FERIN MARTINO'S TOUR: LESSONS FROM ADAPTING THE SAME ALGORITHMIC ART INSTALLATION TO DIFFERENT VENUES AND PLATFORMS

Jeffrey M. Morris
Texas A&M University
morris@tamu.edu

ABSTRACT

Ferin Martino is a piano-playing software algorithm created by Jeff Morris for art installations. It uses nearby motion observed with a camera to shape, but not dictate, its flow of musical decisions. Aesthetically, the work challenges the notions of the composer and the composition by presenting a software program that composes its own oeuvre in such a way that visitors cannot experience the composition without also influencing it. The installation has taken many forms, at times including multiple cameras and speakers, video display, note input by the visitor, a digital player piano, and an outdoor venue with an original sculpture embedded in nature. This algorithm has also proven useful in live performance and computer-aided composition. Case studies of exhibiting this work illustrate the process of finding the most effective venue for an interactive art installation and the process of tuning interactivity for a given venue.

1. INTRODUCTION

Ferin Martino is a piano-playing computer algorithm influenced by motion around it. motion of the viewers, seen by the computer's video camera, influences the intensity of the music created. Like David Cope's algorithm Emily Howell, [1] the striking character of the is algorithm inspired me to give it an anthropomorphic name even though this was not the original intent. Usually presented with the title, *The Collected Solo Piano Works of Ferin Martino, as Conjured by Your Presence*, this work creates a situation that lets us reflect on the ontological nature of music: this music cannot be heard without the audience causing changes in the composition: it is impossible to hear these "collected works" in "unadulterated" form; by listening, you are shaping the music. By extension, it offers a chance to reflect on the way that any composer's music only has its existence in the minds of its audiences, and that the modes of it existence may be as diverse as its listeners. This is an idea suggested by literary theorist Umberto Eco in *The Open Work* [2]. This work consists entirely of software and can generate new material practically indefinitely. The fact that the code fits on one screen indicates the elegance of the approach to generating endless music with pleasing results. This work is an example of something uncommon in technology-based art: it is acoustic computer music. The resulting music is acoustic or synthesized piano, in a capricious, expressive musical style that will not be overbearing for public spaces.

Ferin was not designed to be a "know-it-all" algorithm. Instead, it sets aside sophisticated computer vision to explore hat aesthetically-valuable subtleties can be captured and expressed by simply measuring the amount of motion seen between video frames. It initially presents itself as pure non-interactive spectacle. Given the chance, it will show the attentive visitor that motions can cause musical events to start or stop. The curious visitor is rewarded by discovering the seemingly playful nature of this interaction. The software does not simply convert actions into musical notes, but uses motion to disrupt its "train of thought." The result is a playful interaction partner that may tire of overactive visitors attempting to control it and may begin to please itself for a while. At other times, it may seem to amplify visitor's gestures or play rolling accompaniment instead of maintaining a predictable one-to-one behavior. Whereas it would seem to establish a direct link between motion and music, the experience turns out to be more complex. Since the video is essentially reduced to a single pixel for evaluation (the camera's view is not divided into "hot zones"), many gestures may achieve the same amount of overall motion, and the software is equally sensitive to unintended motions as well as clothing and background colors and changes in lighting (meaning it is somewhat weather-sensitive near windows). Because it uses a simple video camera, positioning the work in view of a window means that it can respond to motion inside and outside a given venue.

Initial experiences running the software in the background and monitoring it as I went about regular office work was quite rewarding. Because this work uses cameras to watch public areas, it can also be seen in relation to surveillance art. In this case, it presents itself as friendly and pleasant, but it can still be unnerving for a visitor to realize how their actions can be watched.

I've found the need to adjust the algorithm to best fit the traffic/motion patterns, distance from the traffic, and lighting, making its behavior in each venue unique. Its varied accompaniment to even mundane rituals can promote

mindfulness throughout extended experiences in which visitors coexist in parallel with the software. As shown in the cases below, the addition of a video display, trackpad input, and digital player piano make it a more striking multisensory experience.

