Composition with Computer Models of the Brain: An Alternative Approach to Music with Artificial Intelligence

Eduardo Reck Miranda

Abstract Artificial intelligence (AI) is aimed at endowing machines with some form of intelligence. Not surprisingly, AI scientists take much inspiration from the ways in which the brain—or the mind—works to build intelligent systems. This chapter proposes a different angle to harness the neurosciences for composition. Rather than building musical ANN to learn how to compose music, I shall introduce my forays into harnessing the behaviour of a type of neuronal model referred to as *spiking neuronal networks* to compose music. The discussion revolves around a piece for orchestra, choir and a solo mezzo-soprano entitled *Raster Plot*.

Keywords Creativity · Music-AI · Spike neurons · Orchestra · Performance

1 Introduction

Artificial intelligence (AI) is aimed at endowing machines with some form of intelligence. Not surprisingly, AI scientists take much inspiration from the ways in which the brain—or the mind—works to build intelligent systems. Hence, studies in philosophy, psychology, cognitive science and more recently, the neurosciences have been nourishing AI research since the field emerged in the 1950s, including, of course, AI for music (Miranda 2021).

The neurosciences have led to a deeper understanding of the behaviour of individual and large groups of biological neurones. We can now begin to apply biologically informed neuronal functional paradigms to problems of design and control, including applications pertaining to music technology and creativity (Magenta 2022). Artificial neuronal networks (ANN) technology owes much of its development to burgeoning neuroscientific insight.

However, this chapter proposes a different angle to harness the neurosciences for composition. Rather than building musical ANN to learn how to compose music, I shall introduce my forays into harnessing the behaviour of a type of neuronal model

E. R. Miranda (✉)
University of Plymouth, Plymouth, England
e-mail: eduardo.miranda@plymouth.ac.uk

C. Vear and F. Poltronieri (eds.), *The Language of Creative AI*, Springer Series on Cultural Computing, https://doi.org/10.1007/978-3-031-10960-7_13

referred to as *spiking neuronal networks* to compose music (Jang et al. 2019). The discussion revolves around a piece for orchestra, choir and a solo mezzo-soprano entitled *Raster Plot*.

2 Description of *Raster Plot*

Raster plot is a tribute to Plymouth-born explorer Robert Falcon Scott. It includes extracts from Scott's diary (Scott 2008) on the final moments of his expedition to the South Pole before he died in March of 1912; the extracts used in the piece are available in Appendix 1.

The mezzo-soprano sings the extracts using *sprechgesang*, a type of vocalization between singing and recitation: the voice sings the beginning of each note and then falls rapidly from the notated pitch, alluding to the endurance of Scott and his companions facing the imminent fatal ending of the expedition. A whispering choir echoes distressed thoughts amidst a plethora of jumbled mental activity represented by the sounds of the orchestra.

2.1 New Models

Inspired by the physiology of the human brain, I devised a method to represent the notion of mental activity musically. I used a computer simulation of a network of interconnected neurones that model the way in which information travels within the brain, to generate patterns that I subsequently turned into music. When the network is stimulated with an external signal (this will be clarified below), each neurone of the network produces sequences of bursts of activity, referred to as spikes, forming streams of rhythmic patterns. A *raster plot* is a graph plotting the spikes (Fig. 1): hence the title of the composition.

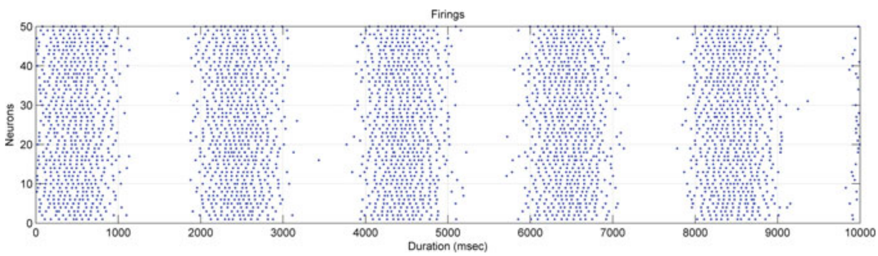

Fig. 1 A raster plot illustrating collective firing behaviour of a simulated network of spiking neurones. Neurone numbers are plotted (y-axis) against time (x-axis) for a simulation of 50 neurones over a period of ten seconds. Each dot represents a firing event

In a nutshell, I orchestrated raster plots by allocating each instrument of the orchestra to a different neurone of the network simulation. Each time a neurone produced a spike, its respective instrument was prompted to play a certain note. The notes were assigned based upon a series of chords, which served as frames to make simultaneous spikes sound in harmony.

