LIGHT IMAGE IMAGINATION

CHAPTER 6

Computer-Aided Musical Imagination

Eduardo Miranda

Perhaps one of the most significant aspects differentiating humans from other animals is the fact that we are inherently musical. Our compulsion to listen to and appreciate sound arrangements beyond the mere purposes of linguistic communication is extraordinary. From the discovery almost three thousand years ago of the direct relationship between the pitch of a note and the length of a string or pipe to the latest computer models of human musical cognition and intelligence, composers have increasingly looked to science to provide new and challenging ways to study and compose music.

Music is generally associated with the artistic expression of emotions, but it is clear that reason plays an important role in music making. For example, the ability to recognise musical patterns and to make structural abstractions and associations requires sophisticated memory mechanisms, involving the conscious manipulation of concepts and subconscious access to intuitive knowledge. One of the finest examples of early rational approaches to music composition appeared in the 11th century, when Guido d'Arezzo proposed a lookup chart for assigning pitch to the syllables of religious hymns. He also invented the musical stave for systematic notation of music and established the medieval music scales known as the church modes.

Any attempt at distinguishing the rational from the intuitive in musical composition needs to take into account the music technology of the time. Between d'Arezzo's charts and the first compositional computer programs that appeared in the early 1950s, countless systematisations of music for composition purposes were proposed. The use of the computer as a composition tool thus continues the tradition of Western musical thought that was initiated approximately a thousand years ago. The computer is a powerful tool for the realisation of abstract design constructs, enabling composers to create musical systematisations and judge whether they have the potential to produce interesting music. A pertinent question comes to mind here: To what extent do composers think differently when composing with computers as opposed to earlier compositional practices, such as the classical picture of the inspired composer working on the piano with pencil and stave paper?

There probably are as many answers to the above question as there are composers. The role of the computer in my own compositional practice has oscillated between two extremes: on the one hand, I have simply assumed the authorship of compositions that were entirely generated by a computer, albeit programmed to follow my instructions. On the other hand, I have composed with pencil on stave paper, using the computer only to typeset the final score. I shall argue that both approaches to composition are not incompatible; instead, they are manifestations of creative processes that are becoming progressively more polarized due to increasingly sophisticated technology. As with the need to understand the state of the art of music technology in order to distinguish the rational from the intuitive in musical composition, I believe we would also need to articulate the notion of cognition in order to discuss the role of technology in musical creativity. I would argue that an important act of cognition in musical creativity is imagination.

Imagination in music can be many things, but here I shall argue that it is something that involves a great deal of abstraction. In a paper I published recently in the journal *Organised Sound* (Miranda 2010), I attempted to shed light on the hypothesis that musical imagination is a by-product of the inherent abstracting and predicting properties of the brain. The processing of music in the brain is an incredibly complex affair, one which is still not well understood. It is generally agreed, however, that the brain employs hierarchical neural structures to process music (Griffiths et al. 1998; Stewart et al. 2008) and these processes may not necessarily happen sequentially. For instance, it has been suggested that some higher-order structure processes the contour of melodies, while some lower-order structure processes their pitches. Assuming that we come to an understanding that the notion of a melodic contour is more abstract than the notion of a sequence of pitch values, this illustrates what abstraction might be. Another example provided in this chapter concerns the notions of beat and metre. The perception of rhythm is structured by beat and metre induction mechanisms. Our brain always tries to infer an underlying regular beat in a sequence of tones. Even in a sequence of absolutely uniform tones (i.e. same pitch, duration, loudness and timbre), the brain would infer a beat by imposing a metric template on the perceived signal. This phenomenon does not seem to be solely dependent on training or attention, which suggests that such metric template is a high-level abstraction emerging from some low-level biological feature of the brain. Such mechanisms for abstracting higher-level musical structures in response to avalanches of lower-level auditory information pervade our brain when we listen to music.

In short, the brain is a complex distributed processing system, with various structures operating concurrently and at different time scales, from short-term to long-term musical forms. Whereas lower-level structures may take care of processing the pitches of a sound sequence, higher-level structures would take care of processing the melodic contour engendered by the pitches of those sounds. But these processes might not necessarily be bottom-up; higher-level structures in the brain may make estimations of how the contour should evolve and this may influence how lower-level structures process pitches.

