When Children Reflect on Their Playing Style: Experiments with the Continuator and Children

FRANÇOIS PACHET, Sony CSL, Paris
AND
ANNA RITA ADDESSI, University of Bologna

The Continuator project concerns the study of interactive musical systems that are able to learn musical styles. The Continuator system was developed with the primary goal of designing new musical instruments in the context of improvised music. The strong subjective impressions created with the use of the system led us to consider its use in a pedagogical context. This article describes experiments conducted with the system and 3- to 5-year-old children. We highlight several dimensions of the study pertaining to music education, including attention span, spontaneous development of playing modes, and capacity to listen analytically. We describe very encouraging preliminary results and stress the importance of using reflective interactive systems for triggering musical interest in children and creating stimulating, nonsupervised music learning environments. We conclude by setting up our research in the context of the theory of flow as an optimal experience.

Categories and Subject Descriptors**:** H.5.5 [**Information Interfaces and Presentation**]: Sound and Music Computing; J.5 [**Computer Applications**]: Arts and Humanities; K.3.1 [**Computers and Education**]: Computer Uses in Education

General Terms: Design, Experimentation, Human Factors

Additional Key Words and Phrases: Music, continuator, children, music education, learning, improvised music, theory of flow, hyperinstruments

1 INTRODUCTION: THE CONTINUATOR PROJECT

The primary goal of the Continuator project is the design of new, interactive musical systems that eliminate the traditional dichotomy between interactivity and learning. There are many justifications for such a project. One trivial reason is to implement Ray Kurzweil's prophecy, which predicts that "Human musicians [will] routinely jam with cybernetic musicians" [Kurzweil 1999]. On a more serious level, this project may be seen as an extension of the work on so-called "hyperinstruments," developed in particular by Tod Machover [Machover 1996]. Hyperinstruments are defined generally as those in which some computer-based mapping is interposed between the actions of the player and the resulting sonic effects [Paradiso 1999]. In principle, such a mapping allows extension of the musical ability of the player in all sorts of ways. In practice, most of the proposed hyperinstruments have focused on control (particularly of timbre) and expressivity. In our context, we were less interested in extending expressivity than in extending the basic

Authors' current addresses: F. Pachel is with SONY CSL, Paris, France. A.R. Addessi is with the University of Bologna, Bologna, Italy.
Permission to make digital/hard copy of part of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date of appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.
© 2004 ACM 1544-3574/04/0100-ART05 $5.00

musical vocabulary of the user. Concretely, our experiments involve the use of a keyboard (considered by hyperinstrumentalists as a rather rudimentary controller). Thefocus of our system is the intelligent generation of musical phrases and structures with conventional expressive characteristics (basically the ones provided by MIDI). In other words, we address the *what* of music generation, rather than the *how*.

The Continuator project started in 2000 at the Sony Computer Science Laboratory in Paris. The initial version was a real-time MIDI-based system that was able to generate linear jazz improvisation from various MIDI inputs such as a MIDI guitar. Several versions were subsequently designed, and the Continuator system benefited from the enthusiastic input of the many musicians, researchers, and people with expertise in all areas who regularly visit Sony CSL. Indeed, initially, the main focus of the Continuator project was to design a system for adults, either beginners or professionals, and feedback from adults was systematically sought and analyzed.

The project had, however, unexpected ramifications in the domain of musical education for the early ages (3- to 5-year-olds). Many projects have been undertaken to enable young children to play music, with the goal of developing musical abilities early by designing musical instruments that are easier to play than conventional ones [MusicLab 2001; Weinberg 1999], or by developing tools that allow children to become instrument designers themselves [Resnick et al. 1996]. Many of the features we thought were exciting for professional musicians, such as the organic capacity of the system to learn musical styles agnostically and its ability to respond in real time proved just as exciting for nonmusicians and young children. In all cases, the main lesson learned from these experiments, and the main claim of this article, is that it is worthwhile to design and use a particular class of interactive systems (we call them reflective, as discussed below) for music education. These are systems in which the user, whatever his skills, competence level, and musical goals, is confronted with a developing mirror of himself. This unusual situation creates strong subjective feelings that we believe can be exploited to enhance musical experience and teach musical skills in general.

This article is organized as follows. First we describe in brief the mechanism behind the Continuator, which is based on an extended Markov model of musical phrases, and report on some experiments with professional jazz musicians. We then describe experiments conducted with children and draw some conclusions regarding the approach by recasting the whole project in the context of Csikszentmihalyi's psychological theory of flow.

2 TECHNICAL CONFIGURATION

The current architecture of the Continuator is very simple, consisting of one MIDI input (typically from a synthesizer) and one MIDI output (typically returning to the same synthesizer). Although there is a graphical interface to tune the many system parameters, its operation in the standard mode involves no interface other than the MIDI instrument itself. The user plays musical sequences of any kind. When the phrase is terminated (we will come back to this), the Continuator, in turn, generates a musical phrase in response. This musical phrase is characterized by its stylistic similarity to the phrases that have been played by the user so far. Technically, it is a continuation of the last input phrase, hence the system's name.

