulted in two issues: learning offsets and timing crawls. Learning offsets occurred most often when the metronome was turned off and the participant taught the system a pattern, only to discover that there was a pause (or rest) introduced into the pattern they input. Timing crawls (Figure 19) occurred often when the participant taught the system a pattern with the metronome initially (on time), and then switched off the metronome to teach the system further similar patterns. Occurring most often during the accompaniment mode, participants would teach

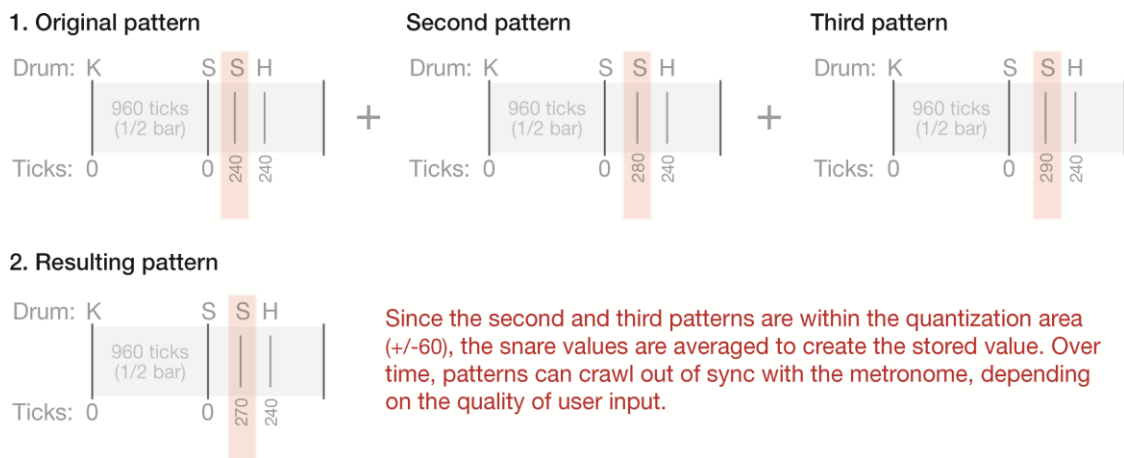


Figure 19: A demonstration and description of how timing crawl occurred during study

the system one pattern using the metronome and then use that drum to base their future patterns on. As they iterated through different patterns on different drums enough latency would be introduced from the original pattern that when the metronome was re-enabled, their resulting pattern sounded offset. There was another concern with the participants not synchronizing with the metronome, one that introduced a clearly audible gap in the recorded pattern. If the user starts a pattern with a gap from the last metronome tick, that gap is introduced into the stored pattern, sounding like an error when the system plays it back to the user. Both of these concerns could be remedied with the introduction of beat tracking into Beatback. By tracking the users' performance live, and adjusting the tempo accordingly, the issues wherein users are not syncing with the metronome would not be a concern: the system would be constantly syncing to the user's performance instead of the user having to follow the system, which also would be of benefit to the usability of Beatback more broadly.

1. User starts pattern with timing space after last pattern reading

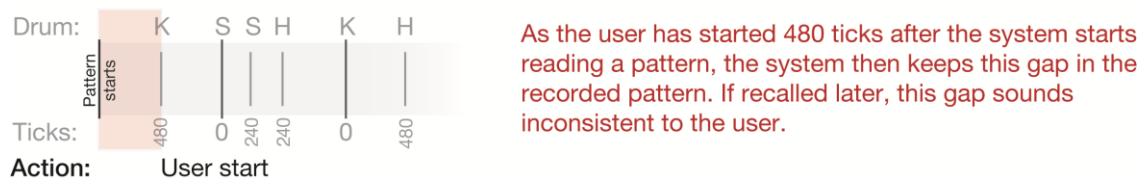


Figure 20: Figure demonstrating the timing gaps that occurred during the study.