The amount of information that flows in the brain is immense. Obviously, the brain is in charge of running our entire body and therefore it will be engaged in a number of other vital tasks while we listen, play or indeed imagine music. It is unlikely that the brain would process such tasks completely unconnected from each other. Brain resources are shared. The brain cannot afford the delay that it would take to wire from scratch billions of neurons for every function it has to perform. We have evolved strategies to react to sensations as quickly as possible. One of the strategies that evolved in the brain to deal with huge amounts of information flow and minimise reaction delays is the ability to make predictions, or anticipation.

Neuroscientists generally agree that the brain is often prepared in advance by the very first incoming signals for how it will react prior to actually processing the whole lot of sensory information that is coming in. Concerning auditory processing, our soundscape is normally composed of several simultaneous sources. It is therefore important to keep track of sound sources by building representations to distinguish between the sounds streaming from the same source and the sounds originating from different sources. The brain needs to evaluate how well incoming sounds fit within the existing representations, because the arrival of a sound that cannot be deemed as a continuation of any of the previously registered streams indicates either the beginning of a new source or a change in the activity of an existing source. In order to do this, the brain needs to build predictive models, whose purpose is to estimate patterns in the incoming stimuli. These predictive models allow the brain to interact with the world efficiently.

The brain is wired up to actively detect patterns in auditory input. As we listen to music, our brain will continuously seek for regularities in the incoming stimuli. A range of features, or combinations of them, define these regularities and they are extracted at many different levels and time scales. The brain may even make up something if necessary; for example, by imposing a metric template on a sequence of entirely uniform tones. Such a metric template is not in the signal; it is in the brain. Building predictive models of the incoming sensory input through the extraction of regularities, towards emergent (and not-so emergent) abstractions, is a fundamental aspect of cognition. By adapting to patterns in the world, the brain becomes more sensitive to stimuli that differ from those implied by the detected regularities. Such different signals excite the brain to refine its representations to more closely match the sensory experience. In this way, we construct models of the world, which are increasingly more specialised. Therefore, intrinsic innate processing strategies combined with evolving experience drive our impelling force to organise sound in the mind.

In a nutshell, the brain is a predictive organ that strives to find or impose structure on sensory information. In order to do this efficiently, it needs to make abstractions to fuel relentless processes of making internal representations of the world. Behind these processes, there is an impelling force to organise sensory information, which is driven by the physiological nature of our brain and its own evolving internal representations, or models, of the world. Therefore, imagination is likely to be a by-product of this mechanism. But how can technology harness musical imagination? – I suggested above that my creative processes involve practices that are becoming progressively more polarised due to use of technology. What does this mean?

One thread that I am currently contemplating to address the question above explores an idea suggested by philosopher Friedrich Nietzsche. Nietzsche (2003) suggested that great artistic creations could only result from the articulation of a mythological dichotomy referred to as the Apollonian and Dionysian. In ancient Greek mythology, Apollo is the god of the sun and is associated with rational and logical thinking, self-control and order. Conversely, Dionysus is the god of wine and is associated with irrationalism, intuition, passion and anarchy. These two gods represent two conflicting creative drives, constantly stimulating and provoking one another. As I understand it, this process leads to increasingly high levels of artistic and scientific achievements. Although dating from the 19th century, this notion still compels me. Ian McGilchrist (2009) has recently discussed this dichotomy in great detail. Below is an attempt at contextualizing it in my own compositional practice.

One side of me is very methodical and objective, keen to use automatically generated music, computer systems, formalisms, models and so on. Conversely, another side of me is more intuitive, emotional and metaphorical. Each side has it own agenda, so to speak, but they are not unrestrained. They tend to inhibit each other: the more I attempt to swing to the Apollonian side, the stronger is the Dionysian force that pulls me to the opposite side, and vice versa.

