Performer interaction and expectation with live algorithms: experiences with Zamyatin

Author:
Bown, O

Publication details:
Digital Creativity
v. 29
Chapter No. 1
pp. 37 - 50
1462-6268 (ISSN); 1744-3806 (ISSN)

Publication Date:
2018-01-02

Publisher DOI:
https://doi.org/10.1080/14626268.2018.1432663

Performer Interaction and Expectation with Live Algorithms: Experiences with *Zamyatin*

Oliver Bown

December 6, 2017

1 Introduction

The scenario of a live improvising musician getting up on stage to perform with a seemingly 'autonomous' software system is becoming an increasingly regular occurrence. Such autonomous software systems have been termed 'live algorithms' (Blackwell et al., 2012), and have also been documented and discussed under a number of other categories such as 'machine musicianship' (Rowe, 2002; Collins, 2016), 'algorithmic composition' (Papadopoulos and Wiggins, 1999) and 'musical metacreation' (Bown et al., 2013; Eigenfeldt et al., 2013a). A growing number of practitioners in and beyond academia have presented, documented, discussed and reported on their experience with their own improvising systems, and there is now a considerable body of practice available to study, warranting a discussion of the range of different approaches and experiences across the field.

An ongoing focus of interest in this body of practice is the issue of how system creators and those that experience interaction with these systems – performers and audiences – conceptualise autonomy and related abstract qualities of the systems as independent agents (discussed previously in Bown et al. (2013)). Although a common starting point is to think of such systems as virtual musicians, and to frame work in terms of the reproduction of human musical competencies, in reality human cognitive capabilities are so out of reach of this current technology as to make for implausible goals for system builders, and ineffective reference points for participants. Real systems are developed by creatively assembling working programs from an existing set of known techniques in computer science, audio analysis and artificial intelligence. Examples include Assayag and Dubnov's use of factor oracles (Assayag and Dubnov, 2004), Brown et al's use of time-series methods (Brown et al., 2013), Eigenfeldt's use of multi-agent systems (Eigenfeldt and Pasquier, 2009), Blackwell and Young's use of swarm dynamics (Blackwell and Young, 2004), and Eldridge's use of cybernetic principles (Eldridge, 2007). Authors working with such applications may make appeals to first order approximations of human behaviour. But in these examples, where the focus is a technical pursuit of creating a system for live performance, grounded in a specific creative practice, the implication is that these systems open up a new space of interactive performance that is not necessarily best grounded in comparisons with human behaviour[1].

Research in computational creativity embraces this complexity of evaluation. Although computational creativity's most widely used definition appeals to human comparison ("the performance of tasks [by a computer] which, if performed by a human, would be deemed creative") (Cardoso et al., 2009), this does not require that the performance in question is itself human-like. More recently Jordanous (2012) has developed an evaluative framework which treats creativity as a cluster concept with 14 different associative terms that participants might experience as creativity, such as "domain competence" or "spontaneity", enabling both greater abstraction and finer detail in evaluations of creativity. Other work in computational creativity evaluation, such as that of Lamb et al. (2015) and Pasquier et al. (2016), continues to investigate how non-trivial issues of familiarity, background, goals and expectation may influence how people judge computer produced creative work.

Building improvising music systems invites the questions of whether performing musicians also enter into performances with such systems without needing to frame their behaviour in terms of anthropomorphic concepts or human reference points (Collins, 2006; Banerji, 2012). Do they employ more abstract conceptions of interaction, either by inferring them or through awareness of computer science and artificial intelligence achievements and limitations? For example, most people know that computers aren't very good at understanding language in open-ended conversational contexts, but musicians will know of algorithms that can detect pitches and beats quite well given a clean isolated audio stream. Such background information enables anyone to potentially approach novel systems with relatively known constraints on their competencies, acting according to the script of a designer with an interaction goal. As in the wider world of human computer interaction, musicians are becoming increasingly familiar with the experience of something being both designed but also exhibiting complex behaviour that may be reminiscent of something 'lifelike'[2].

In an earlier study (Bown, 2015), a workshop was set up in which three improvising musicians performed with the author's system, *Zamyatin* (first introduced in (Bown, 2011)), and also watched each other performing with it, followed by a focus group discussion and survey. This study offered some evidence supporting the idea that musicians are comfortable conceiving of the autonomy and intelligence of interactive music systems in a way that does not need to use human musicians as a reference point, but equally does not diminish the system to more 'inanimate' musical categories such as 'instrument' or 'composition' (see (Bown et al., 2009) for further discussion of these categories in relation to digital music practice). The study also suggested that playing with a system offers quite different

[1] And as Sturm (2016) argues, researchers may want to emphasise the importance of steering audiences away from such comparisons.

[2] A reviewer of this paper comments that there are historical examples of people designing lifelike behaviours going far back in time. This is true, but such designs are nevertheless becoming increasingly familiar and complex.

responses to watching the system being played with.

This paper looks further into this question by investigating how a single musician enters into and experiences interaction with *Zamyatin*, with specific attention to the author's own creative process. It reports on live and recorded performances between *Zamyatin* and a well-known improvising musician, François Houle, along with interviews with the musician concerning their experience. The aim of this study is to begin to build a descriptive understanding of performer responses to illustrate what conceptual tools guide the approach to performing with the system. It is limited in its scope: this is a single user study, meaning that results cannot be generalised, and although it describes an ongoing collaboration, the study is not itself longitudinal.