Exempting the metronome issues mentioned above, participants cited responses by the system as being 'competitive' or 'thought-provoking'; making them consider how to respond. One participant did suggest that they expected greater variation in the responses by the system that the system responses seemed somewhat too expected given their input. A limitation in the use of a

Markov based system is that it only has user input; its ability to vary based on this input is only as varied as is given. As a result, this comment is not surprising, though perhaps it would be worthwhile in future to look at how participants interact with a pre-established library of material, and explore having the system respond with different orders (to introduce further variety) in the Beatback system. Further studies looking at longer participant interaction without the option of resetting the system's learnt patterns⁷ and having a participant interacting with the system over a number of weeks would help garner insight into other avenues of improving the system, and the quality of its output.

Apart from the improvements which stemmed from research results (discussed above), there is a pair of improvements that are intended for future work on Beatback: (1) improved pattern storage, and (2) storage cleaning. Developing Beatback's ability to work with swung or off-beat patterns is important to its ability to function with a wide variety of rhythmic patterns. Ensuring that the system stores all past instances of a pattern – in lieu of averaging them all out – would allow the system to better learn over time. In addition to pattern storage, cleaning is another proposed improvement to the system. The intent of storage cleaning is that as the transition table becomes larger, transitions that only occurred once can be considered as errors. The patterns identified as errors – if left in the table – can still be called on by the system and may provide the user with undesirable responses. By removing the unlikely transitions and associated patterns, the output of the system is more likely to model the user's input.

⁷ Participants averaged resetting the system in call-response mode once every 5.4 minutes in the average 8.9 minute task.

5: CONCLUSIONS

Earlier on in the paper, some design criteria for percussion interaction were outlined (Section 2.3). These goals were meant to guide the design of Beatback, and result in a system suitable to promoting rhythmic exploration and practice. Amongst those goals was that of *augmentation*, which through using existing hardware and minimizing the need for GUI control of the system the user can focus on their musical interactions instead. The Beatback system has succeeded in achieving this: working with existing electronic drum hardware (through MIDI) and the GUI limits the ability of the user to adjust interaction in real-time, which helps to direct user focus back to the interaction. Features such as the accompaniment model of interaction and zoning - as explored in the studies (Section 4) - demonstrated the benefits offered by augmentation over development of a novel interface (Section 2.2.4). Benefits included a base of practised percussionists who with zoning were able to practise and learn a portion of a complex rhythmic pattern while feeling positively motivated about the task. Similarly, a mix of naïve and practised percussionist demonstrated a significant level of enjoyment in interacting with the accompaniment model of Beatback interaction.

Returning to the design goals, when it came to *stylistically appropriate* responses, or the system's ability to provide material rhythmically consistent with the user's input, the system did not succeed as well as it did in other goals. As

was outlined earlier, design decisions embedded in Beatback came to the fore during research (Section 4.3) and demonstrated that before achieving true *conversational* interaction, some further development is required. In particular, the issues of timing crawl and gaps (Section 4.4) which relied on the metronome were of particular hindrance to providing the user with rhythmically appropriate responses. Frequently during research, participants would encounter the gaps in their patterns, or upon re-enabling the metronome would discover the pattern was not synchronized. One reason this issue came to the fore was due partially to the study design, which allowed participants to enable and disable the metronome, meaning they could remove their reference point for meter. At the same time, gaging how much the participants used the metronome is just as important, as it did offer the direction (and the necessity) of improving the system with live beat tracking. The use of a live tracking would further assist Beatback by removing the need for constant metronome feedback, and would make it capable of providing more conversationally appropriate responses.

Even with some considerable concerns with the ability of the system to respond appropriately to naïve users, saving the results of study was that the system at least achieved the design goal of providing *live* responses to participants. Or, in other words, that although the system was not always appropriate in its response, at least it was prompt in providing a reply. This further demonstrates the importance placed by users on having the responses occur in a timely manner when interacting with a musical system.

