

Automated composition using autonomous instruments*

Risto Holopainen

February 2014

A crucial aspect of creativity is to let go of the conscious control one has over the creative process. Composers have since long put into practice the use of oracles or methods that allow for more or less systematically delegating their choices, which in its purest form implies algorithmic composition. Because the roots of algorithmic composition can be found in formalized techniques of composition, examples of such a delegation of choices can be found throughout the Western history of music.

Autonomous instruments have much in common with algorithmic composition in its method as well as its *raison d'être*. It is a way of composing by means of sound synthesis, where high level processes arise as emergent phenomena by applying rules or mechanisms that act on a lower level. Concepts such as self-organization, emergence, complexity and autonomous systems are frequently associated without being synonymous. Feedback is crucial in most contexts where emergent phenomena and self-organization are mentioned. The frequent use of feedback in autonomous instruments, be it acoustic or more indirect forms of feedback, is not a mere coincidence. Composition is no longer a matter of organizing sound, but to create processes that generate self-organized sound.

Limits of autonomy

There is much talk about the affordances of new electronic or digital instruments, and there is possibly even more talk about the absence of affordance in badly designed digital instruments. Already at first sight, we can envisage several things one might do with a violin, say, to pluck its strings, to tune it differently, accidentally sit on it, deposit it at a pawnbroker . . . and correspondingly with other instruments. A synthesizer may have keys for playing, buttons to push and knobs to twist. It is obvious what one can do with the synthesizer. But who would get the idea to fasten one of its keys with tape and let that be the entire composition? Well, a minimalist might do that.

*This version was originally published in the Danish journal *Seismograf* in 2014. An earlier version in Swedish has been published in *Lydskrift*.

Autonomous instruments are systems that one does not interact with as they generate sound. Music making with digital instruments that are not played in the ordinary sense has its roots in the MUSIC N family of synthesis languages. Developed by Max Matthews in the 1960's, these languages are often split into an instrument definition and a score file where the instruments' parameter values at specific points in time are specified. The division into instruments and score files is flexible but presupposes that the music is regarded as composed of individual notes. However, the duration of individual notes may be a few milliseconds and there may be millions of notes collectively giving rise to a granular swarm of sound, or a single note may constitute the entire composition regardless of its length.

Now there are several minimalist works made of a single *tone* (e.g. some of La Monte Young's), but the concept of a note in MUSIC N covers much more. A note does not have to correspond to a tone. The synthesis algorithm itself may engender changes over time, so that a single note can expand into long and complex musical processes.

For such long lapses of time to be musically complex there would have to be variation on several temporal scales. Various strategies can be used to create such variation on multiple temporal levels, such as:

- Stochastic control signals asynchronously updated
- Layering of periodicities each with its incommensurable length
- Self-regulating feedback systems (chaotic systems)

Moreover, these methods can be combined into various hybrids. The method of stochastic signals is exemplified by Xenakis' piece *S.709* realized with the GENDYN program. Periodically recurring material with periods that phase out have been employed, e.g., by Brian Eno in *Music for Airports*. I will return to examples of feedback systems. This strategy of using non-interactive instruments that generate large scale sonic processes without explicitly referring to a note level is typical of autonomous instruments.

*

Autonomous instruments must be automatized in some way, which is easiest to do in the digital or analog domain. What would remain of autonomous instruments had the electricity been turned off? Mechanical autonomous instruments can be imagined; they seemingly would have to be a perpetuum mobile that charges itself with the energy it consumes. I recall reading the novel *Bernhard's magical summer* by the Swedish author Sam J. Lundvall long ago. The main character goes to a night club where he meets the devil. As it turns out, the devil has a collection of seized perpetuum mobiles, which explains why we never see any that work! In absence of perpetual motion one will have to accept that the autonomous instrument receives its energy input from outside. It could be done as in Ligeti's *Poème Symphonique* where the metronomes are wound up and loaded with mechanical energy and the piece lasts as long as the energy is

converted into pendular motion and ticking. As with autonomous instruments, there is no interaction with the metronomes during the piece. “Look ma, no hands!” Very well, but it takes more.