The movement culminates with a transition from the orchestrated raster plots to a concluding passage bearing resemblance to a cathedral psalter chant. I wanted this to represent the moment Scott passed way; musically, it conveys a moment of *poiesis*: a moment of transition.

2.2 Music Neurotechnology

As briefly mentioned above, many recent advances in the neurosciences, especially in Computational Neuroscience, have led to a deeper understanding of the behaviour of individual neurones and their networks. I have coined the term *Music Neurotechnology* in a paper I co-authored for *Computer Music Journal* in 2009 (Miranda et al. 2009), to refer to a new research area that is emerging at the crossroads of Neurobiology, Engineering Sciences and Music. The compositional method described here is one of the outcomes of my continuing research in this field. Another important development in this area includes Brain–Computer Music Interfacing (BCMI) systems to enable persons with severe motor impairment to make music (Eaton et al. 2015).

The spiking neurones model that I used to compose the piece was originally developed by computational neuroscientist Eugene Izhikevich (2007). A biological neurone aggregates the electrical activity of its surroundings over time until it reaches a given threshold. At this point, it generates a sudden burst of electricity, referred to as an *action potential*. Izhikevich's model is interesting because it produces spiking behaviours that are identical to the spiking behaviour of neurones in real brains. Also, its equations are relatively easier to understand and program on a computer, compared to other, more complex models. Izhikevich's equations represent the electrical activity at the level of the membrane of neurones over time and can reproduce several properties of biological spiking neurones commonly observed in the brain.

The simulation contains two types of neurones, excitatory and inhibitory, which interact and influence the behaviour of the whole network. Each action potential produced by a neurone is registered and transmitted to other neurons, producing waves of activation, which spread over the entire network. A raster plot showing an example of such collective firing behaviour, taken from a simulation of a network of neurones, is shown in Fig. 1. Here, the spikes result from a simulation of the activity of a network of 50 artificial neurones over a period of ten seconds: the neurones are numbered on the y-axis (with neurone number 1 at the bottom, and neurone number 50 at the top) and time, which runs from zero to 10,000 ms, is on the x-axis. Every time one such neurone fires, a dot is placed on the graph at the respective time.

Figure 1 shows periods of intense collective spiking activity separated by quieter moments. These moments of relative quietness in the network are due to both the action of the inhibitory neurones and the refractory period during which a neurone that has spiked remains silent as its electrical potential decays back to a baseline value.

The network model needs to be stimulated to produce these patterns of activation. For the composition of *Raster Plot*, I stimulated the network with a sinusoidal signal that was input to all neurones of the network simultaneously. Generally speaking, the amplitude of this signal controlled the overall intensity of firing through the network. For instance, the bottom of Fig. 2 shows a raster plot generated by a network of spiking neurones stimulated by the sinusoid shown at the top of the figure. As the undulating line rises, the spiking activity is intensified. Conversely, as the undulating line falls, the spiking activity becomes quieter. As a gross generalization, if one thinks of the spiking neuronal network model as the brain of some sort of organism, then the stimulating sinusoid would represent perceived sensory information. Albeit simplistic, I find this model inspiring in the sense that it captures the essence of how our brain responds to sensorial information. Of course, a more complex signal could replace the sinusoid; for instance, a sampled sound could be used to simulate the network. In this case, the raster plots would look considerably more complex than the ones I am presenting in this chapter.