Even while considering the issues with output quality of the system, it was capable of *engaging* participants in *focused* interaction. Research reflected the goal of *focused* interaction, first with a 54% to 46% and 42% to 58% user to system performance balance during the study, demonstrating that participants were focused on the systems responses and not simply exploring without regard. Also demonstrating the goal were comments from participants such as feeling 'competitive' or wanting to 'game' the system, which indicated an interest and engagement with the material generated by the system and interacting with it. IMI results further echoed participant comments by showing user enjoyment and competence were significantly benefited by interacting with the accompaniment and call-response modes (respectively) in Beatback.

Looking at one final goal of *iterative* interaction in combination with promoting *engagement* and *focus*, the results from the zoning study are particularly important. Participants in the study were significantly motivated to practice with the zoned as opposed to the un-zoned condition, and as a result is an area that warrants more research. As alluded to (Section 4.4), this study did offer insight into participants' perceived motivations – enjoyment, tension and competence – in the task, and assessed their performance during it, but was limited in its ability to explain their actual learning performance. A review of participant's ability to replicate the pattern post task (without prompting) and the complexity of the pattern to practice are areas that could and should be further assessed to understand the potential of using zoning as a percussion practise or learning feature. The results from this study demonstrating positive motivation

indicate that it would be worthwhile to pursue further research, if only to understand how filtering the pattern may otherwise benefit percussionists who are practising or learning patterns.

Research into Beatback has also echoed other work in the field of music interaction. Pachet's Continuator (2004), from which Beatback received some of its design direction ran similar studies and had similar results: participants are positively motivated when interacting with their own musical material. Pachet's concept of *interactive musical reflexive systems (IMRS)* (Section 2.2.3) is showing the same benefit in the area of percussion interaction as that of piano interaction and the Continuator. Beatback's design helps to further demonstrate the benefits of such a system: augmentation, learning musical material from the performer and providing reasonable responses are all goals that were embedded into Beatback's design. When practicing or learning a pattern, the improvements shown in the zoning study indicate that a user's perceived enjoyment, tension and competence of playing with such a system is improved. In open interaction with Beatback, both accompaniment and call-response interaction models showed a positive influence over enjoyment and competence respectively. Informal or play-oriented learning is essential to musical development (Section 2.4.2) and to continuing one's exploration of rhythm.

As far as we know, the Beatback system is the first interactive reflexive musical system (Section 2.2.3) focusing on rhythm and percussion. Although there are limitations of the system designed (Section 4.5), studies run when using the system have demonstrated that it improves users' intrinsic motivation

(Section 4.3.5 & 4.4.5) by increasing self-perceived enjoyment and competence, while reducing perceived tension. Both practised and naïve percussionists have expressed these improvements to motivations, in learning a pattern and simply interacting with the Beatback system itself. The system also offers the first (to the best of our knowledge) implementation and study of drum zoning (Section 3.2.5), or the filtering out of drums being played by a user. That the IMI results demonstrate significant improvement of motivation when using zoning suggests that it could be a beneficial learning or practise tool for percussionists (Section 4.3). Taking the IRMS model, Beatback branches into percussion and has similarly demonstrated that the IRMS model can be beneficial in promoting and fostering rhythmic exploration.

APPENDICES

Appendix A

IMI Questionnaire used in the Drum-kit Zoning Research (Section 4.2)

Questionnaire (Task One)

P#: _____ C: _____

For each of the following statements, please indicate how true it is for you, using the following scale as a guide.

1. While I was practising the drum-kit, I was thinking about how much I enjoyed it.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

2. I did not feel at all nervous while practicing.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

3. This practice method did not hold my attention at all.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

4. I think I practiced the drum-kit pretty well.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

5. I would describe this practice method as very interesting.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

6. I think I understood this practice method very well, compared to other students.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

7. I enjoyed practicing the drum-kit very much.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

8. I felt very tense while practicing the drum-kit.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

9. The drum-kit was fun to practice.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

Appendix B

IMI Questionnaire and survey used in the Beatback Interaction Research (Section 4.3)

Questionnaire (Task A)

P#: _____ C: _____

For each of the following statements, please indicate how true it is for you, using the following scale as a guide.