Although it is difficult to define the notion of *musical style* very precisely, we have adopted a technical notion of style, consisting of the statistical distribution of notes, chords, and musical elements in general as well as their ordering. The Continuator is based on a Markov model of musical phrases (see the next section), and the model of the style created by the system retains melodic patterns, harmonic progressions, dynamics, and, to a lesser extent, rhythmic patterns of the corpus used for learning. An important consequence of this approach is that the phrases generated by the Continuator are similar to but different from the phrases played by the user. This notion of similarity is of course crucial in music, both from an analytical viewpoint (most music analysis methods are based on some similarity principle (see, e.g., Deliège [2003]) and from a perceptive viewpoint (musical structures are perceived and organized according to similarity judgments of various sorts). The Continuator may therefore be seen as an engine for producing variations of arbitrary musical material.

2.1 Some Musical Examples

To illustrate the workings of the Continuator, several simple musical examples are given below. These examples are notated exactly as they are played, i.e., without rhythmic quantization. They show that the Continuator adapts quickly to arbitrary styles, and is able to generate musical material that "sounds like" the user input on a relatively small scale. Issues related to capturing higher-level structure are not discussed here, as they are not relevant for our purposes (see Pachet [2002a] for more details). More sophisticated examples of music created by the Continuator can be found on the web site of the first author.

The most important aspect of the Continuator is that the musical material that is generated always conforms stylistically to the input. Also (this is more difficult to illustrate graphically), the Continuator keeps on learning from whatever input it gets. As a consequence, the behavior of the system improves over time. If the user produces phrases that are stylistically consistent but unique, the Continuator will learn more faithfully, and will produce musical phrases that are increasingly accurate with respect to the musical style of the user.

Fig. 1. A simple melody (top staff) is continued by the Continuator in the same style.

Fig. 2. A simple chord sequence (top staff) is continued by the Continuator in the same style.

Fig. 3. A jazzy line (top) is continued by the Continuator (bottom).

2.2 Implementation and Design

There has been considerable research in artificial intelligence and information theory on the technical issue of learning a musical style automatically in an agnostic manner. In his seminal paper, Shannon [1948] introduced the concept of information based on the probability of an event's occurrence in communications (messages). Soon after, this notion was used to model musical styles, for instance by Brooks et al. [1957]. These early experiments showed that it is possible to create pieces of music that sound like given styles by simply computing and exploiting probabilities of note transitions. More precisely, given a corpus of musical material (typically musical scores or MIDI files), the basic idea is to analyze this corpus to compute transition probabilities between successive notes. New music can then be produced by generating notes using these inferred probability distributions. One of the most spectacular applications of Markov chains for the generation of music is probably by Cope [1996], although his musical results in their entirety were not produced automatically. A good survey of state-of-the-art Markov-based techniques for music can be found in Triviño-Rodriguez et al. [2001], in particular variable-length Markov models, which capture stylistic information more finely.

The Continuator system is yet another species in the world of musical Markov systems, although with novel features. In our context, we wanted to learn and imitate musical styles in a faithful and efficient manner and make the resulting mechanism usable as an actual musical instrument. This raised a number of technical issues, whose solutions were progressively integrated in the Continuator.

The architecture of the Continuator consists of two modules: an analysis module and a generator module. The analysis module takes MIDI sequences as input, played in real time. The system contains three main parts:

1. *A phrase-end detector* that is able to detect that a musical phrase had “ended. Detection is based on an adaptive temporal threshold mechanism. The threshold is inferred from the analysis of inter-onset intervals in the input sequence. As a result, if the input sequence is slow (or, rather, contains few notes per second) then the threshold is increased, otherwise it is decreased. This simple mechanism ensures that the continuation will be temporally seamless.
2. *A pattern analyzer*. Once detected as complete, the input sequences are sent toa pattern analyzer, which builds up a Markov model of the sequence. The complete algorithm is described in Pachet [2002b], and consists of a left-to-right parsing of the sequence to build a tree of all possible continuations for all possible prefixes of the sequence. To speed up learning, the system also learns all transpositions of the sequence.
3. *A global property analyzer*. Various global properties of the input sequence are also analyzed, such as the density (number of notes per second), the tempo, and the meter (location of strong/weak beats), the overall dynamics (loud or soft), and so on. These properties are used to produce a continuation that is musically seamless with the input.

The generator is responsible for producing the continuation of the input sequence. The actual production of the musical material exploits the Markov graph created by the analysis module [Pachet 2002b]. In essence, it consists of producing the continuation on a note-by-note basis. Each note is generated using the Markov probabilities inferred during the analysis stage. Technically, it uses a variable-order Markov generation that optimizes the relevance of each single note continuation by looking for the longest possible subsequence in the graph. Special care has been taken to perform meaningful segmentations of the input phrases for the learning phase. Indeed, real-world input phrases are never composed of perfectly successive notes or chords. In order to “cut” input phrases into chunks, which are then fed to the learning system, a segmentation process is able to detect note or chord transitions and possibly cut across unfinished notes. The module also stores the possible “residual” discrepancy, and restores it at generation phase so that the material retains the rhythmical “naturalness” of the original style.

This continuation sequence is, however, crude, in the sense that it does not necessarily have the global musical properties of the input sequence. Therefore, a mapping mechanism is applied to transform the brute continuation into a musical phrase that will be played just in time to produce seamlessness. Currently, the properties that are analyzed and mapped are tempo, metrical position, and dynamics (more details can be found in Pachet [2002a]).