2 Autonomy and Stances

2.1 Framing Autonomy

To be autonomous, in simple abstract terms, is to have control of one's future. In complex systems science, an effort is underway to operationalise autonomy in terms of the relative role of external and internal causes. For example, Seth (2010) draws on Granger's definition of causality (G-causality), which states that: "Y causes X if the inclusion of past observations of Y reduces the prediction error of X in a linear regression model of X and Y, as compared to a model that includes only the previous observations of X". Seth's G-autonomy builds on this concept to measure the extent to which a variable is "dependent on its own history and ... these dependencies are not accounted for by external factors", i.e. the extent to which a variable is self-determining. As with G-causality, it asks whether knowledge of the system's own history reduces the prediction error of its own future state, as compared to predictions based only on external elements.

G-autonomy is used here as a guide for thinking about autonomy in terms of causality and predictability, although no mathematical treatment is offered here. It is intuitive. G-autonomy predicts that highly static things are not autonomous because their future states are highly predictable, and highly random things are also not autonomous because their future states are highly unpredictable. In both cases, no *improvement* can be made to prediction error. On the other hand, a complex organism reacting to its environment, something we intuitively think of as autonomous, is G-autonomous because it has internal cognitive states that are key to predicting its future behaviour.

The autonomy of a system can look different depending on the context in which it examined and the depth and breadth with which state is analysed. Autonomy may need to be probed to be properly evaluated (like prodding an animal to tell if it is alive). This matters when considering the autonomy of improvising musicians and related improvising software systems. The system's apparent autonomy is influenced by the way in which any participant interacts with the system, which can in turn be influenced by their expectation. A system may exhibit rich autonomy according to a certain frame such as within a structured jazz improvisation, but be completely incapable of a wider set of actions, such as pushing the musical style into new territory.

Such subtleties mean that the question “is it autonomous?” is preferably exchanged for “how is it autonomous?” Accordingly, different flavours of musical autonomy, are recognised in the literature. For example, Eigenfeldt et al. (2013b) attempt to form a series of levels of autonomy in musical metacreation:

1. Independence: the use of any process on a musical gesture that is beyond the control of the composer.

2. Compositionality: the use of any process to determine the relationships between pre-defined musical gestures.

3. Generativity: the generation of musical gestures.

4. Proactivity: system/agents that are able to initiate their own musical gestures.

5. Adaptability: a) Agents behave in different ways over time due to their own internal evolution; b) agents interact and influence one another.

6. Versatility: agents determine their own content without predefined stylistic limits.

7. Volition: agents exhibit volition, deciding when, what, and how to compose/perform.

Although these are ordered, each of these types of autonomous behaviour themselves constitute a range of possibilities from trivial to sophisticated; lower-numbered elements might involve more sophisticated or complex behaviours than higher-numbered elements. As an exercise one can think of the most flippant way to implement any of these properties in a computer program. The answer could be just a couple of lines of code.

Interestingly, interactive causality and autonomy may not be as important as we might have expected in musical performance. Pachet & Roy (2017) recently studied a jazz ensemble, looking at the G-causality of interaction between players based on their audio signals. The study showed minimal causality between musicians, compared to the high causality of the score on the musicians. This is perhaps a surprising result, pointing to a possible model of structured improvisation in which the score acts as a medium of interaction, and musicians don’t interact so much with each other as with the musical surface.

2.2 Experiencing Systems: Stances and Embodiment

According to Dennett (1989), we assume different stances when approaching different entities in the world. Many things, particularly the abiotic elements of the natural world are best approached with a physical stance. In terms of dealing with and anticipating such objects’ behaviour, we are concerned with the laws of physics and the properties of materials;

how the object flies through the air, bounces or breaks. Another set of exclusively human-made objects are approached instead with a design stance. In this case we are aware that the object has a designer, and a function, and anticipate its behaviour following design heuristics. What does this button 'do'? How am I supposed to hold it? What does that bleep signify? Finally, we can approach animals, people and some human-made systems with an intentional stance. We approach their behaviour knowing that they have goals, and more generally that they perform cognition; they know things and plan their actions. With humans, we grapple with complex theory of mind issues (Tomasello and Call, 1997). For example, is my facial expression giving away my feelings to those around me?

Dennett describes these stances as heuristic methods for framing one's understanding of interaction with different things. One adopts an intentional stance if it is the easiest way to think about the system in question. A physical stance would be useless when trying to predict the action of a tennis opponent, but very useful when trying to predict where the ball will land. Multiple stances may be combined. Motorists simultaneously deal with physics, controls and readouts, as well as predictions about other people's behaviour, when negotiating the traffic.

We can also think about our interaction with physical and digital objects in terms of situated and coordinated action. Malafouris (2013), for example, uses the example of a blind person's stick as being something they experience the world through, rather than experiencing directly; embodied interaction is less to do with what stances we adopt towards other things and more to do with how closely coupled we become with those things, which we may cease to directly notice. Such a description is familiar to musicians, with respect to their relationships to their instruments. Here the notion of stances is replaced with an adaptive embodied view in which any system, simple or smart, can potentially be mapped in interaction to become an extension of an individual's action or perception system. Thus as opposed to stances we can think about the ease of coupling between a musician and the entities they are interacting with.