Nietzsche would not normally be my first choice of philosopher when seeking contemporary explanations for music cognition, but it turns out that the 19th-century Apollonian vs. Dionysian dichotomy resonates remarkably well with the way in which neuroscientists think our brain works. (Springer and Deutsch 1997; Davidson 1996) There are parts of the human brain that are undeniably Apollonian, whereas others are outrageously Dionysian. The Apollonian brain includes largely the frontal lobe of the cortex and the left hemisphere. Generally, these areas are in charge of focusing on attention to detail, seeing wholes in terms of their constituents, and making abstractions. They are systematic and logical. The Dionysian brain includes sub-cortical areas, which are much older in the evolutionary timeline, and the right hemisphere. It is more connected to our emotions. It perceives the world holistically and pushes us towards unfocused general views. The Apollonian brain is concerned with unilateral meanings, whereas the Dionysian brain tends to forge connections between seemingly unrelated concepts.

The notion that the Apollonian and the Dionysian tend to inhibit each other reminds me of the way in which the brain functions. Inhibitory processes pervade the functioning of our brain at all levels, from the microscopic level of neurones communicating with one another, to the macroscopic level of interaction between larger networks of millions of neurones. Indeed, this dichotomy also reminds me of the aforementioned interactions between low-level and high-level brain structures for processing music. In this context, I believe that the further my Apollonian brain pushes me to perceive the world according to its agenda, the stronger the pull of my Dionysian brain to perceive the world differently. Hence, computer technology is of foremost importance for my métier, because it allows me to stretch my Apollonian musical side far beyond my ability to do so by hand, prompting my Dionysian side to counteract accordingly. The composition of 'Evolve', the second movement of my symphonic piece *Mind Pieces* is discussed below as an example of this.

Mind Pieces is a five-movement-long symphonic piece for orchestra, percussion and prepared piano. It premiered at the Peninsula Arts Contemporary Music Festival on 12 February 2011 in Plymouth, by Ten Tors Orchestra, conducted by Simon Ible. Albeit not necessarily obvious to the listener, there was a great deal of Apollonian processes in the composition of 'Evolve'. I started with a set of computer-generated rhythms, which were generated by means of a simulation of evolution and transmission of rhythmic memes; memes are the cultural equivalent of a gene, a term coined by Richard Dawkins (1989). I collaborated with João Martins, then a doctoral student at the Interdisciplinary Centre for Computer Music Research (ICCMR), to develop 'A-rhythms', an A-life-based system to compose rhythms based on a paradigm that we have been working with at ICCMR known as imitation games. (Miranda 2002) In a nutshell, we developed a system whereby a group of software agents evolve repertoires of rhythms by interacting with each other. Software agents are virtual entities – or software robots – programmed to execute tasks. They often are embedded with some form of intelligence and can perform tasks independently from each other, without supervision from a central control.

In 'A-rhythms', the agents were programmed to create and play rhythmic sequences, listen to each other's sequences, and perform operations on those sequences, according to an algorithm referred to as 'the rules of the game'. To begin with, each agent is set up with an initial rhythm stored in its memory. These initial rhythms are randomly generated and are different for each agent. As the agents interact with each other, they can add new rhythms to and/or erase rhythms from their memories, and modify existing rhythms. The aim of the game is to develop a shared lexicon of rhythmic patterns collectively. As the interactions take place, each agent develops a repertoire of rhythms similar to the repertoires of its peers. The agents interact in pairs and at each round, one of the agents plays the role of a player and the other the role of a listener. The agents count the number of times they play each rhythm stored in their memories. This counter is referred to as the popularity of the rhythm. The following algorithm is the core of the rules of the game:

> *Player:*
> P1. Pick a rhythm from its memory and play it.
>
> *Listener:*
> L1. Search the memory for a rhythm that is identical to the rhythm produced by the agent player.
> L2. If an identical rhythm is found, then increase its popularity and give a positive feedback to the agent player.
> L3. If an identical rhythm is not found, then add this rhythm to its memory and give a negative feedback to the agent player.
>
> *Player:*
> P2. If the listening agent's feedback was positive, then increase the popularity of the played rhythm.
> P4. If feedback is negative, then decrease the popularity of the played rhythm.
> P5. Perform memory updates.

After each interaction, the agent player performs a number of updates. For instance, from time to time, the agent may delete the rhythm in question if its popularity remains below a minimum threshold for a given period of time.