2.3 Playing Modes

The basic playing mode of the Continuator is a particular kind of "taking turns" by the user and the system, which is determined by three principles:

1. Automatic detection of phrase endings. The Continuator detects phrase endings by using a (dynamic) temporal threshold (typically about 400 milliseconds). When a time lapse exceeds this threshold, the Continuator takes the lead and produces a musical phrase.
2. The duration of the phrase generated by the Continuator is parameterized; but in most cases the duration is set to be the same as the duration of the last input phrase.
3. Priority given to the user. If the user decides to play a phrase while the Continuator is still playing, then the system will stop and return to listening mode (and eventually apply principle 1 again).

These principles are set without explicitly telling the users. Actually, attempts made to explain them to users have resulted in more confusion than clarification [Pachet 2002c]. Experience with the system has shown that these rules are usually easily learned by the user in an implicit way – the behavior of the system is usually obvious, even to children.

3 SUBJECTIVE IMPACT ON PROFESSIONAL MUSICIANS

Apart from the technical issues related to system design and the evaluation of the "stylistic similarity" of the music produced by the system, it soon appeared that there was an important dimension of the project that was more difficult to describe in a standard scientific paper. This is the dimension of the user's *subjective impressions* when playing with the system or watching it in action.

Indeed, one impressive result of the experiments with professional musicians is that the very use of the system provoked intense subjective impressions and reactions. Rarely but occasionally, these reactions were negative. For instance, a composer from Ircam reacted aggressively against the system as soon as he understood the machine would create (compose) music, and would not even try it. His credo was that machines should only be used to do things humans could not do. Since the Continuator only generates music from human input, he considered the system uninteresting and did not wish to examine it further. Another musician stated *a priori* that the system was not interesting because it was not able to capture stylistic information after the user had only played one note. Indeed, it is a fact that the Continuator analyses transitions between notes or chords, and a one-note sequence does not contain any transitions; but the composer refused to forgive this limitation.

Apart from these two negative and ideological reactions, all the others were very positive. The Continuator captured the attention of the audience beyond the traditional "demo effect" of many computer music presentations. In particular, a systematic *Aha* effect was noticed [Pachet 2002b] for professional users as well as beginners (see Figures 4, 5, 6). In experimental psychology, the *Aha* effect characterizes sudden moments of realization, understanding or inspiration; it is also being studied in neuroscience, where neuronal activity corresponding to the Aha effect has been shown [Mogi 2003].

Exprimentation with the system invariably induces users to reflect on their own musical personalities. Bernard Lubat, a jazz pianist and drummer, at the forefront of progressive jazz in Europe, has played with the system many times. In particular during a performances at the Uzeste festival and a concert at Ircam in October 2002, he evoked with great precision his own evolution in improvisation, aided by the system, which allowed him to virtually "play ahead of his current thinking." György Kurtag Jr., a composer and improviser, described the Continuator as a kind of "amplifying mirror, and said that his regular use of the system throughout 2001 changed his way of improvising and composing. (This collaboration resulted in a composition, "The Hollow of the Deep Sea Wave," performed at the Vienna Festwochen 2002 music festival.)

Fig.4. Bernard Lubat listens to the Continuator during his performance; later, he "launches" the Continuator with a characteristic movement; then gestures in "pretend play" mode while listening to the music produced by the Continuator.

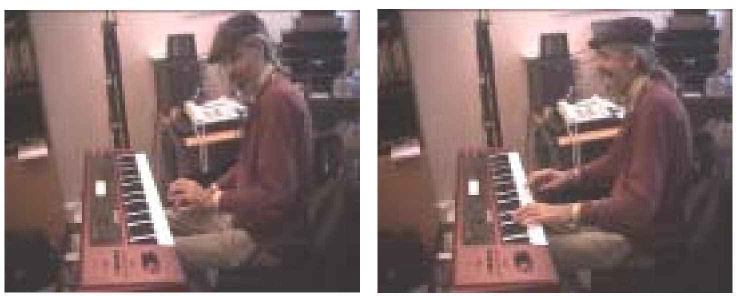

Fig. 5. Alan Silva playing with the Continuator.

Fig. 6. Bernard Lubat using the Continuator during improvisation sessions. As with Silva above, a show of skepticism, followed by an Aha effect after a few minutes experimentation with the system.

4 EXPERIMENTS WITH CHILDREN

4.1 Motivations

The idea of using the Continuator with very young children originated in a personal context (with the daughter of the first author, and then her friends). Like many parents wishing to expand their children's expressive ability as early as possible, we accumulated many toys related to musical performance. Examples of such toys are illustrated in Figure 1, and include approximations of real keyboards as well as various objects that produce sounds in many ways. Additionally, our home is full of musical instruments of many sorts, including a piano, numerous digital synthesizers, guitars, and so on.

None of these objects received any particular attention from this little girl, to the point that her parents even wondered whether she had any interest in music at all. All the attempts to play music for or with her did not trigger the slightest interest. She did not develop any apparent desire to play with either the toys or the instruments, and her attention was very limited in duration. She did not explore many playing modes, and her relation with these musical objects remained extremely passive.

Fig. 7. Three musical toys and two professional keyboards to which 3-year-old children may be exposed, often with limited success.

Fig. 8. A 3-year-old child playing with the Continuator at home.

When the first author started playing with his daughter and the Continuator (she was 3), her positive reaction was quite surprising and significant. She began to be interested in playing the keyboard – she would laugh at the system's "answers" and was able to focus her attention on playing music for longer periods of time.