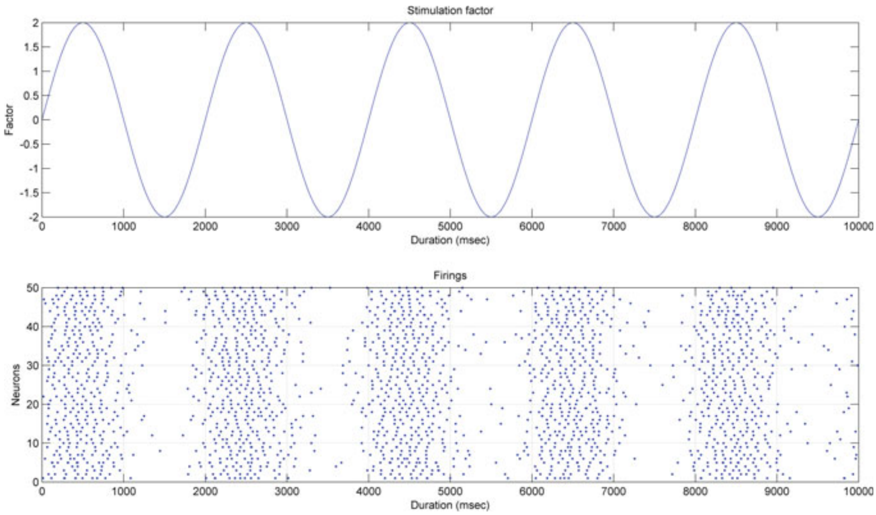

Fig. 2 At the top is a sinusoid signal that stimulated the network that produced the spiking activity represented by the raster plot at the bottom

3 Compositional Process

To compose the piece, I set up a network with 50 neurones and ran the simulation 12 times, lasting for 10 s each. For all runs of the simulation, I set the stimulating sinusoid to a frequency of 0.0005 Hz, which means that each cycle of the wave lasted for 2 s. Therefore, each simulation took five cycles of the wave, which can be seen at the top of Figs. 2, 3 and 4, respectively.

Compositionally, the top of Figs. 2, 3 and 4 suggests musical form to me, whereas the bottom suggests musical content. Hopefully, this will become clearer below as I unpack the process by which I composed this piece.

For each run, I varied the amplitude of the sinusoid, that is, the power of the stimulating signal, and the sensitivity of the neurones to fire. The power of the stimulating signal could be varied from 0.0 (no power at all) to 5.0 (maximum power) and the sensitivity of the neurons could be varied from 0.0 (no sensitivity at all; would never fire) to 5.0 (very sensitive). For instance, for the first run of the simulation, I set the power of the signal to 1.10 and the sensitivity of the neurons to 2.0 (Fig. 3), whereas in the tenth run I set these to 2.0 and 4.4, respectively (Fig. 1). One can see that the higher the power of the stimuli and the higher the sensitivity, the more likely the neurons are to fire and therefore the more spikes the network produces overall. One can observe a considerable increase in spiking activity in Fig. 4, which corresponds to the fourth run. And in Fig. 2, which corresponds to the tenth run, there is a substantial increase in the intensity of spiking activity. Table 1 shows the values for the 12 runs. I had envisaged at this stage a composition where the music

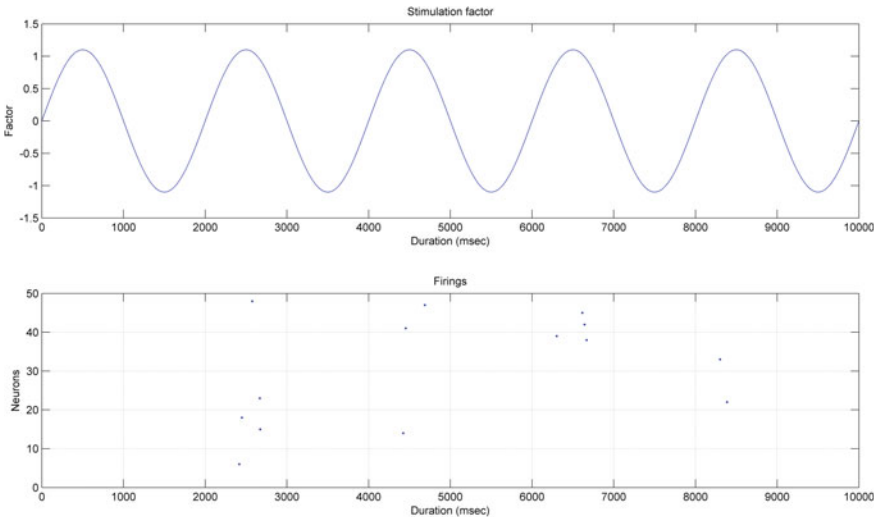

Fig. 3 First run of the simulation produced sparse spiking activity because the amplitude of the sinewave and the sensitivity of the neurons were set relatively low

