CLASSICAL MUSIC FUTURES
PRACTICES OF INNOVATION

EDITED BY
NEIL T. SMITH, PETER PETERS,
AND KAROLY MOLINA

https://www.openbookpublishers.com

Any digital material and resources associated with this volume may be available at https://doi.org/10.11647/OBP.0353#resources.

ISBN Paperback: 978-1-80511-073-6
ISBN Hardback: 978-1-80511-074-3
ISBN Digital (PDF): 978-1-80511-075-0
ISBN Digital eBook (EPUB): 978-1-80511-076-7
ISBN XML: 978-1-80511-078-1
ISBN HTML: 978-1-80511-079-8

DOI: 10.11647/OBP.0353

Cover image: Jean Pierre Geusens, *Focuss22* (2017)
Cover design by Jeevanjot Kaur Nagpal

18. Artificial Intelligence and the Symphony Orchestra

Robert Laidlow

Introduction: The Age of Artificial Intelligence

It is a cold day in December 2020. It is almost a year since British orchestras, and most others around the world, packed up their instruments for the unexpected and unwelcome hiatus caused by the pandemic. I am working in my study on music for orchestra and artificial intelligence (AI); or perhaps, I am working with artificial intelligence on a study for orchestra. To my left is my grandmother's old upright piano, missing the top two notes and the tuning slightly wonky (the tuner was not available to come in the brief period between national lockdowns). To my right my computer is teaching itself to write music, though it doesn't *know* it is writing music. It doesn't know what a sound is. In front of me the window is open. The computer's graphics card, essential for AI algorithms, produces more than enough heat for the little study and the quiet of the woods through the window is only broken by the continual whirr of cooling fans.

When the computer is off, I am no less plugged into a world of AI. It serves me advertisements for upcoming concerts it has learned that I like.[1] It changes the directions on my satnav based on live information from other drivers' phones. It summarises an article I cannot be bothered

1 Emmanuel Mogaji, Sunday Olaleye, & Dandison Ukpabi, 'Using AI to Personalise Emotionally Appealing Advertisement', in *Digital and Social Media Marketing*, ed. by N. P. Rana, E. L. Slade, G. P. Sahu, H. Kizgin, N. Singh, B. Dey, A. Gutierrez, & Y. K. Dwivedi (Cham: Springer, 2020), pp. 137-150.

 https://doi.org/10.11647/OBP.0353.18

to read, corrects the grammar in an email I send before I have had my morning coffee, recommends new films for me to watch.[2] In short, it is everywhere, but more importantly it is, or could be, *anywhere*.[3] That is the reality of a world that increasingly relies on algorithms to command the hidden infrastructures that support our society.

The effects of hidden AI are not always good, or even intentional. In contrast to those examples already mentioned, we have also seen recently that AI can accidentally replicate harmful biases, with recent cases including discrimination against women in the workplace, failure to recognise ethnic minorities in passport recognition software, and the automatic promotion of extremist social media groups to those already at risk of radicalisation.[4]

As a technology, AI also forms an integral part of wider issues that encompass many other technologies and social questions. The proposed 'metaverse', hypothesised to synthesise physical and digital experience seamlessly, relies upon AI in concord with other advanced technologies such as augmented reality and wearable technology.[5] Many researchers are focussed on AI as an essential tool to mitigate climate change,

2 Johann Lau, 'Google Maps 101: How AI helps predict traffic and determine routes', The Keyword, 3 September 2020, https://blog.google/products/maps/google-maps-101-how-ai-helps-predict-traffic-and-determine-routes/; Mico Tatalovic, 'AI writing bots are about to revolutionise science journalism: we must shape how this is done', *Journal of Science Communication* 17:1 (2018); Laksnoria Karyuatry., 'Grammarly as a Tool to Improve Students' Writing Quality: Free Online-Proofreader across the Boundaries', *JSSH* (*Jurnal Sains Sosial Dan Humaniora*) 2:1 (2018), 83–89; Carlos A. Gomez-Uribe & Neil Hunt 'The Netflix Recommender System', *ACM Transactions on Management Information Systems* (*TMIS*) 6:4 (2015).

3 Geoff Cox & Morton Riis, '(Micro)Politics of Algorithmic Music', in *The Oxford Handbook of Algorithmic Music*, ed. by A. McLean & R. Dean (Oxford: Oxford University Press 2018), pp. 603–626.

