The Automation of Surprise: Improvisation, Controlled Hallucinations, and the Predictive Mind

Jeff Kaiser, PhD.
University of Central Missouri, Music Technology and Composition, Warrensburg, MO United States.
jkaiser@ucmo.edu

Abstract

This article takes inspiration from the ideas of predictive mind theory that posits predictive processing: the idea our minds are constantly generating a simulation (or mental model) of the world as we experience it, testing our sensory input from the world against that simulation, and seeking to minimize prediction errors. This model, where the brain simulation is put first, has also been referred to as a controlled hallucination. Automation can be thought of as creating predictability, and in connection with ideas of predictive processing, minimizing prediction errors. I argue that we can also think of automation as enabling moments of surprise that disrupt the controlled hallucination and explore ways to automate surprises that might not be predictable. In live music, for many people, such prediction errors (failures) can be a positive experience, and many performers and audiences seek out surprise in performances. This is where I believe improvisation with technology accels. Specifically, I investigate how the automation of surprises from technological collaborators can disrupt this controlled hallucination in creative practice, in particular, improvisation. I define improvisation as a live interactive construction and ordering of sound where the players/actors (human, technological, and other) are not only constructing and ordering, but are being informed and presented with possibilities as to how to proceed by that which is being interacted with, constructed, and ordered. This creates a feedback loop of possibilities where all involved (including our technological collaborators) are both influenced and influencing, configured and configuring. Inserting surprise into this feedback cycle of influence through ceding decision making to our technological collaborators (what in the past was solely the domain of our human collaborators) through a variety of methods including probability-based algorithms and what I call "deep sonification", now includes our technological collaborators.

Keywords: Improvisation, Automation, Predictive Mind, Controlled Hallucination, Deep Sonification, Probability-Based Algorithms.

Overview

Prediction failure is one of the joys of my life. In particular, I am speaking about music, and like many performers and audience members I seek out surprise and novelty in the sonic arts. The cognitive philosopher Andy Clark calls this characteristic *"attractions of the unexpected"*, saying *"[w]e humans often seem to actively seek out surprising events, deliberately harvesting novel and exciting streams of sensory stimulation. Conversely, we often experience some well expected sensations as unpleasant and to-be-avoided"* (Clark, 2018). There is at least one medical study that suggests that certain people are subconsciously wired to enjoy surprises, i.e., prefer the unpredictable experience over the predictable (Berns et al. 2001)[1].

In predictive mind theory —one of the key concepts I draw upon for this paper— prediction precedes perception. (Also referred to as *predictive coding* or *predictive processing*). This theory posits that our minds are:

1. Constantly generating a "simulation" (or "mental model") of the world as we experience/perceive it, and;
2. Testing our sensory input from the world against that pre-existing simulation/mental model.

Variances between what is predicted and what is experienced are called "prediction error signals". These error signals are not 'bad' (something we frequently associate with the word 'error') but refer to differences, and help us to constantly update our mental model, and to minimize prediction error signals in the future (Hohwy, 2013; Kirchhoff and Kiverstein, 2019; Clark, 2023).

This model, where the mind simulation is put first, has also been referred to as a "controlled hallucination" (Clark, 2016, 2023), in that the prediction that starts inside (mind) is controlled by input from the outside world. As philosopher Anil Seth writes, *"perception is not outside in, but inside out"* (Seth, 2021b). Or as Clarke also writes, *"[i]ncoming sensory information is used to keep the model honest"*. Without that control from the outside world, it becomes an "uncontrolled perception", or hallucination, i.e., an invention of the mind not grounded in sensory input from the world. As Seth explains in a Ted Talk, *"if a hallucination is kind of an uncontrolled perception, then a perception is a controlled hallucination. [...] We have hallucinations all the time, except when we agree on them, then we call it reality"* (Seth, 2021a).

In the history of western art music, we see examples of composers working to surprise people such as Haydn's *Surprise Symphony* and working to mimic/share/communicate the experience of the unpredictable world of drugs (i.e. hallucinations) such as Berlioz's *Symphonie Fantastique.* The allure and desire to surprise people and mimic/share/communicate unpredictable drug experience has of course been present in popular music spawning its own multiple genres of psychedelic music, Psychedelic-[folk, funk, pop, rock, soul], Acid-[house, jazz, rock, techno, trance], Space rock, Stoner rock, Trip

1 The study, *"Predictability Modulates Human Brain Response to Reward",* states, "Certain classes of stimuli, such as food and drugs, are highly effective in activating reward regions. We show in humans that activity in these regions can be modulated by the predictability of the sequenced delivery of two mildly pleasurable stimuli..." and later, "...most of the subjects did not discern any difference between the predictable and unpredictable conditions. If the unpredictable rewards were more pleasurable than predictable ones, or vice versa, then this must be occurring at a subconscious level" (Berns et al. 2001).

hop, and more. But once fixed in either notation and or recording, how does surprise work? If the experience of surprise is the experience of "prediction error signal", which is reliant upon our failure to predict correctly, on the second listen, will our prediction be correct, negating the experience of surprise[2]? Anticipation —which is related to the sensation of predictability and surprise, has been discussed in such works as David Huron's *"Sweet Anticipation"* (Huron, 2006) and can be brought to bear on our understanding of the joy that comes with knowing what is coming, such as a drop in dance music et cetera.