2. ABOUT THE SOFTWARE

Inspired by the free atonal expressionist music of Arnold Schoenberg, musical decisions begin with a choice between a single note treble melody and a three-note harmony below the melody. Pitch intervals are determined by a drunk walk, such that chord pitches are equally spaced according to that interval and are separated from the most recent melody pitch by that interval. A metronome emits a steady (but changeable) pulse triggering new note decisions. Video is converted into monochrome, the absolute differences of each pixel with the corresponding pixel in the previous frame are summed to one value representing overall motion, and that value scales the rate of the decision-triggering metronome as well as the loudness of the playing. This way, more motion yields more active music without giving the visitor direct one-to-one control over note decisions. [3]

Periodically, the software alternates between trigger by video and trigger by a running buffer of recent pitch choices. This loops and recursively shifts pitches to create florid musical sequences. The non-interactive aspect of this feature actually contributes to the appearance of organically developing musical motifs, because it is rare that a whole musical phrase would be captured in this buffer, more commonly just the tail of it that spins off in one direction or the other. It also contributes to the effect pf playfulness because it simply ignores the video at times, unannounced. Finally, it can be very rewarding to visitors who provide note input to it, as it may incorporate and elaborate their musical contributions.

For venues using a video display, the frame-differenced video is displayed with a blurring motion trail. This is helpful in getting visitors to interact with the software, because they can see what the software sees and where its attention flows. See Fig. 1.

3. CROWDED EXHIBIT HALL WITH HEADPHONES AND DISPLAY

The Triennale di Milano is a design and art museum in Milan, Italy, and since 1933, host of the triennial exhibition of the same name. [4] Ferin was positioned in an exhibition room where the 2013 Generative Art International Conference was held. Presented 'shoulder to shoulder" with other works, it was necessary to use headphones. This, along with the compact arrangement of the venue, allowed for minimal movement, mostly waving. This tended to put visitors into a mindset of viewing paintings or a gallery of circus oddities: each is always already doing what it is supposed to do, and one feels as thought one might get "the point" of the work after only a few seconds. This presentation format made it uncomfortable for any visitors inclined to longer encounters to get to know the work better.

Figure 1. Camera input, processed and ready for analysis. This is what Ferin can see. The bright horizontal sweeping line represents the visitor moving a hand quickly from side to side.

4. MAIN LOBBY WITH SPEAKERS AND DISPLAY

The Onassis Cultural Center in Athens, Greece is a performing arts complex build in 2010 dedicated to promoting innovation in the arts. [5] Ferin was exhibited in the main lobby, facing the main entrance, as part of the 2014 joint Sound and Music Conference and International Computer Music Conference, with approximately 300 artist-scholars in attendance. It took the form of a laptop computer (using its built-in camera and display) and stereo speakers, all fit compactly into a 2 ft $\times$ 2 ft footprint encased in a custom wooden stand covered in black fabric.

Positioned against the back wall, the camera had a view of most of the lobby area, including an information desk used to register and assist conference participants; a seating area where participants met, talked, purchased beverages, and awaited scheduled performances; and heavy traffic on Syngrou Avenue outside the front door. To access the elevators, stairs, and concert venues, visitors walked directly up to the installation and continued around the left or right side of the building. Speakers were audible within the lobby, around the sides of the building, and upward one or two floors (before being drowned out by installations on other floors). It was not audible inside the concert venues.

This provided a steady range of motion seen by the camera (caused by cars or people), so the algorithm rarely normalized to the hyperactive level of video noise within the camera. This kept it from becoming too much of a distraction for a public meeting space. The experience was most effective for visitors on their way out of the building, because the connection between their movements and the music became quickly obvious as they suddenly entered the camera's view then paused to look at the screen. This experience was more subtle for visitors entering the space, but it would respond to their gait as they made their way into the building. For these visitors, it made a subtle shift away from indifferent background music, gradually making noticeable connections with motions in the room.

In times when the lobby was empty, the software would

normalize to the level of its video noise and become frenetic. This functioned nicely to gain the attention of people in nearby spaces; they knew something was happening over there without having to be there. Then, when they entered the camera's view and the camera normalized to their movement, the energetic burst followed by silence had the effect of a shy performer who didn't realize he had an audience.

This remains the most effective presentation scenario for the software alone, without user controls.