Another of the characteristics of autonomous instruments is their ability to create sonic processes that were not planned in advance, or that could not have been predicted. Otherwise the CD player would count as an autonomous instrument.

Of course it is pointless to speak of autonomy in some absolute sense. There is only autonomy with respect to something else. Mechanical as well as electric instruments must have their energy supply. Autonomous instruments require a separation towards the composer. They are not supposed to stand under one’s full control. However, the same thing may be said of the combination of an unexperienced musician and a difficult to master instrument, and that is a different matter altogether. Or is it?

Thus, if one is to abstain from playing the instrument, one apparently becomes restricted to some kind of *objets trouvées*. The strategy would not be much different from haphazardly aiming a camera and taking random snapshots or doing field recordings with a recorder left somewhere to register sounds, turning itself on by sounds and turning itself off when nothing happens. Having decided which algorithm to use for the autonomous instrument, the output may be regarded as a found object. The result is a sound file generated by the autonomous instrument. But whence came the algorithm? In order for the instrument to be fully autonomous it ought to have developed its own algorithm. Neither music nor computers grow wild in nature, as Herbert Brün once put it. This is a debatable assertion, but I will leave it for another occasion.

The use of existing algorithms or dynamic systems for sound synthesis or generative music is more related to sonification than to autonomous instruments. In sonification, algorithms actually may be used as found objects. However, the strategy of autonomous instruments is to invent an algorithm for the specific purpose of making it generate an acceptable sound file.

Aesthetic preferences

Now there is the dilemma of quality control. Must one not abdicate from one’s personal taste if assigning the creative act to the machine? However, nothing much is transferred to the machine, save for huge amounts of numerical operations. In particular, one keeps one’s status as the creator of the sound file just as the photographer is the capturer of his photo or the journalist the writer of her interview. Nonetheless, composing well with the help of autonomous instruments is a tricky affair. One must have preferences. Then the autonomous instrument can be programmed in such a way as to make it sound well according to those preferences. The instrument certainly will be resilient exactly to the degree it remains autonomous; there are the most peculiar blocks to improving specific aspects of the composition. On the contrary, having full control over the results by definition means that whatever one is doing, it no longer has anything

to do with autonomous instruments.

It is not always clear in what ways the composer's preferences guide the creative process; the notion of preferences itself is diffuse. Ask a visual artist about her favourite colour (that is a truly elementary preference), ask a composer about his favourite pitch or ask a mathematician about lucky numbers; for most people such questions are meaningless. Working with autonomous instruments leads to another level of choices. The autonomous instrument provides a draft, be it a sketch for a complete work or just some fragmentary material. There are many possible approaches to this material; it can be regarded as the final product or as raw material that will be submitted to a lengthy process of editing. Let us suppose we will use the sound file as it is, exactly in the shape it is outputted from the autonomous instrument. There are numerous possibilities for how it may sound, and the more so the longer it is. This is the context where we speak of preferences. Change a parameter and the sound file will be different. By successively changing a parameter, families of sound files can be obtained that are all more or less similar. The parameter might influence timbre, texture and large scale form without it being possible to separate its effect on different levels. Such convoluted parametric dependencies are to be expected of any complex systems, autonomous instruments in particular.

Aesthetic preferences are not necessarily permanent and stable. Whilst working with an autonomous instrument, using it as a sketching apparatus that comes up with suggestions for a new 20 minute piece, one might perhaps not find the first sketch quite good enough. However, while listening to these suggestions one might forget the original idea if ever there were one. Most likely the Mozart myth does not apply to the majority of composers, i.e., the idea that the composer imagines the complete work from beginning to end and that the actual process of composition consists of nothing but jotting the music down on paper [[Mountain, 2001](#)].

The sketching apparatus may well come up with better ideas than those you had been able to think of by yourself. And when it arrives at something you might fancy using but could not have thought of yourself, is that not a tiny nudging of your preferences?