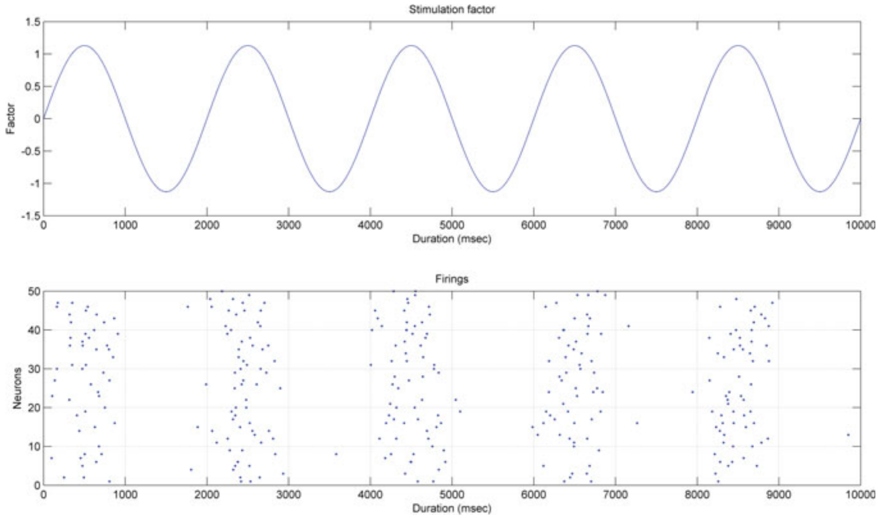

Fig. 4 Sensitivity of the neurons to fire was increased slightly in the fourth run of the simulations, resulting in more spiking activity than in previous runs

Table 1 Parameters for the 12 runs of the spiking neurones network

Run	1	2	3	4	5	6	7	8	9	10	11	12
Power	1.10	1.11	1.12	1.13	1.14	1.2	1.21	1.22	1.3	2.0	2.2	3.0
Sensitivity	2.0	2.3	2.6	2.9	3.2	3.5	3.8	4.0	4.2	4.4	4.8	5.0

would become increasingly complex and tense, culminating with the transition to the psalter-like chant I mentioned earlier.

I established that each cycle of the stimulating sinusoid would produce spiking data for three measures of music, with the following time signatures: 4/4, 3/4 and 4/4, respectively. Therefore, each run of the simulation would produce spiking data for fifteen measures of music. Twelve runs resulted in a total of 180 measures, but as we shall see below, I finished the spiking section at measure number 160. I felt that the resulting music was beginning to linger and loose interest at about this measure. Thus, the time was ripe for the transition to the psalter-like chant.

With the settings shown in Table 1, I noticed that the neurones did not produce more than 44 spikes in one cycle of the stimulating sinusoid. This meant that if I turned each spike into a musical note, then each cycle of the sinusoid would produce up to 44 notes. In order to transcribe the spikes as musical notes, I decided to quantize[1] them to fit a metric of semiquavers, where the first and the last of the three measures could hold up to 16 spikes each, and the second measure could hold up to 12. Next,

[1] To quantize means to restrict a variable quantity to discrete values. For example, an ordinary clock normally quantizes time to seconds; each tick of the clock corresponds to a second.

I associated each instrument of the orchestra, excepting the choir and the mezzo-soprano parts, to a neurone or group of neurones. This is shown in Table 2. From the 50 neurones of the network, I ended up using only the first 40, counting from the bottom of the raster plots upwards. Polyphonic instruments, such as the organ, were associated with a group of neurons because they can play more than one stream of notes simultaneously.

The compositional process progressed through three major steps:

(a) the establishment of a rhythmic template,
(b) the assignment of pitches to the template and
(c) the articulation of the musical material.

In order to establish the rhythmic template, firstly I transcribed the spikes as semiquavers onto the score. Figure 5 shows an excerpt of the result of this transcription for a section of the strings.