3 The Philosophy and Design of *Zamyatin*

Zamyatin is a software system in ongoing development by the author since 2010 (Bown, 2011). Before describing the design of *Zamyatin*, I will explain the creative goal of the project and some related design considerations, including aesthetic decisions related to the above considerations. An earlier detailed description of *Zamyatin*'s design is given in (Bown, 2011)[3].

The primary goal of *Zamyatin* is to produce live performances in which a free improvising musician interacts with a software system, resulting in a duet in which the musician experiences a sense of engagement with another agent in a musical interaction. The system

[3] The source code for *Zamyatin* is also available on the author's Github page at http://www.github.com/orsjb/.

is designed to be autonomous in the sense that it does not need to be controlled by an operator during a live performance, maintaining engaging interaction and development during the piece, but not in the wider sense of choosing musical aesthetic elements to suit the performance. From the taxonomy of Eigenfeldt et al. (2013b) it is intended to be independent, compositional and proactive. The system can be seen as a form of dynamic algorithmic composition by the author. A large number of compositional decisions are made by the author, including selections of scales, timbres, rhythmic patterns and samples. The author also uses evolutionary methods (described below) to search for possible behavioural models that *Zamyatin* might use, and makes the ultimate creative decision about which routines will be used in a performance. In short, the great majority of creative decision making is performed by the author.

Zamyatin is also not designed to enact specific instances of musical intelligence defined in music theoretic terms. It responds only to low-level realtime features derived from the musician's audio. It does not attempt to build statistical models of the performer's actions or even record note sequences into memory to operate upon. It does not attempt to look ahead and generate sequences that satisfy specific requirements. It does not attempt to track the tempo or meter of the performance and has no general model of tempo or meter. It does not implement machine learning algorithms, and does not have musical rules programmed into it (with an exception; see the description of the generative component, (c), below).

These are things that many algorithmic music systems attempt to do to varying degrees of success, but *Zamyatin* has a much simpler objective in terms of autonomous musical interaction, to find interesting pattern-producing couplings with the performer that satisfy dual constraints: to stimulate the interaction, in particular to make the performer and audience feel that the performing system has a live agency; and to be easily adapted by the composer to creative contexts, i.e., to be creatively malleable in a compositional context. Thus *Zamyatin* is far removed from attempts to model human musical behaviour.

Instead, *Zamyatin* is conceived of as a simple artificial organism: a complex non-linear system that embodies dynamic properties that make for musically engaging interaction. One can describe *Zamyatin*'s behaviour by appeal to cybernetic principles of dynamic systems, specifically the liquid analogy used in liquid state machines or the liquid brain model: like a bucket of water, the system has a number of natural resting states when it is not being driven by an input. But when something stimulates it it jumps into action, rippling and oscillating, before settling down again, possibly in a new state. The system might resonate with its input or jump into different dynamic modes. The bucket of water is a simple analogy to this more general set of system behaviours. *Zamyatin*'s dynamic behaviour might also drive itself in indefinite patterns, or perform responses that are of a more cognitive nature, via non-linear computations. Rather than learning from musical examples, these different behaviours are evolved using methods from evolutionary computing, according to abstract descriptions of system behaviours.

Figure [FIGURE 2] shows the system design. *Zamyatin* has a fixed component consisting of realtime audio analysis and a container that can load different dynamic models (left). Low level audio features are extracted and fed into the dynamic model as control data. This data is normalised (mapped to a range 0-1) and is updated at a rate typically set in a range between 10-100hz. Different dynamic models (centre) can be loaded in to *Zamyatin*, representing different responsive behaviours of the system. For example, one model might oscillate regularly, but with the audio input perturbing its oscillations, slowing them down or speeding them up. Another model might exhibit chaotic patterning, and another might exhibit non-linear responses to input. The output from the dynamic model is abstract numeric data. Dynamic models are generated using different processes (see below), and it is an important part of *Zamyatin*'s creative design process that this component is modular and swappable. For the decision trees used by the version of *Zamyatin* described here, the output takes the form of both continuous and discrete values which can be chosen by the composer to map to different musical parameters.

Dynamic models are evolved to exhibit interactive behaviours of interest by running then in a simulation environment, recording their output, and analysing outputs for a number of properties such as overall entropy, divergence of output for different inputs, and presence of regular oscillations. A number of decision trees are evolved to exhibit different combinations of these features. These are stored in a pool which the composer can search by trial and error. The number of combinatoric possibilities is enormous, even from a modest number of generated dynamic models, but the goal is simply to give the composer access to enough non-trivial behaviours that they can search by trial and error to find something that works.

The output of the dynamic model is fed into a number of generative music voices (right of Figure [FIGURE 2]) that the composer creates manually, and that they can turn on or off in real time to orchestrate the performance. These voices might range from continuous synthesiser drones that the dynamic model modulates various parameters of, to more detailed musical sequence generators, that may contain compositional material or generative processes based on musical rules, but that are also controlled by numerous parametric hooks into the system. It is helpful to think of the mapping from the central module (dynamic model) to the output module (voices) as analogous to the operations a musician performs on an instrument (and therefore also analogous to MIDI control). If that instrument were a saxophone this would include continuous micro-control of breath, specific discrete button pushes corresponding to the control of notes, and so on. If the instrument were a drum machine, then this control might take the form of meta-parameters controlling filters, the tempo, and the mix and effect levels of different elements. Figure [FIGURE 3] details the workflow of developing a new piece with *Zamyatin*.