Defining Automation with Room for Surprise

Automation, in a musical recording environment, traditionally has been conceptualized as creating predictability by assigning a particular action to a means of control such as an automation envelope in a DAW, LFO in an analog synth or DAW —even humans— as legendary engineer and author Bobby Owsinski writes,

> *Automation is the process of recording parameter adjustments so that your DAW software can automatically execute them during playback. Before automation, any mixing moves during a complex song had to be made by a combination of the engineer, the assistant, the producer, the band members, and anyone else who had their hands near the console. While this might have been a very "organic" way of mixing, it was not repeatable in the slightest, so engineers everywhere longed for a way that at least their fader moves could be memorized and played back later* (Owsinski, 2020).

Three elements are present in this conception of automation:

1. The desire for technology to do the work.
2. The technology replicates, and minimizes, human activity, and;
3. The goal of being repeatable.

When these traditional elements are situated within predictive mind theory, automation can be seen as minimizing prediction error signals, i.e., creating predictable behavior, from the expected rise and fall of volume, the placement of panning, and enabling/disabling/modifying effects. In other words, automation can create a desired goal of predictability. Yet, handing these tasks off to human assistants runs the risk of them inserting personality by mechanically adjusting parameters in a certain fashion. Automation envelopes added later in software that control things in the DAW predictably, came to mean drawing out exactly what the automation will do, where it will do it, when it will do it, how it will do it, et cetera, erasing unpredictability, surprise, and variation by erasing different personalities being involved.

With the advent of probability-based control in such platforms as Ableton Live in 2021 (Kirn, 2021), AI-based software automating mixing and mastering such as Izotope since 2016 (Stewart, 2023), and the recent RipX DAW, which is billed as the world's "first AI DAW" (Mullen, 2023) adds a fourth element to our current conception of automation,

4. The ceding[3] of decision-making at the point of automation to the technology, further minimizing human interaction while allowing creative variation[4].

This ceding of decision-making —something that diminishes the centrality of humans and human agency— acknowledges what has always been the case: that agency is shared by all participants in the system, including technological participants. A point for further discussion might be: when does automation cease being automation and become an autonomous agent? Or, is it acceptable to view an autonomous agent as a form of automation? For me, the technological autonomous agent can clearly fall under the category of automation as it is about minimizing human interaction, which includes decision-making. I believe this can be referred to, borrowing from Hans-Georg Moeller talking about Niklas Luhmann, as bringing a form of "radical non-anthropocentrism" (Buitendag, 2021) to automation. This dissipates much of the differentiation between automation and generation as we acknowledge and allow for a broader understanding of automation in the audio world. Automation is no longer only about performing the specific will of the engineer in an exact repeatable manner (although that is one option) but about allowing our technological collaborators to make decisions in the process.

Disrupting the Controlled Hallucination

When we incorporate decision-making, i.e. the generation of information, into our definition of automation, we can also think of automation as encouraging and creating prediction failures (surprise) in human observers and participants. This can range from a sudden loud moment in cinematic music or the use of a drop in dance music, which can be surprising the first time, but the second time perhaps rely more on anticipation than surprise. Automating surprises so as not to be predictable allows the computer to make decisions in a range, including making a decision as to when to make a decision. If we view automation as running a set of rules (or algorithms) at a given moment,

2 This has been a concern of musicians for a while, leading to a dismissal of novelty or unexpectedness. As composer and music theorist Edward T. Cone wrote, "Because of our own attempts to create a music of continuous mutation, we suspiciously scrutinize even the classics for traces of tautology. We tend to forget that, of all the artistic effects, novelty is bound to be the least permanent". He also later writes, "[N]o matter how many times we hear the *Eroica*, the moment before the recapitulation never loses its effect. Indeed, the better we know the piece, the more inevitable, and there the more satisfying, the resolution seems to us" (Cone, 1968). This statement can be contrasted with another generalizable aphorism, "familiarity breeds contempt". Cone seems to be wired to enjoy predictability. I believe this is where live music has the potential to excel, by incorporating surprise, even when automated.

3 I struggled with whether to use cede or give in this paper. I went with cede because decision-making possibility is already in the computer, we are not giving it something it already has.