4 Jeffrey Dastin, 'Amazon Scraps Secret AI Recruiting Tool that Showed Bias against Women', in *Ethics of Data and Analytics* ed. by K. Martin (Boca Raton: Auerbach Publications, 2022), pp. 296–299; Leslie, David, 'Understanding Bias in Facial Recognition Technologies', *SSRN Electronic Journal* (2020); Hao, Karen, *He got Facebook hooked on AI. Now he can't fix its misinformation addiction* (MIT Technology Review, 2021), https://www.technologyreview.com/2021/03/11/1020600/facebook-responsible-ai-misinformation/

5 Lik-Hang Lee, Tristan Braud, Pengyuan Zhou, Lin Wang, Dianlei Xu, Zijun Lin, Abhishek Kumar, Carlos Bermejo, Pan Hui, *All One Needs to Know about Metaverse: A Complete Survey on Technological Singularity, Virtual Ecosystem, and Research Agenda* (2021), https://doi.org/10.48550/arxiv.2110.05352

mobilising it alongside more well-known political and social arguments.[6] Climate change is the principal example of what Timothy Morton calls a 'hyperobject': an idea so abstract and vast that it is difficult or impossible for any individual to grapple it in its entirety.[7] Like the case of climate change, it is entirely plausible that AI, which can absorb and analyse truly vast quantities of data, will become a critical tool in understanding other hyperobjects – and may become a hyperobject itself.

In short, I believe the relationship between humans and technology has never been more important, nor more complex. I share this view with others: increasingly there has been a push for closer scrutiny of AI as the techno-utopian vision of a wholly fair, omnibenevolent, quasi-magical algorithm has been disrupted and dismantled. AI is a powerful, exciting, and malleable technology, and it is here to stay. Society has always probed ideas, ideologies and technologies, and art has always been one tool in our arsenal to do so. Music is my way of making sense of these hidden technologies, my chosen method to expose and highlight issues that arise from AI, whether we know it is there or not.

Even taking that for granted, however, I am often asked: 'why classical music?' or 'why the orchestra?'. What relationship does, or can, AI have with the symphony orchestra? These are good questions, and at first glance it may appear that the two have little in common. Dig a little deeper, however, and one finds there exists a kinship between the AI and orchestra, in all its guises as ensemble, community, institution and historical concept. Having said that, I do not imply that other types of music do not share this relationship with AI, nor that classical music and the orchestra is the best type of music to explore the relationship; it is simply my way of doing so.

The relationship is more than one-way: it is not simply using orchestral music as a vessel to comment on a society dependent on AI. AI can, in return, offer the orchestra new perspectives, technologies and approaches to music-making that did not exist even a decade ago, even at the same time as certain actors within the music industry work towards

6 Josh Cowls, Andreas Tsamados, Mariarosaria Taddeo, & Luciano Floridi, 'The AI gambit: leveraging artificial intelligence to combat climate change—opportunities, challenges, and recommendations' in *AI and Society* 1 (2021), 1–25.

7 Timothy Morton., *Hyperobjects: Philosophy and Ecology after the End of the World* (Minneapolis: University of Minnesota Press, 2013).

replacing human musicians with indistinguishable AI performances. AI is both dangerous and full of potential, both good and bad, both hidden behind-the-scenes and on full display as a corporate buzzword. It has these dualities that are both alluring to artists and highly relevant to the orchestra, which is also, I would argue, partially reliant on a foundation of dualities and contradictions. Three particular dualities have stood out to me as important for discourse around designing and understanding AI, which are also areas to which musicians and orchestral institutions already give a great deal of thought. I have termed these:

1. Future and Past
2. Fake and Real
3. System and Secret

Each of these dualities is the focus of an individual movement of my piece *Silicon*, written for the BBC Philharmonic Orchestra as part of my doctoral research at the Centre for Practice & Research in Science & Music at the Royal Northern College of Music, Manchester UK. Through this chapter, I hope to offer an artistic reflection and provocation on the utility of the symphony orchestra as an investigative tool into the technology that is changing our world, as well as a case study of *Silicon*, showing how these abstract ideas might be transformed into music. This case study will also discuss the exciting creative possibilities available when incorporating this technology into a live performance: this is music both *about* AI and *using* AI.

I will give a brief explanation of exactly what I mean by AI, because the term is unhelpfully broad. After this, I will ground my own artistic work and research amongst other fields that it relates to. Outside of fiction and clickbait internet articles, AI in a research context usually means the design of algorithms that, through statistical methods, can 'learn' how to mimic human intelligence. Most AI is designed to excel at one specific task, not to have any form of *general* intelligence (the term for this field of research is AGI). There are several approaches one might take to designing an algorithm, depending on the task, but over the last decade much research focus has been on the development of machine-learning algorithms.[8] Most interaction between art and

8 Terry Sejnowski, *The Deep Learning Revolution* (Cambridge, M.A.: MIT Press, 2018).

AI currently uses machine learning, and for this chapter the terms AI and machine learning can be regarded as interchangeable. Machine-learning algorithms 'learn' from a dataset (in this chapter, the dataset will most often be music of some sort), encoding the rules they learn into a digital 'model'. This model will then be used to carry out the AI's task, be that writing music or directing traffic. Machine learning can be both 'supervised', where a computer programmer provides positive or negative feedback on the results, usually directing towards a desired result, or 'unsupervised', where the algorithm is left totally on its own to learn. It is difficult to speak about AI algorithms, especially regarding their decision-making process, without risking some level of anthropomorphism. For the avoidance of doubt, I want to make it clear that none of these algorithms 'think' or in any way can make conscious decisions as we would understand them.