4.2 Preliminary Experiments: the Continuator at School

To push the experiment further with additional children was therefore quite natural. In most schools, music is still taught using outdated methods, in which children are confronted with formalisms before they experience the joy of playing or listening to music. With a system like the Continuator, basic playing techniques may be learned more easily and earlier than with standard music education practice (piano lessons usually start at the age of 6 at the earliest in most conservatories). Most importantly, the Continuator – or in general a class of systems able to learn and react - could develop a genuine *desire for music* in children, and consequently prepare them for traditional classical training in a more productive way. The need for more fun and interactivity in the classroom has long been called for by various psychological studies [Webster 2002; Delalande 1984]. More precisely, our vision of education as a pleasing experience falls within the boundaries of studies on *flow* [Csikszentmihalyi 1990], as discussed in the last section of this article.

4.3 Preliminary Experiments

Preliminary experiments with children took place at a French kindergarten (Ecole Bossuet Notre-Dame in Paris). The goal was to test basic hypotheses about the effect of the Continuator on the playing abilities of 3- to 4-year-old children. More precisely, we made the following hypotheses:

1. *Increase in attention span*: Children will play longer with the Continuator than with a keyboard alone.
2. *Aha and surprise*: The Continuator will have noticeable Aha effects on children, as it will on professional musicians (children rarely find things surprising, a flying pianoforte would probably not surprise them).
3. *Autonomy*: The Continuator will motivate children to play music. For instance, a child, knowing that the Continuator is there, may express the intention of playing alone more often than with standard musical instruments.
4. *Exploration and playing modes*: The Continuator will push children to explore new playing modes. The results of the experiments with a normal piano suggest that children usually stick to single playing modes, including playing with only one finger, playing clusters with two hands, playing ascending or descending diatonic scales (i.e. white keys), and so on.
5. The Continuator will develop various kinds of *attachment behaviors* in children, similar to what has been observed with the Tamagotchi or Aibo [Kaplan 2001].

These preliminary experiments were made in one session; the protocol follows:

Each child was invited to play with a keyboard (a Korg Karma with a piano sound and no additional effects, connected to a pair of amplified loudspeakers). The Continuator was set to a mode that played phrases in the same style as the child, and approximately the same number of notes as the input phrase. The threshold for triggering the continuation was set to about 500 ms.

The child was alone with the first author in a classroom familiar to both. The protocol consists of two phases: The child was first told to play with the keyboard as he or she wanted, with no particular instruction. When the child stopped playing or expressed significant boredom, he or she was told that the system would now try to play with him or her. At that point, the Continuator would be turned on. The session stopped when the child stopped playing.

A total of eight children took part in the experiment, each was brought in one at a time by their teacher. The sessions were recorded with a video camera, and a number of interesting observations were made:

- There were indeed noticeable Aha effects. The children's reactions ranged from enhanced attention to surprise when the Continuator first started to play.
- The Continuator did seem to augment the attention span of most of the children. On two occasions, the sessions had to be terminated because they lasted longer than 40 minutes. There were two exceptions: one child was very tired and played very little, the presence of the Continuator made no difference; another seemed to enjoy playing the keyboard enormously, again the Continuator made no difference.
- The children who engaged in long interactions (one played 30 minutes and had to be stopped from continuing) also appeared to develop the ability to listen with great attention.
- The Continuator's interaction mode (stopping when the child plays, and playing when the child stops) encouraged the children to "take turns," without explicit directions from instructors to do so. The emergence of this behavior in such a context is not a trivial phenomenon, as shown for instance in the work on complex dynamic systems by Iizuka and Ikegami [2002].
- Some children exhibited a wide variety of playing modes. Apart from the classical playing modes (e.g. playing isolated notes, chords, arpeggios) the children invented new modes such as playing with their sleeves, kissing the keyboard, playing from behind, playing with the palms of their hands, etc. Some of them also theatrically accompanied the launching of the Continuator with gestures like raising their hands at the end of a musical phrase (see Section 4.5.3).

4.4 Systematic Experiments

Experiments involving the second author (at the University of Bologna) started in March 2003 to further investigate the issues brought to light in the first set of experiments. These experiments employed more systematic protocols (e.g., involving cross-experimentations), a larger number of children, and more precise hypotheses on the nature of the children's interactive behavior. The protocol was similar to the previous one, although more systematic. It involved not only single-child sessions, but sessions with two children also (with and without the Continuator). As in the first experiments, the children were 3- to 5-year-olds. A detailed analysis of these sessions from the perspective of experimental psychology is presently under way [Addessi and Pachet 2003]. Only the key observations are mentioned here.

4.5 Phenomena: For Adults and Children

Without committing too much to the precise analysis of the various sessions, we want to comment on a few remarkable phenomena. We describe them in the order of their appearance in a typical session, and were consistent for both adult musicians and children alike.

4.5.1 Surprise and the Aha Effect

The Aha effect, observed in professional musicians, was also seen systematically in children, and was manifested in a variety of facial expressions and gestures, both in single-child and two-children sessions. The term "Aha" is used in a somewhat biased way here to denote the fact that the children came to a sudden realization that the system was somehow trying to analyze and understand the children's own inputs and to speak their language.