5. SIDE LOBBY WITH DIGITAL PLAYER PIANO

The Society for Electroacoustic Music in the United States (SEAMUS) 2015 conference was held in the Moss Arts Center, at Virginia Tech (Virginia Polytechnic Institute and State University). Ferin was placed in the atrium outside one of the concert venues. [6] This was a quiet, low-traffic area, so I felt a need to restrict its maximum activity so it wouldn't be too frenetic for the peaceful space. It proved to be a favorite lounge for some participants to relax. Since it was at a music conference, the addition of the Yamaha Disklavier digital player piano made engagement very rewarding for those that ventured to duet with the software. Besides the rewarding effect of hearing the software mirror and embellish one's own playing, the "struggle" over the keyboard is an especially exciting form of engagement. The visitor, not knowing when or what the software will play, is drawn into a hyper alert improvisation mindset involving intense critical listening as well as continually adjusted plans on when and what to play next. At times, visitors have noted that the software played notes just as the visitor was reaching for them (simultaneously frustrating and gratifying); other times, the software would operate in parallel accompaniment or musical dialogue with the visitor at the opposite end of the keyboard's range.

Even for viewers who didn't play the piano, the effect of watching a human interface (the keyboard) play automatically can evoke the image of a ghost playing the keyboard. They also got the thrill of watching the piano-playing visitors essentially become part of the artwork.

6. OUTDOORS WITH MULTIPLE CAMERAS AND SPEAKERS, WITHOUT DISPLAY

I-Park is a sculpture garden in East Haddam, Connecticut, focusing on ephemeral artworks engaging with nature. [7] For this exhibition, I cleared a circular path in an undeveloped part of the land. Fortuitous developments with other works on the property led me to install a sculpture built from parts of a grand piano and two upright pianos that had been left outdoors for years. I arranged them to resemble the carcass of a large mysterious beast and placed speakers in footlights surrounding it, to make it look like the overgrown site of a forgotten spectacle. Because of the circular path, I embedded four cameras in the sculpture, in order to capture motion in all directions from a suitable distance. I also use opportunity to develop a weather-proof housing for the computer and other electronics.

In order to more fully engage with the site, I incorporated another performing algorithm of mine that captured sounds from the environment, mixed them with the piano sound, and folded them together into rich contrapuntal textures. I originally intended to use plant and tree motion to influence Ferin, but there was a surprising dearth of wind at the time. As the demonstration video shows, this results in a very different character during the night compared to the day. It is busy at night because of increased wildlife activity as well as the software normalizing to the low light levels, in addition to the cameras' infrared illuminators changing the range of things that can be seen. During the day, it plays mildly as visitors walk into the area, exposed to the bright sun, but tended to "stare back" as visitors stopped to ponder the sculpture. It did tend to become more responsive to subtle motions of the visitors, like shifting their weight to one leg, repositioning arms, or turning heads.

For video documentation of this exhibition, see [8].

7. IMPROVISATION WITH SAXOPHONE

In a Disklavier-themed concert at Texas A&M University titled Hot-Wired Piano, I performed with Ferin and colleague Jayson Beaster Jones playing tenor saxophone. It was amusing to reflect on whether this was a duo or trio. Since I was there interacting with the software but not making any sound directly, this trio consisted of only two humans and two instruments, but one of the instruments wasn't being played by either of the humans. In a sense, the saxophone's sound was also (at times) influencing the software, because I replaced the keyboard input with pitch recognition.

At first, having simply replaced the keyboard input with pitch recognition, I quickly discovered that the pitch detection was so accurate that the algorithm would too-often play in unison with the saxophonist. While this was not a musical goal, and it wasn't conducive to building counterpoint, it also starkly highlighted expressive differences in intonation between the saxophone and the piano, which yielded undesirable sonic effects. In response, I added a gradually changing pitch offset between the pitch recognition and the note-triggering parts of the algorithm, and I created a key control that would momentarily unmute the audio input, so that I could have the algorithm listen to the saxophonist only at strategic moments.

With this unmute control in addition to the camera input, I functioned more like a conductor, musically flowing my hands in front of the camera to shape Ferin's musical flow. This type and range of motion felt most comfortable to my human physiology. It did loosely allow me to draw on my rudimentary conducting skills, even though the gestures bore no particular meaning to the software. Rather, my conducting experience informed my movements as a limited kind of dance: patterns that allowed me to sustain certain levels of continuous motion when desired and patterns to organize space to make room for quick, disjunct gestures.