4 Even in the past, when delegating faders to human assistants, decision-making was delegated as we could not control the details of the fader move by the human assistant, i.e., that they moved the faders in a perfect linear or exponential manner et cetera.

why not include surprise as a possibility of one of those rules? We also need to ask what constitutes surprise as we examine the granularity of automation, from small units at moments to gestures over time, to looping motions of different kinds. Can there be a surprise at every moment? Or does surprise only occur with an unpredictable moment after a period of predictability? With the emergence of the use of probability and AI in production environments, we have seen the emergence of automating probability-based events. If every sixteenth note a hi hat is played at a different velocity is that surprising? Or does the irregularity become predictably irregular. This starts to become interesting to me when we have multiple agents introducing multiple musical ideas —traditionally the role of human musicians— to increase the potential for surprising moments and gestures[5]. We can now include other agents from timelines with simple probability-based triggers to other algorithms that incorporate machine learning and more, allowing complex, unpredictable, musical events to emerge from such automation.

The current thought for some on AI and generation is that AI is not good at novel expressions with the ability to surprise, or as scholar Jason Palamara posits, *"that while AI can create music at a high level, it's not yet able to emulate the aspect of choice and surprise that characterizes so much of human creativity"* (Gordon, 2023).

This is rapidly changing. There are examples of AI surprising researchers (Ornes, 2023) and, AI can "hallucinate", which increases potential for creating surprise. AI is trained in a way to validate predictions by mimicking other music/art that has already been made. Hallucination, in the case of AI, is sometimes defined as inventing information not grounded in their external input, much like predictive mind theory posits. According to one group of researchers in a recent study, chatbots may "invent information" from 3 to 27 percent of the time (Metz, 2023).

> *Experts call this chatbot behavior 'hallucination'. It may not be a problem for people tinkering with chatbots on their personal computers, but it is a serious issue for anyone using this technology with court documents, medical information or sensitive business data* (Metz, 2023).

This leads to episodes such as the chatbot that cursed at a customer and criticized the company that employed the chatbot, leading to a comment we will likely hear more and more often in the future, *"[a]n error occurred after a system update on Thursday, Jan. 18. The AI element was immediately disabled and is currently being updated"* (Moench, 2024).

To have the computer as a creative collaborator in the arts, we might want to encourage systems that are more inventive by allowing and encouraging their hallucinations. In other words, we might automate locations where we welcome the invention of information not grounded in external input.

Ceding decision-making to the computer via probabilistic algorithms is an artistic practice I use, a practice that acknowledges the agency of the computer to make behavioral, predictive, decisions not "controlled" by testing against external input. Just allowing it to "invent information", i.e., hallucinate, and act on its own beliefs[6].

Instead of predictability, let's encourage our computers to hallucinate and disrupt the controlled hallucination of predictability in automation[7].

Deep Sonification and Probability-Based Electronica: Made Audible

Personally, I have been building Max patchers for mostly digital audio signal processing since 2005. In 2011 on a guest lectureship at Amherst College, I was asked by a student about how to build a M4L plugin that would automate the generation of bass lines with variations. I had done some work in Ableton Live, but it was there in front of a student that I built what would form the basis of my KaiGen suite of plugins[8]. Up until that time, surprise was provided either by the people I was performing with (improvisers surprise each other all the time) or by exploiting "combinatorial explosions" with technology. Writing about language, cognitive neuroscientist Stanislas DeHaene uses the term "adjustable parameters", which readily translates to music technology:

> *[W]e are unable to fathom the extraordinary number of possibilities that open up as soon as we increase the number of adjustable parameters. This is called the "combinatorial explosion" —the exponential increase that occurs when you combine even a small number of possibilities. Suppose that the grammar of the world's languages can be described by about fifty binary parameters, as some linguists postulate. This yields 2^{50} combinations, which are over one million billion possible languages, or 1 followed by fifteen zeros! The syntactic rules of the world's three thousand languages easily fit into this gigantic space. However, in our brain, there aren't just fifty adjustable parameters, but an astoundingly larger number: eighty-six billion neurons, each with about ten thousand synaptic contacts whose strength can vary. The space of mental representations that opens up is practically infinite* (Dehaene, 2020).

In my case, combinatorial explosions occur after my live processing rig, created using the Max visual programming environment, had enough possible

5 Performers are able to add surprise by encouraging changes in the musical direction by the introduction of musical ideas that contrast rather than compliment. At a rehearsal for a concert I was in, George Lewis encouraged the players, "whenever you hear the call, don't respond" (Lewis, 2010) as complementary ideas can create predictive results, ideas that contrast have the potential to create unpredictable paths.