One important point is that the Aha effect rarely reoccurred. After becoming used to the specifics of the interaction, children concentrated on other aspects of their musical relation with the system.

Fig. 9. Various expressions of surprise or "Aha" in the early stages of the interaction with the Continuator.

Fig. 10. Various expressions of musical excitement. Excitement was mostly provoked by listening to the system, rather than by actually producing music.

4.5.2 Excitement

Here, we separate excitement from surprise, in the sense that the surprise effect is most often short in duration, whereas the excitement phase lasts much longer, sometimes for 20 minutes or more. Excitement was observed in most of the cases. Interestingly, the children were excited mostly by what the system was playing, rather than by what they were doing. Figure 10 shows some expressions of this excitement.

4.5.3 Launching a Virtual Process

An interesting phase in the interaction occurs when children, after having mastered the basics of the system, abstract the concept of a “musical phrase.” This is indicated by a typical gesture demonstrating a pretend “launching” of their own musical phrase (as if it were a golf ball). This gesture determines the end of their musical phrase and also creates the expectation of the system’s response. It is remarkably similar to what professionals do themselves (see Figure 4, Bernard Lubat performing at Ircam).

4.5.4 Concentration and Listening

The ability to listen and concentrate for several seconds and listen to music is remarkable at this age (3 to 5). As pointed out by Rinaldi [2003], listening is probably one of the most important attributes for children to have in discovering the world around them. In Rinaldi’s view, children are researchers constantly making up theories about the world and evaluating them. With the Continuator, we observed this phenomenon rather

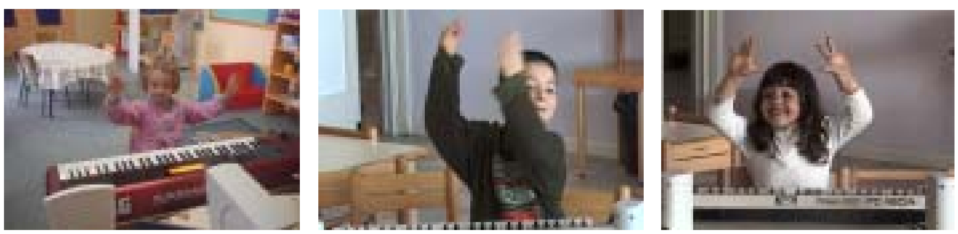

Fig. 11. Children raising their hands and virtually “ launching” the Continuator, after finishing a musical phrase.

Fig. 12. Various gestures showing listening and concentration.

Fig. 13. "*Aspetta*"**:** when one child forces the other to stop in order to listen to the machine.

systematically, i.e., children engaged in deep, concentrated listening to the effects of their playing on the system. Figure 12 shows several examples of this phenomenon.

4.5.5 Joint Attention

A typical situation encountered in sessions involving two children is that of joint attention. More precisely, one child would force the other to stop playing in order to listen to the system. This situation, which we call "*aspetta*" (the Italian word for "wait"), is illustrated in Figure 13. Of course, his behavior was not observed with professional musicians, who so far have experimented with the system as sole individuals.

5 DISCUSSION

The experiments conducted with the Continuator in schools have many implications for the use of interactive technologies in musical pedagogy. Of course, the pedagogical goals for three-to-five-year-olds cannot be expressed in terms of the standard musical curriculum (e.g., learning pitches, scales, chords, harmony, etc.). However, we show that experimenting with the Continuator in such "free" games does allow children to develop nontrivial musical abilities, including that of listening analytically, by experiencing surprise and excitement. We propose to interpret the Continuator's impact by using the theory of flow.

5.1 The Continuator as a Flow Machine

The Continuator appeared as a machine that promotes musical enjoyment in various forms. The system's ability to hold the attention of three-to-five-year-old children for long periods of time (remarkable for this age group) and in general to attract and hold the attention of users of all ages can be interpreted through the *theory of flow*, introduced by the psychologist M. Csikszentmihalyi [1990]. His notion of flow describes the optimal experience as one in which people obtain an ideal balance between skills and challenges. Two emotional states of mind are particularly stressed in this theory: *anxiety* when skills are clearly below the level necessary to meet the challenge, and *boredom* when the challenges are too easy for the skill level. Flow lies in the middle. Other states can also be described in terms of balance between skills and challenges (see Figure 17). We can think of the Continuator as a flow machine in the sense that by definition it produces a response corresponding to the skill level of the user. This approach also allows for the progressive scaffolding of complexity in the interaction, which is not the case for most pedagogical tools designed with a fixed pedagogical goal in mind.

More precisely, Csikszentmihalyi describes the state of flow as consisting of several fundamental traits; the balance between challenges and skills being in our case the most

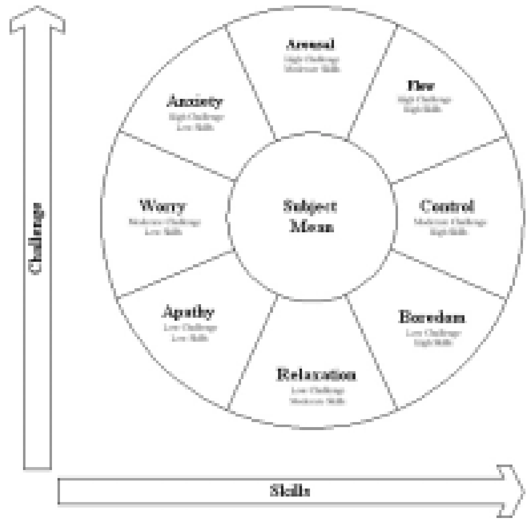

Fig. 14. Csikszentmihalyi's flow diagram describes emotional states such as boredom or anxiety according to the relation between skills and challenges.