For the audience, this experience charged the confusion and (hopefully) curiosity about which sounds and motions

were causing which musical results, as well as tracing the life of some events as the are fed back in loops, changing between sounds and movements, influencing and sometimes echoing or elaborating on each other.

8. COMPUTER-ASSISTED COMPOSITION

Ferin has also been used successfully in computer-assisted composition [9]. In most cases, this workflow has looked much like a "duo" version of the live performance scenario, with me "conducting" in front of the camera and playing occasional brief motives on the keyboard to influence the software in desired directions. The latter has especially been useful in building large-scale form, since as of yet, the software has no long term memory. Left alone, motives only stay active in the music to the extent they happen to be captured in and played from the rolling buffer.

A computer-aided composition work session might typically involve a few "takes" like this and then turn to the editing stage. Given the software's capricious rubato playing, automated transcription from MIDI to music notation is not straight-forward. It is helpful to create a MIDI reference track in a MIDI sequencer to accompany the recorded performance, by tapping out the basic metronome pulse where it should occur in the music, following the tempo as it changes. However, when a sequencer reinterprets the performance by mapping it to the MIDI reference track, all tempo information is lost from the resulting music notation, and needs to be manually interpreted into standard terms such as *allegro* and *ritardando*. Similar work us necessary for its rapid expressive dynamics as well.

This juncture, buried in tedium and practicality, is where Ferin Martino the composer is separated from Ferin Martino the performer. Musical nuances from Ferin's performance are stripped away, leaving room for interpretation by the next (human) performer. As some of the *aura* (after Walter Benjamin [10]) of Ferin's performance is striped away, the composition can gain new aura through its capacity for many variations in interpretations by other performers (cf. Eco [2]).

This algorithm's output is surprisingly pianistic to the ear, especially for its relative simplicity (one page of code) and emergent nature, and curiously, it does not sound characteristic, or even very good, using any other instrument's sound, even other keyboard instruments. One aspect contributing to its pianistic character must be its anthropomorphic design: with trichords played by a virtual left hand with five fingers, melodies played by a virtual right hand, pitch intervals and inter-onset (time) intervals kept within human scale, and "drunk walk" pseudorandom number generators to keep passages from being too disjunct. However, these design factors would make it just as good a fit for any keyboard instrument, including a harpsichord, organ, and electric piano, so there must be more to its good fit for the piano's sound.

It seems that the algorithm's output has been tuned to the modern pianoforte's timbre (the spectrum and amplitude envelope) inadvertently to the exclusion of success with other timbres. Ferin Martino's playing is often dense in pitch and in time. The harpsichord's bright spectrum makes such passages sound too harsh, and temporal density only allows the attack transients to come through while masking the more attractive and interesting decay of each note. In contrast, the electric piano's darker tone prevents individual notes from being discerned, often resulting in muddy sound masses. Organ tones result in similar results by their mostly flat amplitude envelopes, keeping note onsets from standing out, masked by the sustain of previous notes or with attention distracted by abrupt cutoffs—of course, organs are capable of quite lyrical legato playing, but Ferin's emergent proclivity toward compound melodies (à la Hindemith) is made distractingly disjunct with organ-like amplitude envelopes. Experiments with mallet keyboard instrument sounds were more pleasing in regard to these factors, but they lose Ferin's frequent expressive use of the damper pedal (and vibraphone suffers the same as the electric piano), besides the fact that it loses touch with human-scale playability on mallet instruments.

This highlights an important lesson in orchestration: with a long-established discrete pitch system, scales, a canon comprising music that heavily stresses melodies, chords, and human scale rhythms, and an abundance of theoretical analysis focusing on these parameters and not others, it is easy to imagine most instruments, especially the piano, as neutral, general purpose music generators, when really, each instrument's idiosyncrasies shape the musical decisions of the composer, even if unbeknownst to the composer. Since Ferin Martino has no capacity for evaluating its own output, especially with regard to the spectral and temporal features described above, so this is a crucial stage in the development of a performer algorithm. This relates in spirit to William Sethares's work in tuning systems, in which he demonstrates that consonance and dissonance are closely tied to timbre and tuning system more than Fortean interval class vectors. [11]