Previous experiences working with *Zamyatin* have led the author to identify three potential affordances of the system that may support engaging interaction:

- **Low-level sensitivity**. *Zamyatin* allows for low-level sensitivity where very small gestures or changes can trigger more or less predictable responses in the system (satisfying the requirements of G-autonomy when these responses are less predictable or more surprising)

- **Oscillation and quasi-entrainment**. Depending on the evolutionary target behaviour and the success of the evolution of *Zamyatin*'s behavioural models, the system can exhibit interesting oscillatory behaviours. Some models of rhythmic entrainment are based on adaptive oscillators that speed up and slow down and adapt their period of oscillation in response to an input. Although *Zamyatin*'s behavioural models do not exhibit the properties of proper functional adaptive oscillators they do alter their oscillatory behaviour in response to inputs, which makes for an engaging musical discourse that may at times give a sense of entrainment.

- **An octopus of parameters**. *Zamyatin* takes an 'octopus' approach to mapping output control values from the behavioural models to a musical voice. Multiple parameters are controlled simultaneously by different control values, and because the control values are coupled in the pattern generation phase the parameters consequently move in coupled ways that gives the impression of coordinated action. *Zamyatin* can also be stacked up with multiple different voices, each being controlled by multiple control values. Thus the impression of coordinated action can be exaggerated by making the coordination appear to be across multiple instruments, as if *Zamyatin* embodies an entire band (hence the octopus reference).

Each of these affordances, whilst simple to implement, are believed to contribute to the sense of autonomy that performers experience when interacting with the system.

4 Methodology

Pearce, Meredith and Wiggins Pearce et al. (2002) claim that a methodological malaise has led to an all too easy conflation of different motivations driving the creation of autonomous and intelligent music systems. They identify four different categories of research activity: algorithmic composition; the design of compositional tools; computational modelling of musical style; and computational modelling of music cognition. This research is concerned primarily with algorithmic composition, for which a practice-based methodology is warranted. However, it also involves third-party users in the form of performing musicians, who are required to interact with digital systems. Even in the context of an artistic practice, this warrants the addition of components drawn from a design research methodology. This is common in approaches to the creation of interactive digital artworks (e.g., see (Bilda et al., 2008; Candy and Ferguson, 2014)). Whilst respecting the importance of distinguishing between these outcomes and the choice of an appropriate methodology, this research aims for outcomes pertaining to compositional practice, but that may also inform the design of compositional tools as a natural continuation of the work. Here the sense of autonomy and meaningful interaction for those engaging with the system or observing it is important as a creative objective and part of a wider issue of user experience design.

Since the Turing Test (Turing, 1950) is such a common starting point for studying the perception of autonomy in music systems, I first want to address this topic and discuss why I do not take a Turing Test approach. The application of the Turing Test to non-linguistic activities such as the production of music poses additional problems (Pease and Colton, 2011) beyond a number of existing criticisms levelled at the more orthodox use of the test in linguistic contexts (French, 2000; Whitby, 1996; Ariza, 2009). In a seminal paper on the application of the test to music, Ariza (Ariza, 2009) introduces some key terms to the discussion. He begins by distinguishing a more general class of tests called John Henry Tests, in which machines and humans are compared in their ability to perform a specific task, such as playing chess or even more basic industrial tasks such as lifting weights. In these tests, there is a clear outcome that can be measured objectively (even if that objective assessment involves human judges, or the task involves a competition against a human). In such tests, the question is “can the machine do better than the human?”, not “is the machine human-like?”. Such tests do not require blindness or binary forced choice, and can focus purely on performance. Ariza then introduces Harnad’s (1991) discussion of Turing Test variants, organised into levels of sophistication. This includes ‘toy’ Turing Tests, which attempt to evaluate some interaction context that is less complete than Turing’s open conversational paradigm. Ariza defines two sorts of toy test applied to music. The Musical Output toy Test (MOtT - small ‘t’ to distinguish ‘toy’ from ‘Turing’) simply asks judges to listen to instances of music generated by a system, in a blind human-computer guessing game. The Musical Directive toy Test (MDtT) has the user issue directives to the system, which responds by producing outputs, again judged blind. Ariza is dismissive of the value of either restricted form of test, the key point being that they are too easy to win, inviting trivial and deceptive methods to achieve positive results, and consequently telling us little about the qualities of the systems being tested.

4.1 Study Format

This research presents an interview with and observation of a live performer playing with the system, *Zamyatin*. These were used to understand how the performer approached, conceptualised and responded to the system. Inspiration is taken from studies such as those of Collins (2006), who devises performance systems and interviews performers based on their experience of the system’s intelligence and autonomy, and Banerji (2012), who draws on ethnographic research methods, with attention paid not only to how performers evaluated the system given specific questions, but also to how the systems influenced the performer’s playing styles and attitudes. Another source of inspiration in practice-based approach to improvising systems, Lewis' (2000), suggests that a successful systems is one that draws you in and makes you want to respond in certain ways; systems such as his own *Voyager* can therefore be seen as probes that are capable of exploring new non-human interaction strategies. The methods used here aim to elaborate on what makes for successful strategies, whether via capably imitating or 'faking' human strategies, or coming up with new ones.

It is important to acknowledge that within this methodology, as creator of the system and author of this study, I interpret the performer's responses and also provide my own interpretations of what is happening in the piece. This is with full awareness of the potential bias of this interpretation.