1. While I was playing the drum-kit, I was thinking about how much I enjoyed it.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

2. I did not feel at all nervous while playing.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

3. This method of playing did not hold my attention at all.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

4. I think I played the drum-kit pretty well.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

5. I would describe this method of playing as very interesting.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

6. I think I understood this method of playing very well, compared to other participants.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

7. I enjoyed playing the drum-kit very much.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

8. I felt very tense while playing the drum-kit.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

9. The drum-kit was fun to play.

1	2	3	4	5	6	7
Not at all true			Somewhat true			Very true

Questionnaire (Final)

P#: _____

Please answer each question in the space given.

1. How old are you? _____
2. What is your first language? _____
3. Have you ever played any musical instruments? If so, please list them and your experience:
ie. Piano, 4 years with classes

4. Have you ever received formal musical training? If so, please list them and your experience:
ie. Vocal training, 2 years

5. Have you ever played any musical video games? If so, please list them and your experience:
ie. Rockband, guitar, 2 hours

Appendix C: Data Appendix CD-ROM

The CD-ROM, attached, forms a part of this work.

Patch files are Max 5 Collectives (.mxf) and can be opened using Max 5 by Cycling '74.

Patch Files:

Beatback v1r (Jul 2011) (.mxf) 6.32 MB

Version of the Beatback system used in the research study.

Beatback v1t (.mxf) 6.30 MB

Version of the Beatback system last updated before the thesis defence.

REFERENCE LIST

- Wii Drum Machine*. (2006, December 6). Retrieved September 12, 2010, from <http://www.youtube.com/watch?v=8FftLxuYzRA>
- (2007). In N. Collins, & J. d'Esquivan (Eds.), *The Cambridge Companion to Electronic Music* (pp. 171-184). Cambridge: Cambridge University Press.
- More Realistic Wii Drum Kit*. (2007, May 22). Retrieved September 12, 2010, from http://www.youtube.com/watch?v=a8CU1I_8un0
- Guitar Hero Games*. (2010). Retrieved September 14, 2010, from <http://hub.guitarhero.com/games/overview/ghwor>
- Rock Band 3 // Rock Band*. (2010). Retrieved September 14, 2010, from <http://www.rockband.com/games/rockband3>
- Addressi, A. R. (2007). Interactive Reflexive Musical Systems for Music Education. *International Technology, Education and Development Conference*, (pp. 7-10). Valencia.
- Apple. (n.d.). *Apple - Logic Studio - Logic Pro 9*. Retrieved January 25, 2011, from <http://www.apple.com/logicstudio/logicpro/>
- Arfib, D., Couturier, J.-M., & Kessous, L. (2005). Expressiveness and digital musical instrument design. *Journal of New Music Research*, 34(1), 125-136.
- Assayag, G., Bloch, G., & Chemillier, M. (2006). OMax-Ofon. *Proceedings of Sound and Music Computing*. Marseille.
- Avid. (n.d.). *M-AUDIO - Fast Track Pro*. Retrieved January 20, 2011, from http://www.m-audio.com/products/en_us/FastTrackPro.html
- Belcher, J. R. (2007). *Embodied Interfaces for Interactive Percussion Instruction*. Blacksburg: Virginia Polytechnic Institute and State University.
- Bhaskaran, P., Blankenship, E. A., Cohen, J. G., Koparkar, S. S., Rashid, & Nik, R. (2009). Collaboration and Learning in the Video Game Rock Band. *Proceedings of the Emerging Technologies Conference*. Cambridge.
- Bianchi, A., & Yeo, W. S. (2009). The Drummer: A Collaborative Musical Interface with Mobility. *Proceedings of New Interfaces for Musical Expression*. Pittsburgh.
- Biles, J. A. (1994). GenJam: A Genetic Algorithm for Generating Jazz Solos. *Proceedings of the International Computer Music Conference*. Aarhus.
- Blaine, T. (2005). The Convergence of Alternate Controllers and Musical Interfaces in Interactive Entertainment. *Proceedings of New Interfaces for Musical Expression*, (pp. 27-33). Vancouver.
- Blaine, T., & Forlines, C. (2002). Jam-O-World: Evolution of the Jam-O-Drum multi-player musical controller into the Jam-O-Whirl gaming interface. *Proceedings of New Interfaces for Musical Expression*. Dublin.