To finish the description of the transcription process: Ferin's dense playing usually ends up yielding sheet music requiring advanced piano skill to play. This makes it very amenable to treating Ferin's print out as a piano score to be orchestrated for small chamber ensembles, such as piano duo, piano trio (piano, violin, and cello), and string quartet. The manual orchestration process allows me (as the supervising composer, if not the one executing each note choice) to apply the considerations in the preceding paragraphs in order to map the musical content to fit each instrument as idiomatically effective as possible. The fact that Ferin's playing sounds more naturally pianistic than it looks in sheet music is also telling; it suggests that the character of the piano-human combination has its own idiosyncrasies, a subset of the full characteristic possibilities of the piano itself (including other playing scenarios).

One such piece was premiered by the Apollo Chamber Players, a string quartet based in Houston, Texas. The main material was created in a single take with Ferin. I was sculpting Ferin's behavior by moving my hands in front of the camera and occasionally introducing motives of a single repeated pitch via the piano keyboard. This created a dynamic in the piece of striving for stasis and falling away from it, as I would play repeated notes at a few certain

junctures, and Ferin would copy parts of that motive and carry it off in sequences. After the initial performance, I distributed the voices for a string quartet. To build a more coherent form (without trying to cover up Ferin's capricious nature), I added a brief introduction comprising three slow cascading iterations of the repeated-pitch motive, and I inserted a partial recapitulation before the final sequence: at a point where the music seems stuck in a loop, I inserted a break, then a sequence of truncated clips from previous moments, in the order they happened, but always interrupted. This section is played *sul ponticello* to set it off as a copy of previous events, as if retracing one's steps before continuing. This recapitulation leads to the repeated motive where the original music was stopped, and it continues from that point in full voice, going on to the end. I gave it the title, *The Garden of Forking Paths* to reflect on the many musical sequences Ferin tests before moving on, as well as its ultimately through-composed winding path.

9. EXPERIMENTS WITH DANCE

Given these experiences with tuning the software to respond to human movement, it was natural to consider adapting this work for dance. However, these experiments quickly highlighted differences between a camera with video analysis and the human eye with human cognition. As dancers explored the full stage space, they would range from being so close to the camera that most of their bodies could not be seen to being so far at end of the stage that their whole bodies only occupied a small portion of the camera's view. A camera operator would naturally want to zoom in and out and pan around to keep the dancer's full body in the video frame. However, this revealed that moving the camera even slightly can yield a maximal amount of change in the video signal even when there is no movement on stage. Seeing the need for more development for this kind of performance, this exploratory performance was put on hold.

The problem of dancers' distance might be addressed by an automatic spatial normalizing process that crops out any space that contains no movement, leaving a tight-framed rectangle inside to evaluate. However, this will require many more decisions, thereby shaping the character of the work. Additionally, the quick pace of dancers' movement might be outside the software's scope of responsiveness that was effective for causal visitors. Making it more responsive may result in a less interesting one-to-one kind of interaction. The notion of camera operator as performer is an intriguing one worthy of future investigation, but it too will require many character-defining decisions. This line of inquiry may need to use a completely different performing algorithm from Ferin, but informed by these experiences.

10. EXPERIMENTS WITH PAINTING

Working with a painter for visual input has avoided the challenges described with dancers, since a painter is largely engaged with a two-dimensional vertical plane, like the camera and screen. April Zanne Johnson [12] already uses music and sound as a stimulus to influence her mark making vocabulary in her drawing and painting practice. She also sometimes thinks kinetically about the act of painting: making painting gestures in the air before and after touching the brush to the canvas in a process of contemplation, winding up, and following through. Both of these facts have made collaborations with Ferin Martino very fruitful. This creates a fascinating multi-modal creative feedback loop as she paints in response to the music and her motion while painting in turn shapes the music. Ferin's capacity for endless through-composed playing also suits this scenario well. Together, they produce an experience of painting and composition as a realtime process laid out across time, besides the final artifacts that result from the process.