5 Performer Study

François Houle is a Canadian clarinetist who is experienced in free improvised music, new music and electroacoustic music. He first played with *Zamyatin* at a concert in Vancouver in 2013. He subsequently visited Australia in 2015 and recorded sessions with *Zamyatin* that were included in a self-released record, *Playing the Castle*, also featuring Australian composer Kim Cunio (link).

In interview, Houle describes his experience playing with *Zamyatin*, referencing his general philosophy of working with interactive music systems:

> "A programmer, or composer/creator of a machine needs to have a very clear musical vision that encompasses a deep understanding of all the parameters involved in a satisfying musical experience. I am talking about a visceral experience, where the music affects the senses in a very 'physical' and 'emotional' way."
>
> "Interaction implies a two-way, action/reaction principle. If the system only reacts to an action without ever anticipating one, or by ever generating an action of its own, you end up with musical failure, as there is too much inequity at the participatory level. You can forget about what or who you are performing with, if the situation meets one's idea of what constitutes a satisfying musical experience."

Houle states that he regards *Zamyatin* as satisfying this criterion of anticipating or generating actions of its own. *Zamyatin* performs no explicit anticipation in the sense of a predictive model of the performer's action or the musical content, but at moments in these recordings Houle's duetting with *Zamyatin* creates a sense of anticipation from the oscillatory behaviour of the system, playing rhythmically in such a way that *Zamyatin* seems to pick up notes on the beat.

In the first of five duets with the system (link), it is possible to hear the basic design elements described above. Low-level sensitivity is evidenced in the early moments of the piece where Houle's notes perceptibly trigger *Zamyatin* to produce sound, although at other times this cause and effect is more ambiguous. At times the low-level triggering is inverted; *Zamyatin* interjects notes into the silence left by Houle. Later Houle performs rhythmic sequences, and *Zamyatin* momentarily appears to play over these patterns with well-timed rhythmic interjections. The octopus of parameters can be heard throughout, with the tones produced by *Zamyatin* morphing over multiple parameters (such as pitch and loudness of synthesised tones, or the rate, grain size and grain playback position of granular sample players), with multiple coordinated voices coming in at the end.

Playing steady rhythms and modulating intensities can under the right circumstances produce effective couplings where the cause and effect is confused, and can create the impression of *Zamyatin* performing an anticipatory role. It is ambiguous as to whether Houle overrates the system's capacity to perform anticipatory actions, and I think it is reasonable to interpret this as a working model, or 'stance' towards the system, with which Houle can approach the improvisation with effective trust.

Houle goes on to discuss how he sees the system with respect to human improvisers and other systems (not named):

> "I don't think that the 'essence' of human improvisation is easily captured in software, as there are too many un-quantifiable elements at work, but a good system will approximate these elements by generating enough permutations of simple musical parameters. What I am experiencing these days is that not only is the technology becoming more powerful and sophisticated, but the composers/designers have greater skills at developing material that blur the lines between process and reality."
>
> "I found that *Zamyatin* was responding in a much more 'organic' way than most systems I've had the opportunity to work with. This is mostly due to the level of interactive parameters built into the system, I imagine."

This statement supports the process method used to create works with *Zamyatin*. The process of iterating between evolved dynamic models and synthesis and generative music components, as a creative process, made it easy to home in on organic-feeling forms of interaction, which may be harder to achieve working with more rigid and formal models. In this process, the author is never strictly in charge of how the system will behave, and this limitation had a strong impact on the creative decisions made.

> "After the initial 'feeling out' process I was able to predict certain responses from the machine, and to engage in a cat and mouse game, trying to outwit the machine. This lasted a few minutes before I disengaged and tried to tackle the musical elements from an improviser's perspective."

This suggest Houle could not precisely operate the system, but had a good enough model of how it would behave to experience some sense of control over the system. The reference to outwitting the machine has a weakly anthropomorphic sense, but perhaps more important is the establishment of a 'game' of control and anticipation played by the performer (not by the system). The reference to an improviser's perspective suggests that this earlier mode is not a standard mode of human improvisation, but is specific to performance with a computational system. Houle elaborates that the experience is comfortable but engaging:

> "My feeling was that I was able to go to my personal musical language without 'worrying' about how the system would respond, and to interact playfully with what was being offered. The only frustrating part was that some of the sonorities the system would generate were directly related to basic parameters such as attack, velocity, dynamics, and frequency range. However, there were other tones which were more difficult to pin down, as to what triggered them. I never felt frustrated, as there were plenty of ideas work with."

This relates to another point of tension for designers of improvising systems, to create a situation where the performer stops being overtly conscious of the system but acts 'naturally' with it, perhaps not treating it as a human improviser, with the associated expectations for interaction, but also not treating it as a mere backdrop to a solo performance. If Houle has understood the system correctly, then the last statement refers to the way that certain actions of the performer will trigger events somewhat too predictably, for example if there is silence then playing a note will likely trigger *Zamyatin* to play a note.

Houle does then explicitly appeal to a 'human-ness', and the attribution of an attitude, with specific musical qualities:

> "...it lends itself very well to the type of music I have an affinity towards. I don't particularly like to improvise over grooves, or simple rhythmic stratas. I like shiftiness, permutations, and sudden caesuras in music. *Zamyatin* seems to be built to behave with that kind of attitude. It could gain perhaps by adopting more 'predictable' rhythmic behaviour spread over longer time periods, but that would be at the expense of what it does best. I would also say that its 'human-ness' comes from this 'stream-of-consciousness' feel that you get from interacting with it."