The working process can be schematically illustrated with two elevators; the sketching apparatus is one and the composer is the other. They take turns lifting the sketches to new heights, each of which are plateaus from where the other party can lift it further up, till it floats in mid-air, where it finally reaches the threshold of originality¹.

Footnote on repertoire

'Autonomous instruments' was not an established term with any connection to music when I began my research on them. Consequently, works for or by autonomous instruments are not easy to find. On the other hand, there are pieces

¹ Here 'threshold of originality' is to be understood in its legal sense. The corresponding Swedish term literally means 'height of work', a pun that is lost in translation.

that arguably fit more or less into the description, pieces that can be said to be exemplars of music made with autonomous instruments. However, there is a nexus of pieces connected by family resemblance that are often mentioned in surveys. Feedback is a crucial element [Sanfilippo and Valle, 2012], other keywords include cybernetics, complex systems, emergence and self-organization. Whenever one or more of these terms are used to describe a piece of music, it may indicate that ideas related to autonomous instruments are being explored.

Whereas the search for musical works involving autonomous instruments may not yield much at all, there is the similar term, semi-autonomous instruments, which occasionally has been used to designate certain kinds of interactive systems that usually combine machine listening with some kind of algorithmically generated response [Jordà, 2007]. A few canonical examples are George Lewis' *Voyager*, Agostino Di Scipio's *Hörbare Ökosysteme* and Gordon Mumma's *Hornpipe* (which is an analog interactive system). Network ensembles like The League of Automatic Composers and later The Hub should also be mentioned; in these networks computers are interconnected and take care of different musical functions. David Tudor's *Rainforest* series of installations and many other of his works are brilliant examples of how analog electronic components can be patched together into complex networks exhibiting a complicated behaviour which may be influenced, although it might be hard to control.

In these cases the computer or the analog circuits function as an improvising partner that, to some extent, introduces its own musical impulses, hence the autonomy. However, attempting to achieve a total autonomy within such interactive systems would mean that the computer did not have to respond to what the musician is doing, which would obviously spoil the interactivity. Nevertheless, semi-autonomous instruments were an important source of inspiration in my work on autonomous instruments. The crucial connection can be illustrated with the following metaphor. A musician, say, a flutist, plays a tone and immediately notices that its pitch is too high and adjusts his or her embouchure accordingly. This very rapid process depends on feedback from what you hear to what you do, and from what you do to what the instrument sounds like. My autonomous instruments utilize a similar feedback mechanism. A signal is generated and analysed, then signal descriptors are mapped to control functions that control the synthesis parameters. Although I have not tried to model the functioning of musicians, this kind of feedback mechanism is an important component in all kinds of automated self-regulating control systems. A musician, then, can be thought of partially as a self-regulating system just as much as an autonomous instrument can be self-regulating.

Judging from the apparent lack of repertoire of works produced with autonomous instruments, it is perhaps not a common strategy of music making. On the one hand, there is a lot of interactive music making where improvisation plays a decisive role. On the other hand, there is the acousmatic school where fixed pieces without live performers has been the main musical form. In acousmatic music, it has been rare to completely assign the composition to algorithms. If one should highlight some pieces that exemplify the principles of autonomous instruments, then those of Xenakis' pieces realized with his GENDYN program

are among the most convincing ones.

Digression on Feature Extractor Feedback Systems

In my dissertation, I focused on a particular class of feedback systems which, to the best of my knowledge, had not been studied before. These systems are composed of three components (see Figure 1): a signal generator (G), an analysis module (A) and a mapping or transformer (T). Systems of this kind are called *Feature Extractor Feedback Systems* (FEFS).