11. OTHER FUTURE WORK

As highlighted throughout the discussion above, there are many areas where future developments can be done. Further work will investigate possibilities for building larger formal structures over time, including buffers on medium and long time scale, as well as memory cues as in Butch Morris's Conduction technique for conducting ensembles of improvisers. The software could also generate its own MIDI reference track for beat mapping during transcription. Further software routines could be created to automatically interpret tempo and dynamics changes into standard words and symbols. Synthetic timbres can be explored, sharing the key features of the piano and avoiding problematic features described above. Also in this vein, the software can be tuned to suit other instruments. This would be an especially interesting inquiry for thinking about composition, since it turns out that many of its most interesting melodic lines incorporate both virtual hands to build compound melodies or sequences. Finally, the software could be taken to another level of sophistication and self governance if it could monitor the sonic features of its own output and adjust its note decisions accordingly.

However, as these opportunities are opened up above, the discussion also highlights distinctive features of the character oft his algorithm that would be significantly changed, perhaps to become less intriguing or idiosyncratic, or at least becoming something that should be considered a different process altogether. These advances will be considered carefully with this potential sacrifice in mind. It will more likely spawn a number different performer agents rather than numerous upgrades to this one.

12. CONCLUSIONS

It is delightfully gratifying that something originally meant to be an afternoon project to demonstrate a Disklavier digital player piano, with a modicum of interactivity, for a tour of university dignitaries would provide so rich an output, become so fascinating a collaborator, open up deep questions of the nature of composition and performance, and reach audiences in Milan, Athens, and in the United States, Texas, the Southeast, the Mid-Atlantic, and New England.

These experiences show that simple, perhaps even simplistic solutions can reach new levels of elegance and open

up new dimensions to contemplate, as with Alexander's solution to the Gordian knot and with Cantor's diagonal. In the end, an impatient attempt to simply "make music" has revealed myriad questions that are answered in that process, whether the composer takes conscious responsibility for them, leaves them to tacit intuition, or lets the governing technology fill in those answers according to what comes most naturally to it—for better or worse.

13. EXAMPLES

Video, audio, and musical scores resulting from this work can be found at [13].

Acknowledgments

Thanks to saxophonist Jayson Beaster Jones and painter April Zanne Jonhson for joining me in developing these productions, and to the I-Park Foundation and the Atlantic Center for the Arts for sponsoring this work.

14. REFERENCES

[1] D. Cope, *Computer Models of Musical Creativity*. Cambridge, MA, USA: MIT Press, 2005.

[2] U. Eco, *The Open Work*. Cambridge, MA, USA: Harvard University Press, 1962/1989.

[3] J. M. Morris, "Ferin Martino: A small piano algorithm and its lessons on creativity, interaction, and expression," in *3rd International Workshop on Musical Metacreation (MUME 2014)*. Artificial Intelligence and Interactive Digital Entertainment, 2014, pp. 40–44.

[4] T. di Milano, "History and mission," September 2016, http://www.triennale.org/en/chi-siamo/storia-e-mission/.

[5] A. S. Onassis Public Benefit Foundation, "Onassis cultural center athens," Decemeber 2015, http://www.onassis.org/en/cultural-center.php.

[6] J. Morris, "Ferin Martino demo," Video, September 2015, https://www.youtube.com/watch?v=iaDjezcrzHA.

[7] I. Park Foundation, "I-Park international artist-in-residence program," March 2003, http://www.i-park.org/.

[8] J. Morris, "Where Ferin was," November 2015, http://www.morrismusic.org/2015/where-ferin-was.

[9] J. M. Morris, "Collaborating with machines: Hybrid performances allow a different perspective on generative art," in *Proceedings of the Generative Art International Conference*. Polytechnic University of Milan, Italy, 2013, pp. 114–125.

[10] W. Benjamin, *Illuminations: Essays and Reflections*. Berlin, Germany: Schocken Books, 1936/1969, ch. The Work of Art in the Age of Mechanical Reproduction, pp. 217–252.

[11] W. A. Sethares, *Tuning, Timbre, Spectrum, Scale*, 2nd ed. New York, NY, USA: Springer, 2004.

[12] "Neurologically produced synesthetic forms and color: Studio visit with April Zanne Johnson," http://artmazemag.com/april-zanne-johnson/.

[13] J. Morris, "Ferin Martino," July 2014, http://www.morrismusic.org/ferinmartino.