This means that other agents probably do not share this rhythm and therefore should no longer be used. Also, from time to time the agent may transform the rhythm. This is decided based on a number of factors; for example there is a variable, referred to as the transformation counter, which is updated in terms of its popularity. The more popular a rhythm is, the more likely the agent would transform it. Furthermore, the agents are programmed with a memory-loss mechanism, whereby after each interaction all the rhythms have their popularity decreased by a specified amount.

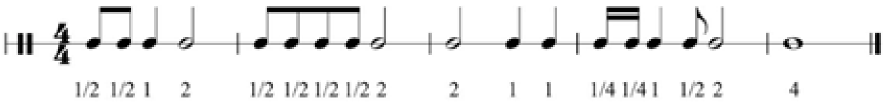

Fig. 1
Standard music notation of a rhythmic sequence and its corresponding inter-onset representation.

The agents store rhythms as a sequence of inter-onset intervals, represented in terms of small-integer ratios of an isochronous (regular) pulse [Figure 1]. At the core of the mechanism by which the agents develop rhythmic sequences are transformation operations. The transformation operations are as follows:

Divide a rhythmic figure by two (e.g. ½ = ¼ + ¼)
Merge two rhythmic figures (e.g. ½ + ½ = 1)
Add one element to the sequence
Remove one element from the sequence

The agents are programmed with the ability to measure the degree of similarity of two rhythmic sequences. This measurement is used when a listener searches for identical rhythms in its repertoire. If the degree of similarity of two rhythms is within a given threshold then the rhythms are considered identical. This threshold is set beforehand. Therefore, two rhythmic sequences do not necessarily need to be exactly equal to be considered identical. This method to measure the similarity is detailed in a paper presented at the 10th Brazilian Symposium on Computer Music. (Martins et al. 2005) We ran ‘A-rhythms’ systematically with different parametric values in order to observe the behaviour of the agents under a number of different conditions. We observed the emergence of repertoires across the agents, some of which were more coherent than others. Also, the size of the repertoires varied. These varied according to the number of agents in a group, and thresholds for probing the popularity and transformation counters mentioned earlier. (Martins et al. 2008)

For the composition of 'Evolve', Martins and I ran simulations with 3, 10 and 50 agents, for 5,000 interactions or so each. At the end of the simulations we opened the memories of the agents and picked those rhythmic patterns that all of them had evolved in common. Then I loaded these patterns into a music notation editor and sequenced them. I had no plans for how the composition would develop from here. I auditioned the sequence on various timbres hoping for an idea to emerge. When I played them on a snare drum, my Dionysian brain somehow connected it to Maurice Ravel's orchestral piece

Fig. 2
An excerpt from 'Evolve', bars 234-239. Only the upper part of the full orchestral score is shown.
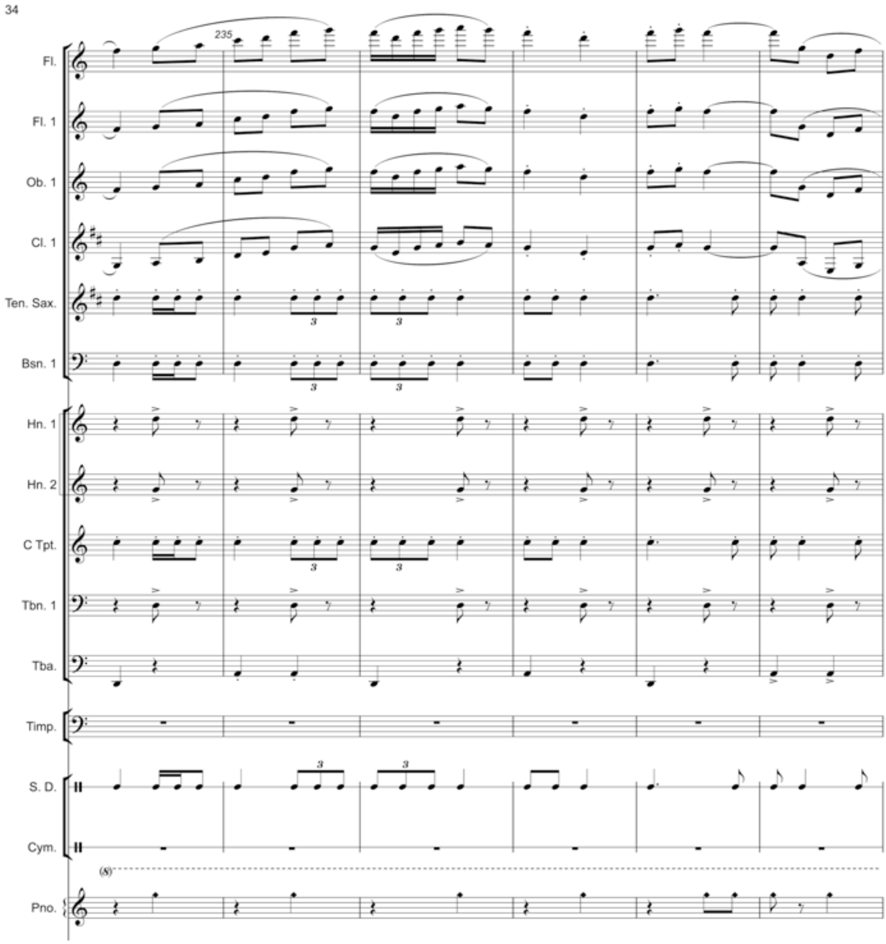
Boléro and made a split-second decision: to use the rhythmic sequence to form the backbone of the entire movement and to base the orchestration of the entire movement on that of *Boléro*. As my Apollonian side strived to be as systematic as possible by following the orchestration scheme laid out by Ravel, my Dionysian brain brought in melodic lines and themes whose origins I am unable to ascertain. I speculate that they were musical ideas lurking deep in my memory. Figure 2 shows an excerpt of 'Evolve'. The computer-generated rhythm played on the snare drum (S.D.) is doubled by the saxophones (Ten. Sax.), bassoons (Bsn.1) and trumpets (C Tpt.).