Although I could have written a piece of software to transcribe the spikes, I ended up transcribing the spikes manually. I printed the raster plots for each cycle of the stimulating signal (Fig. 6). Then, I used a template drawn on an acetate sheet to

Table 2 Instruments are associated with neurones

Neurones	Instruments	Neurones	Instruments
1	Contrabass 2	17	1st Violin 1
2	Contrabass 2	18, 19, 20, 21	Organ
3	Cello 3	22, 23, 24, 25, 26	Celesta
4	Cello 2	27, 28, 29	Vibraphone, Timpani
5	Cello 1	30	Snare drum, Cymbal, Tam-tam
6	Viola 3	31	Tuba
7	Viola 2	32	Trombone 3
8	Viola 1	33	Trombone 2
9	2nd Violin 4	34	Trombone 1
10	2nd Violin 3	35	Trumpet 2
11	2nd Violin 2	36	Trumpet 1
12	2nd Violin 1	37	Horn 3
13	1st Violin 5	38	Horn 2
14	1st Violin 4	39	Horn 1
15	1st Violin 3	40	Clarinet bass clarinet
16	1st Violin 2		

Each instrument plays the spikes produced by its respective neurone or group of neurones

Fig. 5 Transcribing spikes from a raster plot as semiquavers on a score

establish the positions of the spikes and transcribe the information into the score (Fig. 7).

To forge rhythmic patterns that would be recognized as such by performing musicians, I altered the duration of the notes and rests, while preserving the original spiking pattern as much as I could. Figure 8 shows the new version of the score

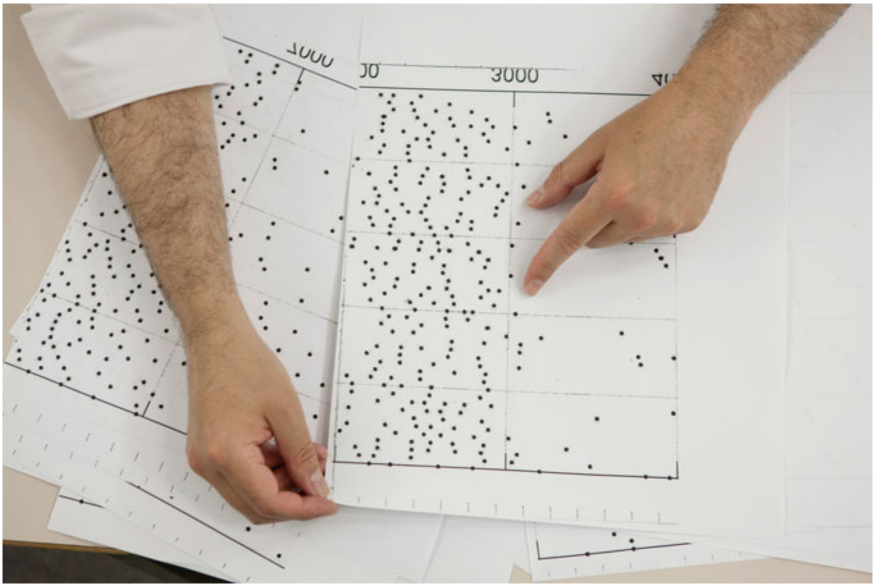

Fig. 6 A raster plots for each cycle of the stimulating signal produce spiking data for three measures of music

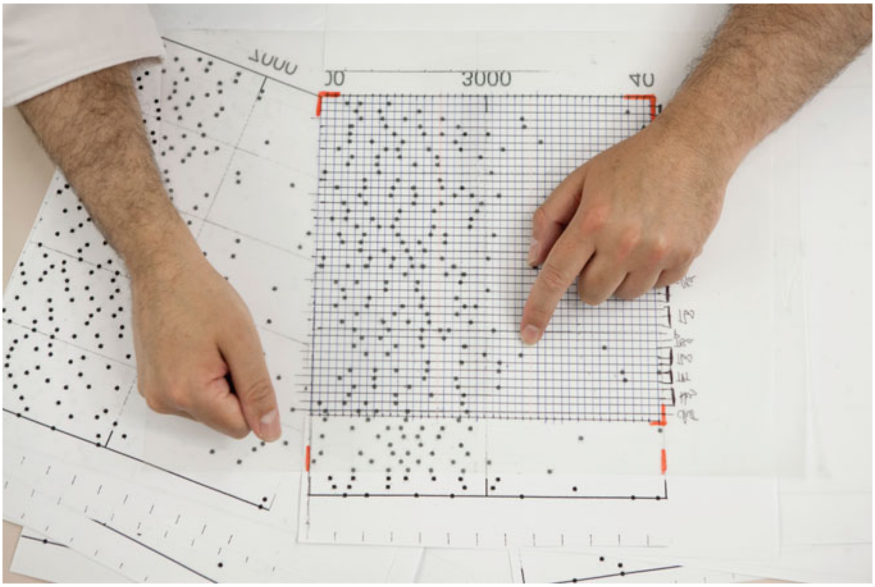

Fig. 7 A template drawn on an acetate sheet was used to transcribe the spikes into the score

shown in Fig. 5 after this process. Figure 9 shows the result of the compositional process, with pitches and articulation.