The signal generator is an oscillator or some compound synthesis model that generates an audio signal x_n which depends on the synthesis parameters π_n , where $n = 0, 1, 2, \dots$ is the time. The signal x_n is split and written to a sound file (or sent through the sound card to loudspeakers), and it is also fed into the analysis module. The purpose of the analysis module is to get access to various signal descriptors that relate to perceptual features [Peeters et al., 2011] such as the fundamental frequency, spectral centroid, amplitude, degree of noisiness, etc. The analysis module has access to some segment of the recently generated signal, say, the last L samples, and it outputs a stream of signal descriptors ϕ_n . For simplicity we may assume that the signal descriptors are generated at the audio sampling rate. Then the analysis module is a function

$$\phi_n = A(x_n, x_{n-1}, \dots, x_{n-L+1}) \quad (1)$$

where $\phi_n \in \mathbb{R}^p$ may be a vector of p signal descriptors analysed in parallel. Next, this stream of signal descriptors is converted to a format that is fit for use as synthesis parameters. If the signal generator has q synthesis parameters, the mapping has to be some function $T : \mathbb{R}^p \rightarrow \mathbb{R}^q$ such that

$$\pi_n = T(\phi_n). \quad (2)$$

Finally, the signal generator outputs the signal x_n in ways that depend on the current value of the synthesis parameters. It will be practical to include an internal state variable θ_n to represent, e.g., the phase of an oscillator which must be stored in memory until next time the generator calculates an output sample. Then the generator is a function

$$\begin{aligned} x_{n+1} &= G(\pi_n, \theta_n) \\ \theta_{n+1} &= \theta_n + f(\pi_n). \end{aligned} \quad (3)$$

The three equations (1, 2, 3) taken together define a dynamic system, or rather a huge class of systems, because now these abstract equations must be given a concrete interpretation by assigning to them specific signal generators, feature extractors and mappings. Several examples of autonomous instruments of the FEFS type are discussed in detail in my dissertation and in some papers².

² The dissertation and sound examples are archived at <http://urn.nb.no/URN:NBN:no-39603>.
See also <http://ristoid.net/research.html>.

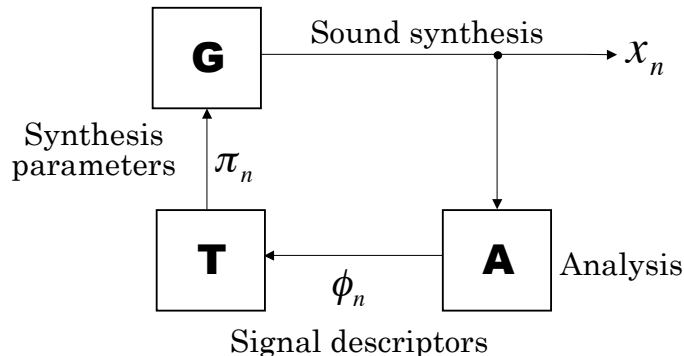


Figure 1: Generic Feature Extractor Feedback System (FEFS).

It is difficult to summarize all the peculiarities that can be found in this wide-ranging class of dynamical systems. However, it is worth emphasizing that these are deterministic systems where randomness plays no role, although their complexity can render the systems effectively unpredictable³.

Physicists and mathematicians have studied systems far simpler than the FEFS, such as cellular automata or low-dimensional chaotic systems, and with far greater rigour. The advantage of simplified models is that they can be more easily comprehended than a more complicated and realistic model. As far as scientists care about aesthetics and elegance, they often prefer the simple model. Although I have full sympathy for this search for simplicity, it turned out to be hard to construct any autonomous instruments that were at once simple and capable of exhibiting a complex and musically compelling behaviour.

Autonomous instruments of the FEFS variety can be designed in numerous ways, but there is a general observation that seems to hold in most cases. It turns out that the analysis window length L often has a significant impact on the sonic character. Using long windows, a long segment of the signal will be analysed, which tends to smooth out irregularities and sudden changes. Shorter windows follow details in the signal and easily give rise to wild, noisy, crackling sounds. An introductory transient usually can be heard as the system stabilizes. Intuitively, the transient's duration might be expected to be correlated with the

In this paper I refrain from discussing those of my own compositions that are related to autonomous instruments. Suffice it to mention that attempts in that direction have been carried out in some of the pieces in the *Signals and Systems* cycle, the most radical example of which is, perhaps, the 16 minute piece *Écriture Automatique*.