6 Beliefs can be referred to as "probabilistic states" (Bottemanne, Longuet, and Gauld, 2022).

7 To be clear, I am talking about the arts, not when flying airplanes, performing surgeries, et cetera.

8 While Ableton introduced probability features in 2021 with Live 11, Max users have been making patchers using probabilities since the beginnings of that visual programming language that has captured the imagination of so many musicians looking for an alternative to text-based coding environments. Followed by the release of Max for Live (M4L) in 2009 (Ableton Public Relations, 2009), many of us started porting patchers to M4L plugins. My examples will refer to my own Max patchers and M4L plugins. The M4L suite is called KaiGen and is freely available at http://jeffkaiser.com/kaigen/.

variations that hitting multiples would create surprising results. Substantially different from the variation of five or six effects pedals on my old hardware-based performance rig. Treating devices as binaries, being on or off, five or six devices would yield $2^5=32$ or $2^6=64$ combinations, now with 16 processing modules in my software rig there are $2^{16}=65536$ possible binary combinations. That number yields quite a bit of sonic variety. Faking randomization by just blindly punching at the controller selection pads on my rig creates surprising combinations for varied reasons: sometimes surprising by the fact that nothing comes out, other times by unique sounds that I would have a hard time imagining, creating, let alone recreating. Automating the randomization of the 16 binary states (on/off) became an easy way to create surprise. And this is only treating those 16 processing modules as having binary parameters (on/off) when you add the variable control parameters within each individual module, then it becomes a number difficult to imagine.

After I started making the KaiGen M4L plugins, I got more and more into groove-based music (music that Ableton Live was built to privilege) and formed a duo with Trevor Henthorn of *Sweat Engine and Aesthetic Meat Front* fame. I worked for Trevor in the Summer of 2008, assisting him in his Programming position at the University of California, San Diego, where (among other things) he kept the computers in the music studios alive, while performing and creating music as often as possible. In addition, he was a professor of music technology at the Art Institute of California, San Diego. We called ourselves *Made Audible* (http://madeaudible.com/) after the concept that "music is numbers made audible" a quote variously attributed to Boethius, Aristotle, Pythagoras, and others. For us, music was probability made audible, and our website from the beginning proudly declared us as performing "Probability-Based Electronica". We focused on building sonification and probability-based Max for Live plugins. For us, sonification and probability were ways to cede decision-making to the computer, making the computer a full participant in the music making. This was very much inspired by the work of professor, author, musician, and technologist George Lewis, and we wanted to create our own version of a *"nonhierarchical, interactive musical environment that privileges improvisation"* as he described his Voyager in *Too Many Notes: Computers, Complexity and Culture in Voyager* (Lewis, 2000).

Recently, in conversation, Henthorn described the music of *Made Audible* as,

> *...the juxtaposition of the predictable with the unpredictable. Predictable data that would be heard for what it was, like physics and nature, that could be associated with something real via sonification. For us, this included data sets such as public restroom locations in Amsterdam and their spatial relationship to restaurants, and bars. This would then be juxtaposed with the unpredictable, generated surprise through probability gates performing different functions. When performing live, there were moments of solo, with the addition of the 'I have no idea of what is coming next' from Jeff and his probability-based plugins and with me where you 'might know what is coming next'. Sonic surprise is always the challenge* (Henthorn, 2024).

Sonification is one way of using data to automate performance parameters. In our experience, much of the sonification we heard from others had at its foundation a one-to-one relationship between sound and data (information). Our goal was to move beyond that one-to-one relationship to what I like to *call deep* sonification, which draws upon the "conversational" model of "deep mapping", the core of which author, professor, and digital humanist Maureen Engel writes "is a way to open up questions rather than resolving them, to communicate *knowledge* rather than simply information" (Engel, 2018). In our own way —sometimes humorous, sometimes serious— we were opening up conversations about the places, people, and politics we brought into our musical conversation via data sets, audio samples, probabilities, and deep sonification.

Working in Ableton Live, Trevor used classic drum machines and synthesizers to create the "recognizable" while I worked mostly in the realm of the "juxtaposed" employing samples of politicians and more, creating what we began to understand was the automation of unpredictability. This predated Ableton's 2021 incorporation of probability-based functions in Ableton Live. (Although since 2009 you could do it with M4L plugins). Methods used included Trevor playing consistent, recognizable dance beats, while I worked with my laptop to create floating probability-driven rhythm parts that were given the opportunity to play unpredictably, while the laptop would intervene/decide "do I play it?" and "if I play it, how loud do I play it?" We added pitch-based instruments that would start to make decisions on pitch, harmony, durations, velocities, scales, rhythmic variety. Here is an example of one of the KaiGen plugins (KaiGen-R) where the creation of dance music is automated while interacting with the human player and laptop, creating surprising results:

https://youtu.be/bGFgFUbePDo

I used the same probability ideas to create melodies and harmonies with KaiGen-M:

https://youtu.be/IJcJ1n2dg48

We combined these with a KaiGen plugin that used the I-Ching to create synthesizer patch changes:
https://youtu.be/PuaQT1-oyB4

And finally, added probabilistic sample shuffling with politicians and other samples using the KaiGen-Smaple-Palyer:

https://youtu.be/MhK7H72z6Io

There is also an included plugin with the suite, KaiGen-C, that generates MIDI controller information that can be smoothed and manipulated in different ways.