I would say that in many ways the compositional method that I developed for raster plot draws on Pierre Boulez serialism (Griffiths 1979). In order to assign pitches to the rhythmic template, I defined a series of 36 chords of 12 notes each, as shown in Fig. 10. These chords sprang on the back of the napkin after a conversation I had with composer Peter Nelson on a rainy afternoon in a café in Edinburgh. I had mentioned to him that I was struggling to find a decent way to assign pitches to the spiking rhythms. Peter suggested using matrices representing harmonic topologies. It was a eureka moment.

I started by creating 12 chords based on the harmonic series. Then, I established additional 24 chords firstly by inverting only a portion of those 12 chords (e.g., only the notes on the G clef) and then by inverting chords entirely. I do not remember the exact rationale for the different key signatures; most probably, I defined them in haste in order to avoid having to write all accidents next to the respective notes on the score.

To begin with, I used the first chord of the series to furnish pitches for the first 21 measures of music. As the spiking activity up to this point was not so intense, I decided to use only this chord to begin with. Then, from measure 22 onwards I used each subsequent chord of the series to furnish pitches for every three measures, and so on. Once I had furnished the pitches for measures 124–126 with the 36th chord, I subsequently selected chords unsystematically to continue the process until measure number 160.

Fig. 8 Resulting rhythmic figure

Fig. 9 Resulting music

Fig. 10 Series of chords for the harmonic structure of *raster plot*

The actual allocation of pitches of the chords to notes of the rhythmic figures was arbitrary. I did this differently as the movement progressed. In general, those figures to be played by instruments of lower tessitura were assigned the lower pitches of the chords and those to be played by instruments of higher tessitura were assigned the higher pitches, and so on. An example is shown in Fig. 11, which shows the allocation of pitches from the G clef portion of chord number 22 to the rhythmic figures for

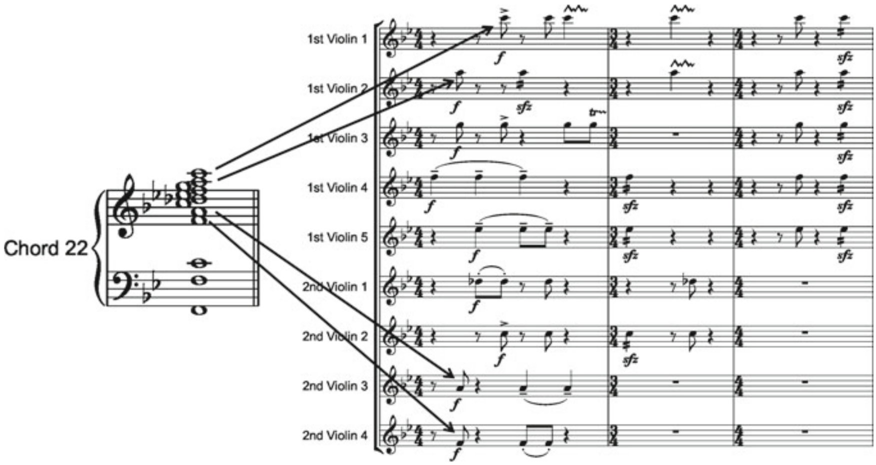

Fig. 11 An excerpt from *raster plot* illustrating the assignment of pitches to rhythmic figures

the violins in measures 82–84. There were occasions where I decided to transpose pitches one octave upwards or downwards in order to best fit specific contexts or technical constraints of the respective instrument. Other adjustments also occurred during the process of articulating the musical materials.

3.1 Limitations

A caveat of my method to turn raster plots into music is that it limits my ability to compose with the parameters composers would normally expect to work with, that is, duration and pitch. In a way, this limitation forced me to work with other musical parameters, such as articulation and timbre, to fashion the materials. To this end, I employed several non-standard playing techniques to forge new musical gestures.