important one. The other traits can be discussed in relation to our experiments as follows:

- *Focused attention*. The experiments clearly show that children are engaged in focused activity while both playing and listening.
- *Ease of concentration*. This is particularly clear given the fact that no instructions whatsoever are given to the children. They play with the system in a self-motivated way, without any external constraints.
- *Clear-cut feedback*. The Continuator produces clear feedback (in fact this is the only thing it does). In some sense the interaction is reduced to the analysis of the feedback produced by the machine.
- *Control of the situation*. Children are in control of the situation most of the time. They quickly understand that they can interrupt the system whenever they want. The limitations on control are due to the difficulties that may arise when interpreting some of the system's outputs (see an example in the next section).
- *Intrinsic motivation*. The most striking result of the experiments is related to the children's intrinsic motivation, who were not told anything about the rules of the system.
- *Excitement*. Excitement is clearly shown most of the time, in particular in the early phases of the sessions.
- *Change in the perception of time and speed*. A systematic study on interaction times is underway, but it is already clear that at least for some of the children time passed very quickly: some sessions had to be terminated by the experimenters when the time limit was reached.

There is, however, one flow characteristic that does not apply directly to the Continuator experiments:

- *Clear goals*. No explicit goal was given to the children, except to play until they were bored. Indeed, improvisation is generally not goal-oriented. It could be argued, however, that children did create goals spontaneously during their interactions with the system. For instance, several sessions involved children trying to push the system to replicate a particular frantic musical style that they had played some minutes ago.

5.2 Phases to Reach Flow

In order to precisely define the role and importance of the Continuator in developing musical abilities, the phenomena described above deserve careful analysis on an individual basis (e.g., attention span, listening and concentrating, etc.). However, It is also worth studying the "life cycle" of these phenomena, or, in other words, the progression of these flow traits over time. We observed that they always progress in the same order (see Figure 14).

5.3 Other Remarkable Observations

During the analytical stage (i.e., after the initial periods of surprise and excitement), there were clear indications of many other behaviors. For instance, several children developed spontaneously innovative playing modes. Besides rediscovering standard modes such as playing individual notes, arpeggios or chords, they would sometimes produce remarkable arpeggios and clusters, as well as playing with elbows, with hands in back, with lips, and so on.

Another phenomenon worth mentioning – and probably studying in more depth – is the children's ability to take turns in a spontaneous way. As mentioned before, the "rule of the game" (in this case a particular kind of switching of master/slave roles) was not explained to the children. But somehow they learned it implicitly and extremely quickly, after only a few interactions.

The ability to learn the rule of taking turns is nontrivial. In particular, a child who knows how it works must also understand the beginnings and endings of musical phrases, his own as well as the system's. A few cases show that skill in ascribing an intentional ending to phrases generated by the system develops very early in the sessions. Figure 15 shows a typical scenario in which a child is confronted with an interesting situation

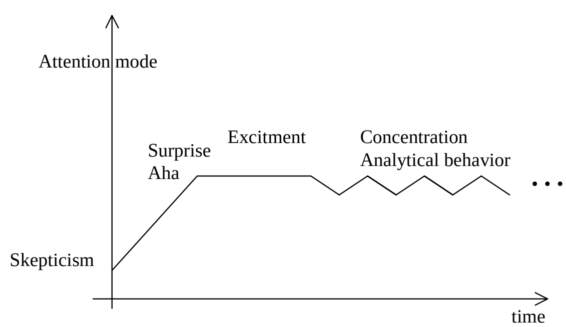

Fig. 15. A tentative sketch of the "life cycle" of the interaction mode with the Continuator.

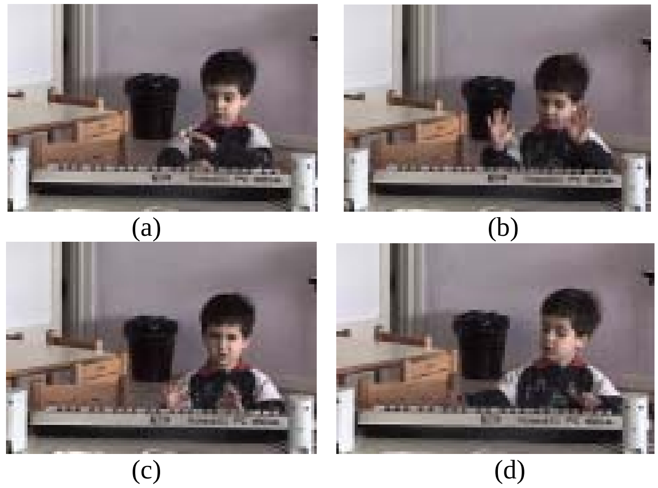

Fig. 16. A child clearly showing his interpretation of the Continuator's phrase endings: (a) shows him just finishing a phrase and waiting for the answer. The answer generated by the Continuator turns out to contain the last part of the child's input phrase followed by some additional notes. At the end of the first part of the continuation (i.e., repetition of the child's ending phrase) the child gets ready to play again (b), assuming that the phrase played by the Continuator has finished. It has not, and the child shows his misunderstanding by his facial expression (c). Eventually, the Continuator ends its phrase, and the child resumes playing (d).