- Blaine, T., & Perkis, T. (2000). The Jam-O-Drum Interactive Music System: A Study in Interaction Design. *Proceedings of the Designing Interactive Systems Conference* (pp. 165-173). Brooklyn: ACM.
- Borchers, J. O. (1997). WorldBeat: Designing a Baton-Based Interface for an Interactive Music Exhibit. *Proceedings of the Computer Human Interaction Conference*. Atlanta: ACM.
- Burgess, D., & Mynatt, E. (1994). 3-D Interactive Percussion: The Virtual Drum Kit. *Proceedings of the Computer Human Interaction Conference*, (pp. 45-46). Boston.
- Chadabe, J. (1977). Some Reflections on the Nature of the Landscape within Which Computer Music Systems are Designed. *Computer Music Journal*, 1(3), 5-11.
- Chadabe, J. (1984). Interactive Composing: An Overview. *Computer Music Journal*, 8(1), 22-27.
- Ching, W. K., & Ng, M. K. (2006). *Markov Chains: Models, Algorithms and Applications*. New York, USA: Springer Science + Business Media Inc.
- Chun, S., Hawryshkewich, A., Jung, K., & Pasquier, P. (2010). Freepad: A Custom Paper-based MIDI Interface. *Proceedings of New Interfaces for Musical Expression*, (pp. 31-36). Sydney.
- Croft, J. (2007). Theses on liveness. *Organised Sound*, 12(1), 59-66.
- Cycling '74. (n.d.). *Max - Cycling '74*. Retrieved 10 24, 2010, from <http://cycling74.com/products/maxmspjitte/>
- Dannenber, R. B. (2007). New Interfaces for Popular Music Performance. *New Interfaces for Musical Expression*, (pp. 130-135). New York.
- Davidson, J. W., Howe, M. J., & Sloboda, J. A. (1997). Environmental factors in the development of musical skill over the life span. *The Social Psychology of Music*, 188-208.
- Davidson, J., & Jordan, N. (2007). Private Teaching, Private Learning: An Exploration of Music Instrument Learning in the Private Studio, Junior and Senior Conservatories. In *International Handbook of Research in Arts Education* (Vol. 16, pp. 729-744). Springer International Handbooks of Education.
- Deci, E. L., Koestner, R., & Ryan, R. M. (2001). Extrinsic Rewards and Intrinsic Motivation in Education: Reconsidered Once Again. *Review of Educational Research*, 71(1), 1-27.
- Dobrian, C., & Koppelman, D. (2006). The 'E' in NIME: Musical Expression with New Computer Interfaces. *Proceedings of New Interfaces for Musical Expression*, (pp. 277-282). Paris.
- Dunlop, M. D., & Crossan, A. (2000, June). Predictive text entry methods for mobile phones. *Personal and Ubiquitous Computing*, 4(3), 134-143.
- Eerola, T., & Toiviainen, P. (2004). *MIDI Toolbox: MATLAB Tools for Music Research*. Retrieved April 16, 2010, from University of Jyväskylä: Kopijyvä, Jyväskylä, Finland.: <http://www.jyu.fi/hum/laitokset/musiikki/en/research/coe/materials/miditoolbox/>.