Note that the 'human-ness' is expressed in terms that invoke the coupled nature of the interaction. The stream-of-consciousness is not attributed directly to *Zamyatin*, and could be interpreted as belonging to Houle or to the musical material itself.

He then adds:

> I strongly believe that a minimum of interaction is best in order to achieve a true state of musical spontaneity, where all participants have a clean slate to draw upon.

This appears to contrast with his initial statement about satisfying interaction, implying that interaction is conceived as occurring at least two different levels, or that interaction should be non-obvious, making space for autonomous actions.

Some of the interaction concepts discussed above can be picked up in the musical excerpts. In the fifth excerpt (link), at 5 minutes 19 seconds in, there is a section where *Zamyatin* holds a couple of long notes. The second is longer than the third, a long note that holds through Houle's faster jittery playing. Houle is able to play against this and ready himself for any change. This is an example where the system unexpectedly asserts stability against expectations that it would be more imminently reactive. The playful nature of this state, demonstrated by Houle's slightly exaggerated exit when *Zamyatin* does end the note, evokes the "cat-and-mouse game" Houle refers to.

A telling example of the inequality between system and performer that Houle refers to presents itself in another excerpt. At the end of a very coherent period of well-coupled playing immediately following the section described above (at 5:56), *Zamyatin* rounds off a phrase with three assertive, regularly spaced square-wave beeps which come as a total surprise given the previous sound palette. Houle immediately picks up this rhythmic cue and establishes a melodic pattern over the resulting pulse. *Zamyatin*, meanwhile, does nothing for some time, before coming in with something completely different. This is a conspicuous giveaway of *Zamyatin*'s musical incapacity, and also of Houle's capacity, covering for the system. But in this case Houle's quick adaptivity seems to expect too much of the system, it backfires and he is left to go it alone. Thus whilst the system does a good job of keeping things moving and prompting the player, the cycle of mutual influence is sometimes broken.

As Houle reflects,

> "It might not be able to think 'long term' but that has its advantages when you are working 'in the moment'. ... As *Zamyatin* is incapable of locking in to a moment, it allows the music to move forward constantly, that is if the human at the other end is ok with that! My reaction to this was to work in miniatures, stating an idea, toying with it for a bit, then look for a transformation process that would gradually push the machine to react differently to the given material."

Houle is right that *Zamyatin* is "incapable of locking into a moment", but one of the objectives of the design of the dynamic models is that they do lock in to specific states before switching. This remains a goal that has not yet been achieved.

6 Conclusion

This paper appeals to using practice-based and design research methods to return to a more cybernetics-inspired approach to creating interactive music systems. The design features of *Zamyatin* of establishing low-level sensitivity, oscillations and an octopus of parameters are proposed here as having some general applicability to the design of interactive music systems that are not explicitly designed to be human-like, but to create engaging musical interactions. The musician's interview and discussion of recordings have been used to try to indicate how these affordances of the interactive system relate to properties of the musician's performance. One way that these results can be used is to feed them back into the design of evolutionary targets that have been used to direct the evolution of *Zamyatin*'s dynamic models.

This process broadly supports the results from Bown (2015) that the performer readily approaches the system using abstract notions of musical interaction, without recourse to direct human comparison. This suggests a more fluid approach to stances, accommodating specific attitudes towards a synthetic agent; physical and design stances are present, and although there is no clear indication of an intentional stance, the system is at times characterised as exhibiting autonomy and anticipation, such as in the reference a game of cat-and-mouse. Aspects of an embodied interaction with the system are present throughout the discussion, such as Houle's reference to forgetting who or what he is performing with. The performer makes frequent references to the predictability of the system and the question of cause-and-effect, relating to the concept of G-autonomy, and to the taxonomy of Eigenfeldt et al. (2013b).

It remains challenging to derive precise design goals from the outcomes of this discussion and analysis, but it is this author's hope that the terms used can be further probed through iterations of this research process, leading a clearer model of the performer's approach and to more engaging interactive systems. The familiarity of the performer with the system also introduces ambiguity into the results, and more systematic studies looking at performers with differing levels of exposure to the system will be of interest. In future work it will also be of interest to apply more detailed models of how performers engage with an unfolding musical surface, such as the expectation-based models of Huron (2006) and Wiggins and Forth (2015). It may also be of interest to measure G-autonomy and G-causality in documented performances.

References

Ariza, C. (2009). The interrogator as critic: The turing test and the evaluation of generative music systems. *Computer Music Journal 33*(2), 48–70.

Assayag, G. and S. Dubnov (2004). Using factor oracles for machine improvisation. *Soft Computing-A Fusion of Foundations, Methodologies and Applications 8*(9), 604–610.

Banerji, R. (2012). Maxine's turing test - a player-program as co-ethnographer of socio-aesthetic interaction in improvised music. In *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment (AIIDE'12) Conference.*

Bilda, Z., E. Edmonds, and L. Candy (2008). Designing for creative engagement. *Design Studies 29*(6), 525–540.

Blackwell, T., O. Bown, and M. Young (2012). Live algorithms. In J. McCormack and M. d'Inverno (Eds.), *Computers and Creativity*. Springer-Verlag.

Blackwell, T. and M. Young (2004). Self-organised music. *Organised Sound 9*(2), 137–150.