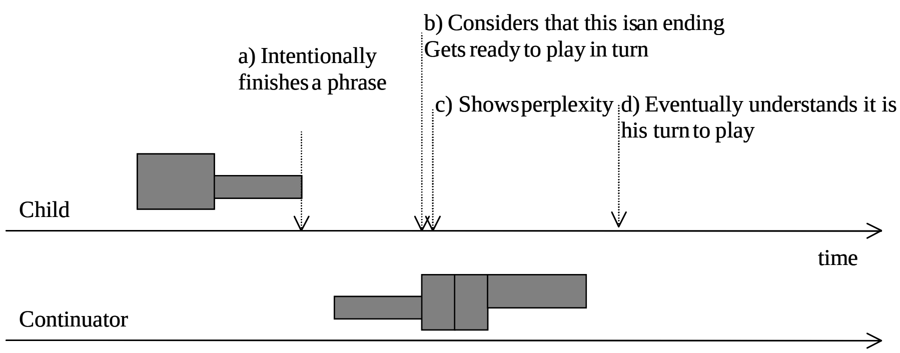

Fig. 17. A short scenario demonstrating the complexity of interpreting phrase endings. The phrase structure is represented schematically using rectangle of various shapes. The important aspect of the scenario is that the last chunk of the input phrase is the first one of the continuation, thereby creating a false sense of an ending for the child.

regarding phrase endings. In this case, a particular continuation by the system began with the ending of the child's phrase followed by additional notes. This situation, although rare, can occur because the Continuator has no particular notion of beginnings or endings. Figure 16 shows this graphically. The phrase played by the child is represented schematically by two consecutive chunks that last a few seconds. The last chunk is repeated by the Continuator, and is followed by another one.

5.4 Reflective Systems

The issue of what makes interactive software really usable has been a matter of investigation for a long time. The notion that an ideal interactive system should create a sense of intimacy with the user was pointed out by, for instance, Sidney Fels [Fels 2002], in his analysis of several interactive systems designed to create novel aesthetic experiences. He argues that intimacy, by creating the sensation that the system is an extension of the user's body, assures the efficiency of the user's actions.

The Continuator is characterized by another interactive dimension, related to the fact that it has a learning component. Because it is able to learn and imitate the user's musical personality, the Continuator acts as a dynamic mirror, and we assert that most of the interesting behaviors in our experiments stem from this feature. As such, the Continuator is only one instance of a larger class of systems that we call "reflective," i.e. systems in which users can interact with virtual copies of themselves, or at least agents that have a mimetic capacity and can evolve in an organic fashion. There are a number of characteristics that can be pointed to in defining reflexivity in interactive systems. We propose the following list, by no means exhaustive or even prescriptive, but to be taken as a starting point:

- *Similarity or the mirroring effect*. What the system produces "looks" like what the user is able to produce.
- *Agnosticity*. The system's ability to reproduce the user's personality is learned automatically and agnostically, i.e., without human intervention. In our case for instance, no preprogrammed musical information whatsoever was given to the system.
- *Incrementality*. Interactive systems are not only designed for short demos. Because the user is constantly interpreting the system output and altering his playing in response, it is important to consider the longer-term behavior of the system. Incremental learning ensures that the system keeps evolving all the time. Each interaction with the system contributes to changing its future behavior. Incremental learning is a way of endowing the system with an organic feel, typical of open, natural systems (as opposed to preprogrammed, closed-world systems).
- *Seamlessness*. The system produces material that is virtually indistinguishable from the user's input. This characteristic does not apply in the case of "classic" hyper-instruments, where the sonic effects are produced entirely through the system, and therefore do not result in any material directly produced by humans.

One important consequence of reflective systems is that the "center of attention" in the interaction is not the *end product* (the music), but the *subject* engaged in the interaction. Thus engaging in an interaction with a reflective system is a means of discovering oneself, or at least exploring one's ability in a domain (in our case musical improvisation). The natural, deep interest in self-discovery, particularly during the early years of childhood, may explain why the Continuator is so appealing. So generalization to a whole class of reflective systems could be productive if we design other successful instances of this class using the same principles.

Implementing reflective systems is now possible using artificial intelligence and machine-learning techniques and real-time architectures. In some sense, these systems are an extension of the "second self" [Turkle 1984], where not only does the machine seem to "think," but also thinks like the user. An interesting consequence is the reversal of roles: the student becomes the teacher and the user teaches the machine about himself.

6 CONCLUSION

Experiments with the Continuator and 3-to-5-year-old children show that many improvements can be made to enhance the musical capabilities of young children. Systematic analysis of these experiments should allow us to draw scientific conclusions about the various issues raised in this article, in particular on attention span. However, these experiments are sufficiently convincing to merit further investigation of the subjective effects of the Continuator as well, and more generally the effects of interactive musical systems on young children. The global goal is to exploit the attachment effect of certain interactive systems to improve music education.