Great for automated knob twiddling that goes beyond LFOs.

Automaton as Automation

Drawing upon my dissertation research on improvisation and technology, alongside the creative practice and scholarly collaboration[9] with the other half of the duo *KaiBorg* (http://kaiborg.com/) —my good friend and professor/author David Borgo— I wrote that improvisation can be,

> *defined as a live interactive construction and ordering of sound where the players/actors (human, technological, and other) are not only constructing and ordering, but are being informed and presented with possibilities as to how to proceed by that which is being interacted with, constructed, and ordered. This creates a feedback loop of possibilities where actors are both influenced and influencing, configured and configuring* (Kaiser, 2018).

The earlier versions of this definition led to my development of a probability-based improvisation algorithm that would cede decision-making to the computer based on pitch estimation, a method I was introduced to by my dear friend Ritwik Banerji while we were both graduate students. Banerji, now a composer/performer and professor focusing on the anthropology of sound and experimental ethnography, showed me how he would, as he later wrote, *"generate creative output in interaction with human musicians by exploiting a pitch detection algorithm's idiosyncratic interpretation of a relatively noisy and pitchless sonic environment"* (Banerji, 2016). The creative output of Banerji's improviser, named Maxine, was constantly surprising and incredibly enjoyable to me. I remember the first time I heard Maxine perform with Banerji at his home in Berkley, just the three of us; it was a pure joy to hear. It inspired me to create my own improvisor/technological collaborator that would use a similar pitch estimation algorithm with its own system of probability gates that work a bit differently[10].

George Lewis, whom Banerji cites in the above mentioned paper, also uses pitch detection, where with Lewis's players, as he writes in his paper *Interacting with Latter-Day Musical Automata, "These behavior specifications are in turn determined in part by the analysis and development of pitch and velocity data, which is taken from the improvisor's playing via a so-called 'pitch follower' —a device known to exercise its own creative options from time to time".* This is followed by feature extraction ("volume, velocity, sounding duration, interonset duration, octave, register, interval width, pitches used, volume range, interonset duration range, frequency of silence, articulation and other important features"), and then random numbers that "provide much of the 'personality' of the system, and include melody and harmony, orchestration, ornamentation, pacing, transposition, rhythm, and internal behavior decisions". Notably, when there was no input from the human player, the musical "behavior specifications" are created by the computer (Lewis, 1999).

This employment of an automaton —a "player" in Lewis's terminology, a "virtual free improviser" for Banerji, or for me, a "technological collaborator"— can be seen as automation through the creation of an automaton. Much like the early days in the studios of handing off tasks to human assistants/collaborators (that might introduce personality), we can now hand off tasks to automata.

Similarly to both Lewis and Banerji, "exploiting" the pitch detection software and, to a lesser extent, feature extraction, are used in my software. The algorithm also builds off my above definition of improvisation using the feedback loop as automation to insert surprise into the feedback cycle. What can be the role of humans in a human ensemble, now goes to my technological collaborator.

This grew out of a desire, as it did with Banerji, for a technology-based collaborator that fit more of a "free improvisation" model. For me, it became a player that would be in the moment, a simplified software algorithm that didn't build, store, and analyze databases, but where memory existed in the feedback loop: where remnants of what was played were being looped and changed with what was currently being played. The loop, drawing on Professor Rodrigo Quiroga, becomes a kind of "forgetting machine" that forgets more than it remembers. As Quiroga writes about storage of memories,

> *So how do we do it? How do we store all this information? The surprising answer is that we basically do not. We remember almost nothing. The idea that we remember a great deal of the subtleties and details of our experiences, as if we are playing back a movie, is nothing more than an illusion, a construct of the brain[11]* (Quiroga, 2017).

The following example is of an improvising algorithm where the decision-making is automated, handed off to the computer[12]. This, as noted in my first listening to Maxine, is one way to automate surprise in music, allowing our technological collaborators to improvise. If automation is perceived as the execution of another's will, it becomes hierarchical, much like pointing a finger and commanding, but ceding decision-making allows room for nonhierarchical roles of automation. In ceding decision-making to an algorithm, as in the following examples, and treating this as a form of automation by minimizing human interaction, we also further the idea of "radical non-anthropocentrism".