³There is no reason why autonomous instruments and FEFS should not include stochastic systems, but for simplicity only deterministic systems are considered here.




Figure 2: Transient durations in minutes as a function of a control parameter.

analysis window length, but that is not always the case.

In one particular and rather complicated FEFS, transients of various durations could be observed, sometimes lasting up to several minutes before the system eventually settled on a stable state in which the oscillator just generates a steady pitch. Figure 2 shows the result of an experiment in which the system was run at many different values of a control parameter for a maximum of 20 minutes. At some parameter values there are short-lived transients that last less than a minute, whereas at nearby parameter values the process may last over 20 minutes. Whether the system will ever stabilize cannot be known unless the system is run until it does—this peculiar situation closely parallels the halting problem in the theory of computation. Zooming in and varying the parameter over a smaller interval results in more of the same irregular changes between short transients and long lasting processes.

Unfortunately, the system that gives rise to these fascinating transients is relatively complicated; for a complete description, see my dissertation [Holopainen, 2012, p. 255 ff]. Could not similar phenomena emerge in far simpler systems? That seems possible, but the question is how to simplify the system. However, some of the simplest conceivable instances of FEFS show no remarkable behaviour at all lest you add further complicating mechanisms.

The physicist J. C. Sprott systematically investigated a large number of ordinary differential equations in order to try to find the most elegant chaotic systems [Sprott, 2010]. It is by no means evident how the notion of elegance is to be used in a rigorous study of dynamic systems. However, Sprott defines it more or less such that a system with fewer variables and simpler constants is more elegant than a system with more variables and more complicated constants. Using that definition, which is more stringent than it may appear from my description, he has been able to perform an automated search for elegant chaotic systems. A set of FEFS may perhaps be searched in a similar way. As far as elegance is concerned Sprott’s criteria will do, but if you are also looking for musically viable results there is no shortcut to listening to it all and making

your own decision.

Conceivably, the system’s dimensionality could be related to elegance, although that would be misleading. Some systems have many dimensions, yet can be concisely described, as is the case if there is some symmetry that simplifies them. In a FEFS, the state space typically has a very high dimension, because it has to have at least as many dimensions as the number of samples in the analysis window (e.g. $L = 1000$ samples). Although eq. 1 does not have to be rewritten just because a different window length is used, a shorter window may make the system more prone to become chaotic.

I have proposed as a quality criterion that an autonomous instrument should exhibit some sort of perceptually complex behaviour. In order to make the system’s dynamics more complex the transformer T needs to be designed in such a way that it amplifies small variations in the signal descriptors and then in turn causes greater variation in the synthesis parameters. Yet the variations should not be allowed to grow without bound, which would only cause the system to blow up and collapse. For a deterministic autonomous system to have those properties, it should be chaotic. However, it is even more interesting to try to design a system capable of switching between a range of behaviours and to exhibit variation over several temporal scales.

The system should surprise you, it should do more than expected. There is a catch, however, because as you refine a single autonomous instrument you may come to learn how it works. Then there will be less to be surprised by.

Fear-mongering

Norbert Wiener, the founder of cybernetics, warned against the unreflecting use of artificial intelligence early on, such as putting blind trust in a translation program [Wiener, 1961]. Computer music, particularly using machine listening, has seen a growing trend toward the use of artificial intelligence and autonomous systems. Although computer music may be a harmless domain for experiments, there has always been those who have feared that “the machines are taking over”⁴. Being able to play an instrument (in the sense of having achieved a high level of motor control) is no longer a requisite for creating music. Mental skills have replaced motor skills. Nevertheless, no more than digital printing has extinguished the art of etching have electronic musical instruments put an end to acoustic instruments. The remaining question is how far are we willing to pursue the automation of musical machines?

Nick Collins’ Autocousmatic program ⁵ is a nice example of automated composition [Collins, 2012]. The program loads a number of sound files from a folder; all it asks for is the duration of the composition and the number of chan-

⁴ This problem is not taken lightly by some researchers who are well aware of the things at stake. The [Machine Intelligence Research Institute](#) (MIRI) is a recently started academic centre that will deal with, among other things, the potential risks of self-improving artificial intelligence that might not be benevolent.