The process of articulating the musical material is a difficult one to explain objectively because it was much less systematic than the processes described thus far. The vocal part was composed at the same time as I worked on the articulations. But it was not directly constrained by the spiking neurones method. The mezzo-soprano, which sings in *sprechgesang mode*, appears in measures corresponding to periods of rarefactive spiking activity. Musically, I wanted to create an effect akin to responsorial singing. Metaphorically, I wanted to allude to an imaginary process, whereby the neurones were sending commands to control the muscles of the vocal mechanisms of a hypothetic singer, but not so efficiently. Hence, the undefined effect of hearing neither clear singing nor clear speaking. The bass clarinet often doubles the mezzo-soprano, representing the hypothetic singer's mind's ear; it plays the melodic lines she intends to sing. Technically, this aids the singer to find the right pitch to enter passages that are difficult to ascertain the pitch unaided.

4 Reflection on Process

By way of introspection, I often find myself confronting the following dichotomy whenever I attempt to articulate my compositional practice. On the one hand, I think of music as the intuitive expression of ineffable thoughts, highly personal impressions of the world around me, and the irrational manifestation of emotions. On the other hand, I am keen to maintain that music should be logical, systematic, and follow guiding rules. In general, I think that rationality does play an important role in music composition, especially classical music. Hence, formalisms, rules, schemes, methods, number crunching, computing, and so on, are of foremost importance for my *métier*: But I also think that music that is totally generated automatically by a machine is rather meaningless. Music needs to be embedded in cultural and emotionally meaningful contexts, which composers express in subtle, often ineffable ways. A computer would not be capable of composing a piece such as Beethoven's *Symphony No. 9*. Its backstory, myriad of references, drama, and so on, are aspects of musicianship that computers, as we know them today, cannot grasp. The composition of *Raster Plot* is a good example of this dichotomy.

All the same, one of the reasons I find it exciting working with artificial intelligence, and computers in general, is because they can generate musical materials that I would not have produced on my own manually. This mindset is akin to John Cage's thinking when he preferred to set up the conditions for music to happen rather than composing music set in stone. Cage liked being surprised by the outcomes of such happenings (Cage 1994). By the same token, I enjoy being surprised by the outcomes of a computer. But I am not willing to just leave these materials intact I am afraid.

A recording of the premiere of *raster plot* by Ten Tors Orchestra, under the baton of the late Simon Ible, is published by Da Vinci and a free version is available on YouTube.[2] A short excerpt of the score is shown in Appendix 2.

Appendix 1 The Lyrics for Raster Plot

For God's sake,

look after our people.

Had we lived,

I should have had a tail to tell,

of the hardihood endurance and courage

[2] *Raster Plot* on YouTube: https://www.youtube.com/watch?v=xEywlAbP8Vs

of my companions,
which would have stirred
the heart of every Englishman.
These rough notes
and our dead bodies
must tell the tale.
We shall stick it out
to the end,
but we are getting weaker.
Of course
and the end cannot be far.
It seems a pity,
but I do not think
I can write more.
For God's sake,
look after our people.

Appendix 2 Excerpts from the Score

See Figs. 12 and 13.

Fig. 12 Page 17 from the score

Fig. 13 Page 20 from the score

References

Cage J (1994) Silence: lectures and writings. London, Marion Boyars. ISBN 978-0714510439

Eaton J, Williams D, Miranda E (2015) The space between us: evaluating a multi-user affective brain–computer music interface. Brain Comput Interfaces 2:103–116. https://doi.org/10.1080/2326263X.2015.1101922

Griffiths P (1979) Boulez. Oxford University Press. ISBN: 978–0193154421.

Izhikevich ER (2007) Dynamical systems in neuroscience. Cambridge (MA), The MIT Press. ISBN 978-0262090438

Jang H, Simeone O, Gardner B, Grüning A (2019) An introduction to spiking neural networks: probabilistic models, learning rules, and applications. https://arxiv.org/abs/1910.01059

Magenta Studio. https://magenta.tensorflow.org/studio-announce

Miranda ER (ed) (2021) Handbook of artificial intelligence for music. Springer. ISBN: 978–3–030–72116–9

Miranda ER, Bull L, Gueguen F, Uroukov IS (2009) Computer music meets unconventional computing: towards sound synthesis with in vitro neural networks. Comput Music J 33(1):9–18

Scott RF (2008), in: Jones M (ed) Journals captain Scott's last expedition. Oxford University Press. ISBN: 978–0199536801