9 Configurin(g) KaiBorg: Interactivity, ideology, and agency in electro-acoustic improvised music (Borgo and Kaiser, 2010).

10 Banerji and I have talked about getting our systems to play together, something we have not yet done but hopefully will.

11 To be clear, my work aims to open up creative spaces inspired by the ideas of cognitive science and neuroscience, in particular ideas coming from those working in the area of 4E cognition which grew out of the "realization that cognition was not limited to processes in the head, but was embodied, embedded, extended, and enactive" (Newen, Bruin, and Gallagher, 2018). I am not trying to create software that imitates how the mind or brain works: I firmly believe the brain is not a computer and the computer is not a brain. For a thought-provoking essay on this topic, I suggest Robert Epstein's, The Empty Brain: Your brain does not process information and it is not a computer (Epstein, 2016).

12 Like Lewis, this algorithm was created while I was in residency at STEIM in Amsterdam.

In automation, we are dealing with different granularities of decision-making. In the following algorithm, there is a high level of granularity of occurrence, with a decision being made almost constantly. Whether this is incorporated with lower —or higher— difference thresholds limits or restricts the range of decisions that the computer can make.

When I first presented this to a computer music researcher at a research university his first response was, “It should not work, it should just start playing octaves and play nothing but octaves”. But because the acoustic world is not as pristine as, let’s say, isolated MIDI note numbers in a database, and because microphones allow in all of the glorious sonically imprecise information from the environment —including possible bleed from the technological collaborators themselves— the computer starts to make decisions that do not exactly mimic the human. The pitch estimation portion of the algorithm likes to “exercise its own creative options” and “idiosyncratic interpretation” of what is going on. By letting it decide to listen (or not) to the various inputs of the human player(s), itself, and other technological collaborators, this automation of small moments becomes a combinatorial explosion of sonic possibilities.

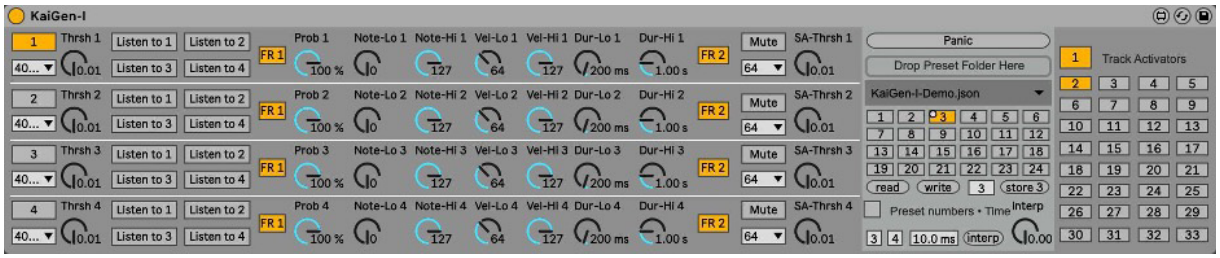

Figure 1. KaiGen-I, four players version.