- Eerola, T., & Toiviainen, P. (2004). MIR in Matlab: The MIDI Toolbox. *International Conference on Music Information Retrieval*, (pp. 22-25). Barcelona.
- Eigenfeldt, A. (2006). Kinetic Engine: Toward an Intelligent Improvising Instrument. *Proceedings of Sound and Music Computing*. Marseille.
- Eigenfeldt, A. (2009). The Evolution of Evolutionary Software: Intelligent Rhythm Generation in Kinetic Engine. *Proceedings of EvoWorkshops 2009* (pp. 498–507). Berlin: Springer.
- Eigenfeldt, A., & Kapur, A. (2008). An Agent-based System for Robotic Musical Performance. *Proceedings of New Interfaces for Musical Expression*. Genova.
- Eigenfeldt, A., & Pasquier, P. (2010). Realtime Generation of Harmonic Progressions Using Controlled Markov Selection. *Proceedings of the First International Conference on Computational Creativity*, (pp. 16-25). Portugal.
- Fels, S. (2000). Intimacy and Embodiment: Implications for Art and Technology. *Proceedings of the International Multimedia Conference* (pp. 13-16). Los Angeles: ACM.
- Ferguson, S. (2006). Learning Musical Instrument Skills Through Interactive Sonification. *Proceedings of New Interfaces for Musical Expression*, (pp. 384-389). Paris.
- Fredericks, J. A., Blumenfeld, P. C., & Paris, A. H. (2004). School Engagement: Potential for the Concept, State of the Evidence. *Review of Educational Research*, 74(1), 59-109.
- Gray, R., Lindesll, S., Minster, R., Symonds, I., & Ng, K. (2009). An Augmented Snare Drum. *Proceedings of the International Computer Music Conference*. Montreal.
- Gurevich, M., & Treviño, J. (2007). Expression and Its Discontents: Toward an Ecology of Musical Creation. *Proceedings of New Interfaces for Musical Expression*, (pp. 106-111). New York.
- Hawryshkewich, A., Eigenfeldt, A., & Pasquier, P. (2010). Beatback: A Real-time Interactive Percussion System for Rhythmic Practice and Exploration. *Proceedings of New Interfaces for Musical Expression*, (pp. 100-105). Australia.
- Hsu, W., & Sosnick, M. (2009). Evaluating Interactive Music Systems: An HCI Approach. *Proceedings of New Interfaces for Musical Expression*, (pp. 25-28). Pittsburgh.
- Hunt, A., & Hermann, T. (2004). The Importance of Interaction in Sonification. *Proceedings of the International Conference on Auditory Display*. Sydney.
- Hunt, A., Wanderley, M. M., & Kirk, R. (2000). Towards a Model for Instrumental Mapping in Expert Musical Interaction. *Proceedings of the International Computer Music Conference*. San Francisco.
- Johnston, A., Candy, L., & Edmonds, E. (2009). Designing for Conversational Interaction. *Proceedings of New Interfaces for Musical Expression*, (pp. 207-212). Pittsburgh.

- Jordà, S. (2004). Digital Instruments and Players: Part 1 - Efficiency and Appretniceship. *New Interfaces for Musical Expression*, (pp. 59-63). Hamamatsu.
- Jordanous, A. (2011). Evaluating Evaluation: Assessing Progress in Computational Creativity Research. *Proceedings of the Second International Conference on Computational Creativity*, (pp. 102-107). Mexico City.
- Kapur, A., Singer, E., Benning, M. S., & Tzanetakis, G. (2007). Integrating HyperInstruments, Musical Robots & Machine Musicianship for Noth Indian Classical Music. *Proceedings of New Interfaces for Musical Expression*, (pp. 238-241). New York.
- Kaye, J., Boehner, K., Laaksolahti, J., & Staahl, A. (2007). Evaluating Experience-Focused HCI. *CHI Extended Abstracts on Human Factors in Computing Systems* (pp. 2117-2120). New York: ACM.
- Kelley, R. (2003). *Markov Music - A Generative Music Program*. London: University of London.
- Kiefer, C., Collins, N., & Fitzpatrick, G. (2008). HCI Methodology for Evaluating Musical Controllers: A Case Study. *Proceedings of New Interfaces in Musical Expression*. Genoa.
- Knowles, M. S. (1975). *Self-directed learning: A guide for learners and teachers*. London, England: Cambridge Book Company.
- Kochanski, G. (2005, March 8). *Markov Models, Hidden and Otherwise*. Retrieved July 11, 2011, from Greg Kochanski: <http://kochanski.org/gpk/teaching/0401Oxford/HMM.pdf>
- Konishi, Y., & Miura, M. (2010). Estimating musical score and proficiency. *Proceedings of 20th International Congress on Acoustics*, (p. 7). Sydney.
- Lehman, A. C. (1997). The Aquisition of Expertise in Music: Efficiency of Deliberate Pracitce as a Moderating Variable in Accounting for Sub-Expert Performance. In I. Deliege, & J. A. Sloboda (Eds.), *Perception and Cognition of Music* (pp. 161-187). Hove, East Sussex, United Kingdom: Psychology Press.
- Levisohn, A., & Pasquier, P. (2008). BeatBender: Subsumption Architecture for Rhythm Generation. *Proceedings of the ACM International Conference on Advances in Computer Entertainment*, (pp. 51-58). Yokohama.
- Machover, T. (2002). Instruments, Interactivity, and Inevitability. *Proceedings of New Interfaces for Musical Expression*. Dublin.
- Magnusson, T., & Mendieta, E. H. (2007). The Acoustic, the Digital and the Body: A Survey on Musical Instruments. *New Interfaces for Musical Expression*, (pp. 94-99). New York.
- Maki-Patola, T., Hämäläinen, P., & Kanerva, A. (2006). The Augmented Djembe Drum - Sculpting Rhythms. *Proceedings of New Interfaces for Musical Expression*, (pp. 364-370). Paris.
- MIDI Manufacturers Association Incorporated. (n.d.). *MIDI Message Table 1*. Retrieved January 17, 2010, from <http://www.midi.org/techspecs/midimessages.php>