Bown, O. (2011). Experiments in modular design for the creative composition of live algorithms. *Computer Music Journal 35*(3), 73–85.

Bown, O. (2015, June). Player responses to a live algorithm: Conceptualising computational creativity without recourse to human comparisons? In *Proceedings of the International Conference on Computational Creativity*, Park City, pp. 126–133.

Bown, O., A. Eigenfeldt, A. Martin, B. Carey, and P. Pasquier (2013). The musical metacreation weekend: Challenges arising from the live presentation of musically metacreative systems. In *Proceedings of the Musical Metacreation Workshop, AIIDE conference*, Boston.

Bown, O., A. Eldridge, and J. McCormack (2009). Understanding interaction in contemporary digital music: from instruments to behavioural objects. *Organised Sound 14*(02), 188–196.

Brown, A., T. Gifford, and B. Voltz (2013). Controlling interactive music performance (cim). In *Proceedings of the Fourth International Conference on Computational Creativity*, pp. 221.

Candy, L. and S. Ferguson (2014). Interactive experience, art and evaluation. In *Interactive Experience in the Digital Age*, pp. 1–10. Springer.

Cardoso, A., T. Veale, and G. A. Wiggins (2009). Converging on the divergent: The history (and future) of the international joint workshops in computational creativity. *AI Magazine 30*(3), 15.

Collins, N. (2006). *Towards Autonomous Agents for Live Computer Music: Realtime Machine Listening and Interactive Music Systems*. Ph. D. thesis, Centre for Science and Music, Faculty of Music, University of Cambridge.

Collins, N. (2016). Towards machine musicians who have listened to more music than us: Audio database-led algorithmic criticism for automatic composition and live concert systems. *Computers in Entertainment (CIE) 14*(3), 2.

Dennett, D. C. (1989). *The Intentional Stance*. MIT press.

Eigenfeldt, A., O. Bown, P. Pasquier, and A. Martin (2013a). Towards a taxonomy of musical metacreation: Reflections on the first musical metacreation weekend. In *Ninth Artificial Intelligence and Interactive Digital Entertainment Conference*.

Eigenfeldt, A., O. Bown, P. Pasquier, and A. Martin (2013b, October). Towards a taxonomy of musical metacreation: Reflections on the first musical metacreative weekend. In *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment (AIIDE'13) Conference*, Boston, MA.

Eigenfeldt, A. and P. Pasquier (2009). A realtime generative music system using autonomous melody, harmony, and rhythm agents. In *Proceedings of the 12th Generative Art Conference GA2009*.

Eldridge, A. (2007). *Collaborating with the behaving machine: simple adaptive dynamical systems for generative and interactive music*. Ph. D. thesis, University of Sussex.

French, R. M. (2000). The turing test: the first 50 years. *Trends in Cognitive Sciences 4*(3), 115–122.

Harnad, S. (1991). Other bodies, other minds: A machine incarnation of an old philosophical problem. *Minds and Machines 1*(1), 43–54.

Huron, D. (2006). *Sweet Anticipation*. MIT Press.

Jordanous, A. (2012). A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cognitive Computation 4*(3), 246–279.

Lamb, C., D. G. Brown, and C. Clarke (2015). Human competence in creativity evaluation. In *Proceedings of the Sixth International Conference on Computational Creativity June*, Volume 102.

Lewis, G. E. (2000). Too many notes: Computers, complexity and culture in voyager. *Leonardo Music Journal 10*, 33–39.

Malafouris, L. (2013). *How things shape the mind*. MIT Press.

Pachet, F. Roy, P. (2017). Do jazz musicians really interact? In P.-J. M. Micheline Lesaffre and e. Marc Leman (Eds.), *Routledge Companion to Embodied Music Interaction*. Routledge.

Papadopoulos, G. and G. A. Wiggins (1999). AI methods for algorithmic composition: a survey, a critical view and future prospects. In *Proceedings of the AISB'99 Symposium on Musical Creativity*, pp. 110–117. Brighton, UK: SSAISB.

Pasquier, P., A. Burnett, N. G. Thomas, J. B. Maxwell, A. Eigenfeldt, and T. Loughin (2016). Investigating listener bias against musical metacreativity. In *Proceedings of the Seventh International Conference on Computational Creativity*.

Pearce, M. T., D. Meredith, and G. A. Wiggins (2002). Motivations and methodologies for automation of the compositional process. *Musicæ Scientiæ 6*(2), 119–147.

Pease, A. and S. Colton (2011). On impact and evaluation in computational creativity: A discussion of the turing test and an alternative proposal. In *Proceedings of the AISB symposium on AI and Philosophy*.

Rowe, R. (2002). Machine musicianship. *Journal of the Acoustical Society of America 111*(5).

Seth, A. (2010). Measuring autonomy and emergence via granger causality. *Artificial Life 16*, 179–196.

Sturm, B. L. (2016). The "horse" inside: Seeking causes behind the behaviors of music content analysis systems. *Computers in Entertainment (CIE) 14*(2), 3.

Tomasello, M. and J. Call (1997). *Primate Cognition*. Oxford University Press.

Turing, A. M. (1950). Computing machinery and intelligence. *Mind 59*(236), 433–460.

Whitby, B. (1996). *The Turing Test: AI's Biggest Blind Alley?*, pp. 519–539. Oxford University Press.