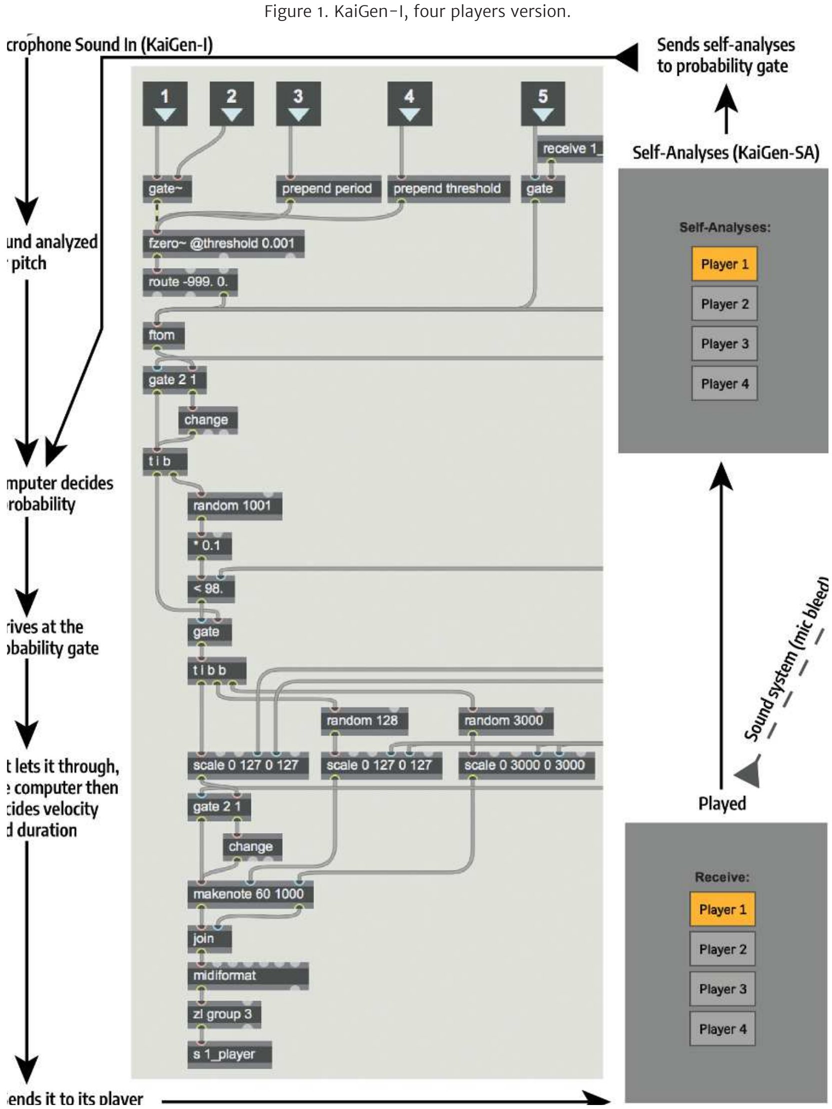

Figure 2. KaiGen-I internal processes.

Here is an example that happened when I was working on the software at my house while the refrigerator was running, cars passing by, the fan overhead was whirring, me walking to get a cup of tea, and KaiGen-I was listening. High levels of restraints are set as the algorithm responds to a literal version of what Banerji described earlier as a "noisy and pitchless sonic environment". You can hear the feedback loop at the beginning, and then it begins to vary:
https://youtu.be/d6txBadyjiM

Here is another example of a less constrained instance of each player listening to itself and each other. The starting instance was a single "whoop" into a microphone:
https://youtu.be/whvDUrTBsUI

Finally, here is a video of me improvising live on flugelhorn with KaiGen-I. Changes in constraints are controlled by 'dummy clips' in Ableton Live that contain automation envelopes that change the settings of the KaiGen-I plugin. These clips are played using 'follow actions' on the scenes that automate the advance to the next scene, creating the sectionality of the work. We can hear KaiGen-I follow me at first, but slowly starts making its own creative decisions, where about one minute in it is playing its own thing and I am following it. Then it comes back to a more close following of me at the end. For the purposes of this demo, scenes are short and follow each other in linear fashion. Probability can be added both to deciding whether to advance to the next scene —or not— as well as making the next selection different (next, previous, specified jump, random, et cetera). In this example, KaiGen-I is playing the IRCAM prepared piano sample library, which has also been randomized:
https://youtu.be/dauGnfLu5do

By combining probability in the automation of parameters that encourage or constrain behaviors of the technological collaborator —with what might now be seen as more traditional ideas of control— we can start to get both predictable and unpredictable responses: allowing for the automation of surprise.

Bibliography

Ableton Public Relations. 2009. "Ableton Releases Max for Live". November 23, 2009. https://www.ableton.com/en/press/press-archive/press-archive-mfl-release/.

Banerji, Ritwik. 2016. "Balancing Defiance and Cooperation: The Design and Human Critique of a Virtual Free Improviser". In International Computer Music Conference.

Berns, Gregory S., Samuel M. McClure, Giuseppe Pagnoni, and P. Read Montague. 2001. "Predictability Modulates Human Brain Response to Reward". The Journal of Neuroscience 21 (8): 2793–98. https://doi.org/10.1523/JNEUROSCI.21-08-02793.2001.

Borgo, David, and Jeff Kaiser. 2010. "Configurin(g) KaiBorg: Interactivity, Ideology, and Agency in Electro-Acoustic Improvised Music". In Beyond the Centres: Musical Avant-Gardes Since 1950, edited by Costas Tsougras, Danae Stefanou, and Kostas Chardas. Thessaloniki, Greece: Aristotle University.

Bottemanne, H., Y. Longuet, and C. Gauld. 2022. "The predictive mind: An introduction to Bayesian Brain Theory". L'Encephale 48 (4): 436–44. https://doi.org/10.1016/j.encep.2021.09.011.

Buitendag, Nico, dir. 2021. Luhmann & Cybernetic Identities w/ Hans-Georg Moeller. Undisciplined. https://www.youtube.com/watch?v=j5SaBSFApwc.

Clark, Andy. 2016. Surfing Uncertainty: Prediction, Action, and the Embodied Mind. New York: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780190217013.001.0001.