- Miranda, E., & Wanderley, M. (2006). *New Digital Musical Instruments: Control and Interaction Beyond the Keyboard*. Middleton: A-R Editions Inc.
- Moore, F. R. (1988). The dysfunctions of midi. *Computer Music Journal*, 12(1), 19-28.
- Murray-Rust, D. (2003). *VirtuaLatin: Agent Based Percussive Accompaniment*. Edinburgh: University of Edinburgh.
- Neuhoff, J. G., Kramer, G., & Wayand, J. (2002). Pitch and loudness interact in auditory displays: Can the data get lost in the map? *Journal of Experimental Psychology*, 8(1), 17-25.
- Ng, K., & Nesi, P. (2008). i-Maestro: Technology-Enhanced Learning and Teaching for Music. *Proceedings of New Interfaces for Musical Expression*. Genoa.
- Orio, N., & Wanderley, M. M. (2002). Input Devices for Musical Expression: Borrowing Tools from HCI. *Computer Music Journal*, 26(3), 62-76.
- Pachet, F. (2003). The Continuator: Musical Interaction With Style. *Journal of New Music Research*, 32(3), 333-341.
- Pachet, F. (2004, July 5). *The Continuator*. Retrieved October 1, 2010, from <http://www.csl.sony.fr/~pachet/Continuator/>
- Pachet, F. (2006). Enhancing Individual Creativity with Interactive Musical Reflective Systems. In I. Deliege, & G. Wiggins (Eds.), *Musical Creativity: Current Research in Theory and Practice*. London, England: Psychology Press.
- Pachet, F., & Adessi, A. R. (2004, January). When Children Reflect on Their Playing Style: Experiments with the Continuator and Children. *Computers in Entertainment*, 2(2).
- Pauletto, S., & Hunt, A. (2004). A Toolkit for Interactive Sonification. *Proceedings of the International Conference on Audio Display*. Sydney.
- Poepel, C. (2005). On Interface Expressivity: A Player-Based Study. *Proceedings of New Interfaces for Musical Expression*, (pp. 228-231). Vancouver.
- Price, S., & Rogers, Y. (2004). Let's get physical: The learning benefits of interacting in digitally augmented physical spaces. *Computers & Education*, 43, 137-151.
- Rae, A. (2009). *Generative Rhythmic Models*. Atlanta: Georgia Institute of Technology.
- Roh, J. H., & Wilcox, L. (1995). Exploring tabla drumming using rhythmic input. *Proceedings of Human Factors in Computing Systems*, (pp. 310-311). Denver.
- Roland. (n.d.). *Roland TD-9KX V-Drums V-Tour Series :: Overview*. Retrieved January 25, 2011, from <http://www.roland.com/products/en/TD-9KX/index.html>
- Rowe, R. (2001). *Machine Musicianship*. Cambridge: MIT Press.
- Russell, S. J., & Norvig, P. (2003). *Artificial Intelligence: A Modern Approach* (2nd ed.). New Jersey, USA: Pearson Education Inc.
- Ryan, R. M., & Deci, E. L. (2009). Promoting self-determined school engagement: Motivation and well-being. In K. R. Wentzel, & A. Wigfield, *Handbook on motivation at school* (pp. 171-196). New York: Routledge.