Wiggins, G. A. and J. Forth (2015). Idyot: a computational theory of creativity as everyday reasoning from learned information. In *Computational Creativity Research: Towards Creative Machines*, pp. 127–148. Springer.

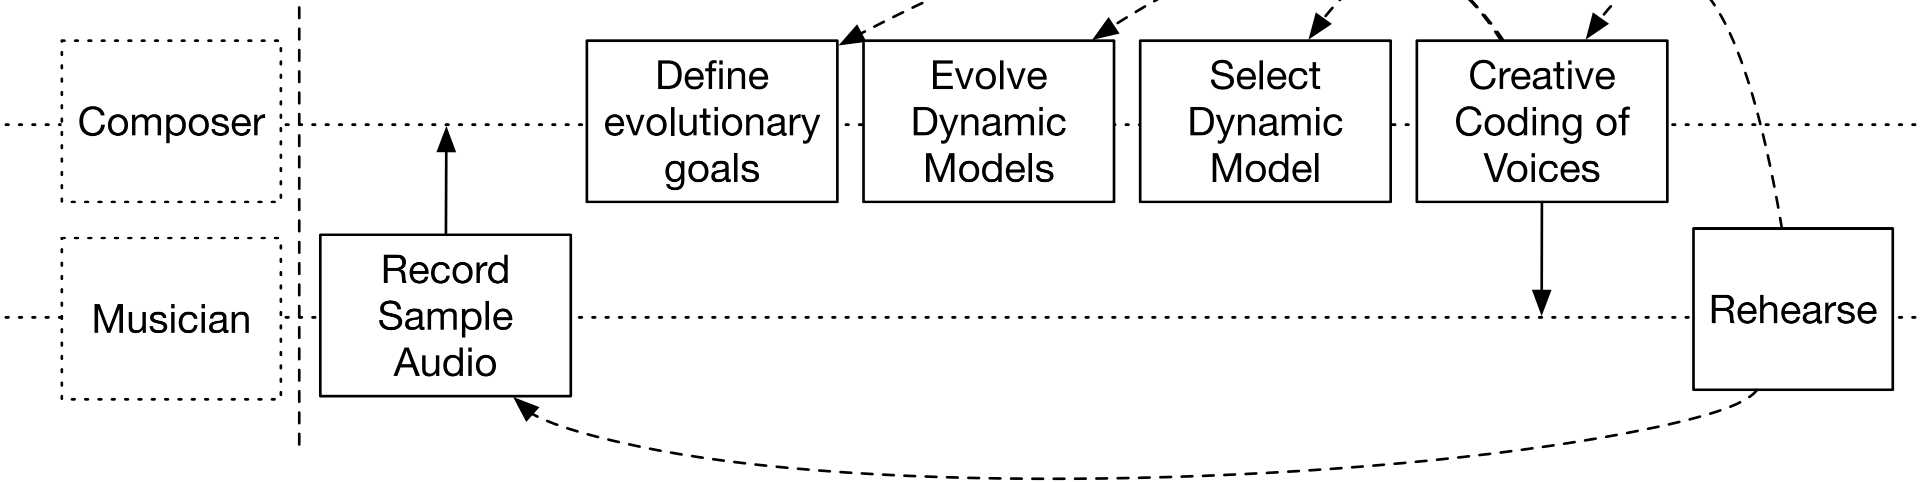
Composer
Musician
Record Sample Audio
Define evolutionary goals
Evolve Dynamic Models
Select Dynamic Model
Creative Coding of Voices
Rehearse

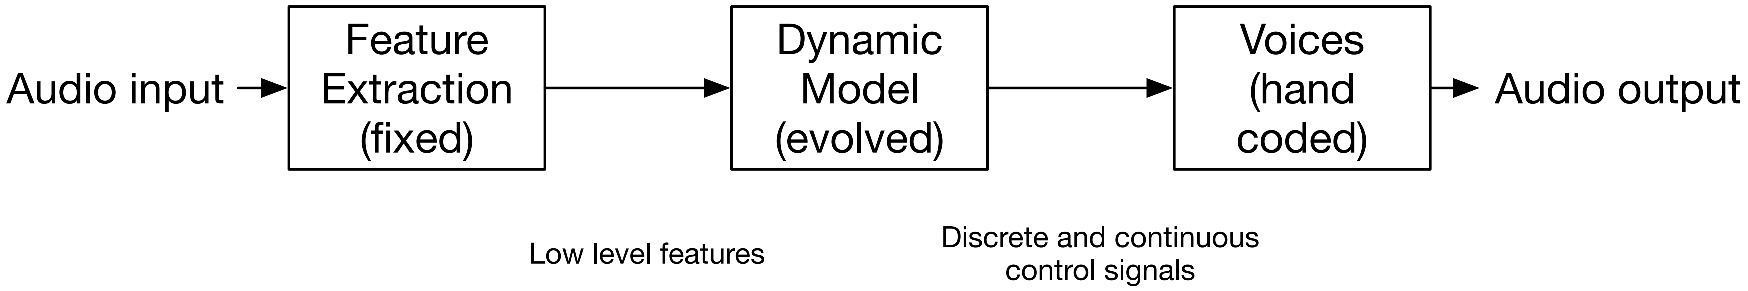
Audio input
Feature Extraction (fixed)
Dynamic Model (evolved)
Voices (hand coded)
Audio output
Low level features
Discrete and continuous control signals