⁵ <http://www.sussex.ac.uk/Users/nc81/autocousmatic.html>

nels (the output can have up to eight channels) and then the program generates an acousmatic composition. The single and decisive choice left to the user is what sound files to put in the folder, but what the program will do with them cannot be influenced in any way. Whether or not a program like Autocousmatic will become popular among composers probably will depend less on its musical qualities than the amount of participation it affords on its way towards the finished work. In the case of Autocousmatic, the composer's participation is almost negligible although the result will depend on the selection of sound files that are fed into the program. For that reason, I should find it surprising if Autocousmatic or other similar programs would become widespread tools for composition. Most composers that I know would not be satisfied with building their music out of ready-made kits that they can assemble according to the instructions and put their signature on. However, being the creator of the software that automatically generates the music is an entirely different matter. In algorithmic composition, the composer becomes an author twice, first as programmer and then as the selector of the results that one accepts.

The complement of automatically generated music is artificial listening. In fact, Collins has a 'listening' agent in his program that analyses sketches of the final mix. In addition, Collins has analysed a few canonical acousmatic works in order to inform the shaping of the overall form. The artistic merits of Autocousmatic were evaluated by submitting some automatically generated compositions to festivals, but unfortunately all attempts were rejected. Thus, there is scope for improvements and every reason to expect better results before long.

Algorithmic composition aiming at the imitation of established styles, be it acousmatic music or Chopin, can be easily evaluated since there is an original to which the results can be compared. The algorithmic creation of something truly novel is more challenging both to achieve and to evaluate. The Russian cybernetician Zaripov suggested that the future, utopian project of artificial composition would not be limited to the imitation of existing styles, but would also involve the prediction of novel styles [Zaripov, 1969].

To make autonomous instruments more innovative and to diverge outside established stylistic norms, it would perhaps take a collaboration with a new kind of listener, artificial listeners who have a different perception than human listeners. At that point we might as well close the circuit, let the autonomous instruments play their music for their artificial audience and leave them alone to mind their own business.

References

- [Collins, 2012] Collins, N. (2012). Automatic composition of electroacoustic art music utilizing machine listening. *Computer Music Journal*, 36(3):8–23.
- [Holopainen, 2012] Holopainen, R. (2012). *Self-Organised Sounds with Au-*

- tonomous Instruments: Aesthetics and experiments*. PhD thesis, University of Oslo, Norway.
- [Jordà, 2007] Jordà, S. (2007). Interactivity and live computer music. In Collins, N. and d’Escriván, J., editors, *The Cambridge Companion to Electronic Music*, chapter 5, pages 89–106. Cambridge University Press.
- [Mountain, 2001] Mountain, R. (2001). Composers and imagery: Myths and realities. In Godøy, R. I. and Jørgensen, H., editors, *Musical Imagery*, chapter 15, pages 271–288. Swets and Zeitlinger.
- [Peeters et al., 2011] Peeters, G., Giordano, B., Susini, P., Misdariis, N., and McAdams, S. (2011). The timbre toolbox: Extracting audio descriptors from musical signals. *Journal of the Acoustic Society of America*, 130(5):2902–2916.
- [Sanfilippo and Valle, 2012] Sanfilippo, D. and Valle, A. (2012). Towards a typology of feedback systems. In *Proc. of the ICMC 2012*, pages 30–37, Ljubljana, Slovenia.
- [Sprott, 2010] Sprott, J. C. (2010). *Elegant Chaos. Algebraically Simple Chaotic Flows*. World Scientific, Singapore.
- [Wiener, 1961] Wiener, N. (1961). *Cybernetics: or Control and Communication in the Animal and the Machine*. The MIT Press, second edition.
- [Zaripov, 1969] Zaripov, R. K. (1969). Cybernetics and music. *Perspectives of New Music*, 7(2):115–154.