———. 2018. "A Nice Surprise? Predictive Processing and the Active Pursuit of Novelty". Phenomenology and the Cognitive Sciences 17 (3): 521–34. https://doi.org/10.1007/s11097-017-9525-z.

———. 2023. The Experience Machine: How Our Minds Predict and Shape Reality. Knopf Doubleday Publishing Group.

Cone, Edward T. 1968. Musical Form and Musical Performance. Illustrated edition. New York, NY: W. W. Norton & Company.

Dehaene, Stanislas. 2020. How We Learn: Why Brains Learn Better Than Any Machine ... for Now. Penguin.

Engel, Maureen. 2018. "Deep Mapping: Space, Place, and Narrative as Urban Interface". In The Routledge Companion to Media Studies and Digital Humanities, edited by Sayers, Jentery. Routledge.

Epstein, Robert. 2016. "The Empty Brain: Your Brain Does Not Process Information and It Is Not a Computer | Aeon Essays". Aeon. May 18, 2016. https://aeon.co/essays/your-brain-does-not-process-information-and-it-is-not-a-computer.

Gordon, Eden Arielle. 2023. "A New Beatles Song Was Just Released Thanks to AI. This Is Just the Beginning". News. POPSUGAR. November 9, 2023. https://www.popsugar.com/amphtml/entertainment/music-and-ai-beatles-drake-essay-49314403.

Henthorn, Trevor. 2024. Interview with Trevor Henthorn.

Hohwy, Jakob. 2013. The Predictive Mind. OUP Oxford.

Huron, David. 2006. Sweet Anticipation: Music and the Psychology of Expectation. Cambridge: MIT Press.

Kaiser, Jeff. 2018. "Improvising Technology, Constructing Virtuosity". Cuadernos de Música, Artes Visuales y Artes Escénicas 13 (2): 87–96. https://doi.org/10.11144/javeriana.mavae13-2.itcv.

Kirchhoff, Michael D., and Julian Kiverstein. 2019. Extended Consciousness and Predictive Processing: A Third Wave View. Routledge.

Kirn, Peter. 2021. “Ableton Live 11 Has a Release Date on 23rd - What’s New, What Do You Want to Know?” CDM Create Digital Music. February 23, 2021. https://cdm.link/2021/02/ableton-live-11-has-a-release-date-on-23rd-whats-new-what-do-you-want-to-know/.

Lewis, George. 1999. “Interacting with Latter-Day Musical Automata”. Contemporary Music Review 18 (3): 99–112. https://doi.org/10.1080/07494469900640381.

———. 2000. “Too Many Notes: Computers, Complexity and Culture in Voyager”. Leonardo Music Journal 10: 33–39.

———. 2010. “George Lewis: Compositions, Improvisations, Conversations”. Concert, University of California, San Diego, Conrad Prebys Music Center, February 10.

Metz, Cade. 2023. “Chatbots May ‘Hallucinate’ More Often Than Many Realize”. The New York Times, November 6, 2023, sec. Technology. https://www.nytimes.com/2023/11/06/technology/chatbots-hallucination-rates.html.

Moench, Mallory. 2024. “Delivery Firm’s AI Chatbot Curses at Customer”. News. TIME. 20 2024. https://time.com/6564726/ai-chatbot-dpd-curses-criticizes-company/.

Mullen, Matt. 2023. “Hit’n’Mix Says the New Version of Its RipX Software Is the World’s ‘First AI DAW”. News. MusicRadar. November 8, 2023. https://www.musicradar.com/news/ripx-ai-daw.

Newen, Albert, Leon De Bruin, and Shaun Gallagher. 2018. The Oxford Handbook of 4E Cognition. Oxford University Press.

Ornes, Stephen. 2023. “The Unpredictable Abilities Emerging From Large AI Models”. Quanta Magazine. March 16, 2023. https://www.quantamagazine.org/the-unpredictable-abilities-emerging-from-large-ai-models-20230316/.

Owsinski, Bobby. 2020. “A Look At The History Of Console Automation”. Bobby Owsinski’s Music Production Blog. June 17, 2020. https://bobbyowsinskiblog.com/console-automation-history/.

Quiroga, Rodrigo Quian. 2017. The Forgetting Machine: Memory, Perception, and the Jennifer Aniston Neuron. BenBella Books.

Seth, Anil. 2021a. “Anil Seth: How Your Brain Invents Your ‘Self’ | TED Talk”. October 2021. https://www.ted.com/talks/anil_seth_how_your_brain_invents_your_self.

———. 2021b. Being You: A New Science of Consciousness. Dutton.

Stewart, Ian. 2023. “AI & Automated Mastering: What to Know”. Software. Izotope. April 19, 2023. https://www.izotope.com/en/learn/ai-mastering.html.