- Ryan, R. M., Connell, J. P., & Plant, R. W. (1990). Emotions in Nondirected Text Learning. *Learning and Individual Differences*, 2(1), 1-17.
- Salter, L., & Hearn, A. M. (1996). *Outside the Lines: Issues in Interdisciplinary Research*. Quebec City: McGill-Queen's University Press.
- Serra, X. (2005). Towards a Roadmap for the Research in Music Technology. *Proceedings of the International Computer Music Conference*. Barcelona.
- Sharp, H., Rogers, Y., & Preece, J. (2007). *Interaction Design*. New York: Wiley.
- Tanaka, A., & Parkinson, A. (Performers). (2010). *4 Hands iPhone*. Cube Festival; Charm of Sound Festival; Mois Multi Festival; Music with a View Festival, Flea Theatre; Electron Festiva, Paris; Helsinki; Quebec; New York City; Geneva.
- Tidemann, A., & Demeris, Y. (2008). A Drum Machine That Learns to Groove. *German Conference on Advances in Artificial Intelligence*, (pp. 144-151). Heidelberg.
- Tsigilis, N., & Theodosiu, A. (2003, August). Temporal Stability of the Intrinsic Motivation Inventory. *Perceptual and Motor Skills*, 97(1), 271-280.
- University of Rochester. (2010, March 4). *Self-Determination Theory*. Retrieved March 4, 2010, from University of Rochester: <http://www.psych.rochester.edu/SDT/measures/>
- Waterworth, J., & Fallman, D. (2005). Dealing With User Experience and Affective Evaluation in HCI Design: A Repertory Grid Approach. *Proceedings of the CHI Workshop on Evaluating Affective Interfaces - Innovative Approaches*. Oregon: ACM.
- Weinberg, G., & Driscoll, S. (2007). The Interactive Robotic Percussionist - New Developments in Form, Mechanics, Perception and Interaction Design. *Proceedings of the Human Robotic Interaction Conference* (pp. 97-104). Arlington: ACM.
- Weinberg, G., Aimi, R., & Jennings, K. (2002). The Beatbug Network: A Rhythmic System for Interdependent Group Collaboration. *Proceedings of New Interfaces for Musical Expression*. Dublin.
- Weinberg, G., Beck, A., & Godfrey, M. (2009). ZOOZBeat: A Gesture-based Mobile Music Studio. *Proceedings of New Interfaces for Musical Expression*, (pp. 312-315). Pittsburgh.
- Weinberg, G., Driscoll, S., & Thatcher, T. (2006). Jam'aa - A Middle Eastern Percussion Ensemble for Human and Robotic Players. *Proceedings of the International Computer Music Conference*. New Orleans.
- Weinberg, G., Godfrey, M., Rae, A., & Rhoads, J. (2007). A Real-time Genetic Algorithm in Human-robot Musical Improvisation. *Proceedings of Computer Music Modeling and Retrieval*, (pp. 351-359). Denmark.
- Weinberg, G., Raman, A., & Mallikarjuna, T. (2009). Interactive Jamming with Shimon: A Social Robotic Musician. *Proceedings of the Human-Robot Interaction Conference*, (pp. 233-234). La Jolla.
- Young, D., & Fujinaga, I. (2004). Aobachi: A New Interface for Japanese Drumming. *Proceedings of New Interfaces for Musical Expression*, (pp. 23-26). Hamamatsu.