To further this aim, a project is under way at schools in several European countries to expand the issues, scale-up the endeavor, and integrate social and cultural factors in the experiments. Another topic worth exploring includes the influence of the instrumental interface on the playing style. Indeed, each interface constrains style in a particular way. For instance, a beginner on a keyboard will tend to play with the white keys only (scale of C); whereas a beginner on a guitar would probably strum the instrument, producing mainly "open string" notes (E, A, D, G, B, E).

The main difference between the Continuator and classic edutainment environments is that musical goals are not stated explicitly in the Continuator. From the "flow" viewpoint, this is a limitation, or at least a point of divergence, as "clear goals" are a key ingredient in achieving flow. However, the lack of precise goals allows the interactions to grow in complexity, something that is not common in standard, fixed-goal approaches.

Furthermore, these experiments provide many ideas for developing innovative musical games in the style of the Continuator. Elaborating on observations concerning turn-taking, we plan to implement and test software games in which the goal is for children to infer the playing rules rather than develop specific musical abilities. Understanding the rules of a game is as challenging and musically relevant as actually trying to follow them. Pushing this idea further, we are now working on a system that is able to invent rules in real-time, and conversely interpret and adapt to rules created by the user, that is, a system able to demonstrate the same flexibility that good music teachers show when interacting with students.

REFERENCES

ADDESSI, A.-R. AND PACHET, F. 2003. Musical style replication in 3/5 years old children: experiments and analysis. Submitted to the *British J. Music Education.*

BROOKS, HOPKINS, NEUMANN, AND WRIGHT. 1957. An experiment in musical composition. *IRE Trans. Electronic Comput. 6*, 1(1957).

COPE, D. 1996. *Experiments in Musical Intelligence*. A-R Editions, Madison, WI, 1996.

CSIKSZENTMIHALYI, M. 1990. *Flow, the Psychology of optimal Experience.* Harper & Row, 1990.

CSIKSZENTMIHALYI, M. 1996. *Creativity: Flow and the Psychology of Discovery and Invention.* Harper Collins, 1996.

DELALANDE, F. 1984. *La musique est un jeu d'enfant.* Buchet/Chastel, Paris. 1984.

DELIÈGE, I. 2003. Special issue on music similarity. *Musica Scientae* (Fall 2003).

FELS, S. 2000. Intimacy and embodiment: Implications for art and technology. *In Proceedings of the ACM Conference on Multimedia* (Marina del Sol, CA, Oct. 2000). ACM New York, 13-16.

IIZUKA, H. AND IKEGAMI, T. 2002. Simulating turn-taking behaviours with coupled dynamical recognizers. In *Proceedings of the 8th International Conference on the Simulation and Synthesis of Living Systems* (University of New South Wales, Sydney, NSW, Australia , 9th-13th Dec. 2002).

KAPLAN, F. 2001. Artificial attachment: Will a robot ever pass Ainsworth's strange situation test?. In *Proceedings of Humanoids 2001: IEEE-RAS International Conference on Humanoid Robots.* S.Hashimoto, ed. 2001. 125-132.

KURZWEIL, R. 1999. *The Age of Spiritual Machines*. Viking Press, 1999.

MACHOVER, T. 1996. Brain opera, In *Memesis: The Future of Evolution*. Ars Electronica Editions, Linz, Austria.

MOGI, K. 2003. Onceness. In *The Future of Learning*. M. Tokoro and L. Steels, eds. IOS Press.

MUSICLAB. 2002. 6 interactive music applications for music teaching in national education. http://www.ircam.fr/produits/technologies/multimedia/musiclab-e.html.

PACHET, F. 2002a The Continuator: Musical interaction with style. In *Proceedings of the ICMC Conference* (Sept. 2002), 211-218. Extended version in *J. New Music Research 31,* 4 (2003).

PACHET, F. 2002b. Interacting with a musical learning system: The Continuator. In *Music and Artificial Intelligence*. C. Anagnostopoulou et al. eds. Lecture Notes in Artificial Intelligence vol. 2445, Springer Verlag, 2002. 119-132.

PACHET, F. 2002c. Playing with virtual musicians: The Continuator in practice. *IEEE Multimedia* (April-June 2002).

PARADISO, J. 1999. The brain opera technology: New instruments and gestural sensors for musical interaction and performance. *J. New Music Research 28*, 2 (1999), 130-149.

RESNICK, M., BRUCKMAN, A., AND MARTIN, F. 1996. Pianos not stereos: Creating computational construction kits. *Interactions 3*, 6 (Sept./Oct. 1996).

RINALDI, C. 2003. The joys of preschool learning. In *The Future of Learning: Issues and Prospects*. M. Tokoro and L. Steels eds. IOS Press, 2003.

SHANNON, C. E. 1948. A mathematical theory of communication. *Bell System Tech. J. 27* (July and Oct.1948), 379-423 and 623-656.

TRIVIÑO-RODRIGUEZ, J. L. AND MORALES-BUENO, R. 2001. Using multiattribute prediction suffix graphs to predict and generate music. *Computer Music J. 25,* 3 (2001), 62-79.

TURKLE, S. 1984. *The Second Self: Computers and the Human Spirit*. Simon & Schuster, New York.

WEINBERG, G. 1999. Expressive digital musical instruments for children. M.S. thesis. MIT Media Laboratory, 1999.

WEBSTER, P. R. 2002. Computer-based technology and music teaching and learning. In *The New Handbook of Research on Music Teaching and Learning*. R. Cowell and C. Richardson eds. Oxford University Press, 416-439.