My musical imagination therefore does seem to be driven by a push-and-pull embodied by the aforementioned dichotomy between reason and intuition. I would probably never have had the idea of basing the orchestration of 'Evolve' on that of Ravel's *Boléro* if I had not worked with those computer-generated rhythms. However, I feel that whereas my Apollonian side might probably be able to compose music on its own right, my Dionysian side would not be able to do so. The latter needs the aid of the former. In this sense it can be proposed that technology mediates the embodiment of imagination.

BIBLIOGRAPHY

Davidson, Richard. 1996. *Brain Asymmetry*. Cambridge, MA: MIT Press.

Dawkins, Richard. 1989. *The Selfish Gene*. Oxford: Oxford Paperbacks (2nd Revised Edition).

Griffiths, T., et al. 1998. ‘Analysis of temporal structure in sound in the brain’. *Nature Neuroscience* 1: 422-27.

Martins, J., and Miranda, E.R. 2008. ‘Engineering the Role of Social Pressure: A New Artificial Life Approach to Software for Generative Music’. *Journal on Software Engineering*, Vol. 2, No. 3: 31-42.

Martins, J.M., et al. 2005. ‘Similarity measures for rhythmic sequences’. Proceedings of the 10th Brazilian Symposium on Computer Music (SBCM), Belo Horizonte (Brazil).

McGilchrist, Ian. 2009. *The Master and His Emissary: The Divided Brain and the Making of the Western World*. New Haven, CT: Yale University Press.

Miranda, E.R., and Biles, J.A. 2007. *Evolutionary Computer Music*. London: Springer.

Miranda, Eduardo Reck. 2010. ‘Organised Sound, Mental Imageries and the Future of Music Technology: A Neuroscience Outlook’, *Organised Sound*, 15(1): 13-25.

—. 2002. ‘Mimetic development of intonation’, in *Proceedings of the 2nd International Conference on Music and Artificial Intelligence – Lecture Notes on Artificial Intelligence*. London: Springer.

Nietzsche, Friedrich. 2003. *The Birth of Tragedy out of the Spirit of Music*. London: Penguin Classics (New Edition).

Springer, Sally P., and Deutsch, Georg. 1997. *Left Brain, Right Brain: Perspectives from Cognitive Neuroscience*. Cranbury, NJ: W.H. Freeman & Co Ltd (5th edition).

Stewart, L., et al. 2008. ‘fMRI Evidence for a Cortical Hierarchy of Pitch Pattern Processing’, *PLoS ONE* 3(1): e1470. doi:10.1371.