There are many composers using algorithms in their music and the algorithmic music field is a lively and fascinating area of practice-based research.[9] There is a rich history of musicians engaging specifically with classical music utilising algorithms, from twentieth-century composers such as Iannis Xenakis[10] and Karlheinz Stockhausen, to musician-programmers such as Francois Pachet,[11] stretching back to the earliest pioneers in the field of what would become computer science, including Ada Lovelace. However, these algorithms are generally coded top-down, which means the composer/programmer has had complete control over the creation of the algorithm. Even if the composer leaves in elements of randomness or instability, this remains a creative decision they have made. By definition, this is not an AI algorithm, which in some way teaches itself and can subsequently make decisions responding to new information based on what it has previously learned. Since AI technology is relatively nascent, the area of the orchestra and AI remains an exciting and comparatively untapped field, with space for many perspectives and artistic approaches.

9 *The Oxford Handbook of Algorithmic Music*, ed. by Alex McLean, & Roger T. Dean (Oxford: Oxford University Press, 2018).

10 Iannis Xenakis, *Formalized Music* (Bloomington, I.N.: Indiana University Press, 1971).

11 Francois Pachet, 'The Continuator: Musical Interaction with Style', *Journal of New Music Research* 32:3 (2010), 333-341, http://dx.doi.org/10.1076/jnmr.32.3.333.16861

By contrast, there is fascinating AI being developed by computer scientists that relate to music (music-related AI), though these tools are not always primarily intended to be used by composers. Often, music-related AI's purpose is to emulate existing music as accurately as possible, using music as a yardstick through which we can audibly measure the progress of the machine-learning field as a whole. Of course, there are several inspiring exceptions, including the work of Google Magenta, which releases intuitive AI tools for use by amateur and professional musicians in the studio, the Wekinator, which explores AI as a way to control digital instruments, PRiSM's reimplementation of SampleRNN, which many composers have used in recent years to generate electronics tracks for new pieces, and recent work from IRCAM such as the RAVE algorithm.[12]

Future and Past (*Silicon Mind*)

The first duality to be discussed is the relationship between the future and the past. Specifically, this relationship is one of legitimacy: how does the past legitimise the future? In the field of music-related AI, this relationship is often implied through the data and evaluation methods frequently employed to judge the success of AI. Here I will focus on AI that *generates* music.

AI needs data to learn from and, in the case of music, that data usually comes either in the form of audio files or sheet music. In both cases, AI cannot create music out of nothing, it needs to learn from music that already exists. When coding generative AI, researchers usually use existing music both as a dataset from which the AI learns the rules of music and also as a yardstick against which to judge the quality of the AI's outputs. Marcus du Sautoy writes that 'Bach is the composer most composers begin [learning] with, but he is the composer most computers

12 On the Wekinator see: Rebecca Fiebrink, & Perry R. Cook, 'The Wekinator: a system for real-time, interactive machine learning in music', *The 11th International Society for Music Information Retrieval Conference* (2010), http://code.google.com/p/wekinator/; for further discussion of PRiSM see: Christopher Melen, 'PRiSM-SampleRNN', *RNCM PRiSM* (2019), https://www.rncm.ac.uk/research/research-centres-rncm/prism/prism-collaborations/prism-samplernn/

begin with too',[13] and indeed Bach is often the choice of dataset and generation for much recent research.[14] The more indistinguishable from Bach, the more successful the AI is deemed to be. This might lead to the impression that one of the main uses of generative AI could be to complete unfinished pieces by dead composers, as indeed we have seen in recent years with AI 'completions' of Beethoven and Schubert.[15] Relatedly, several recent AI algorithms outside of academia have been developed that automatically generate music for soundtracks or other media uses.[16] Crucially, these are designed to replace composers who write this music, not necessarily to create more interesting music. Once again, AI needs to sound like what already exists as much as possible, in this case to avoid the costs of human labour.

This leaves us in the situation that this technology, which promises the future, is consistently looking to the past to prove its legitimacy. Of course, there are exceptions to this rule, as some researchers do place genuine novelty and creativity at the heart of their research. Overall, however, it is easier for AI to function in this way because currently the dominant AI methodology is machine learning, which relies on existing data. I worry that such a relationship between past and future has the potential to hamper innovation. As this technology becomes more prevalent, it will become even more essential to have the skills to use it for more interesting and creative purposes than simple repetition and recreation.

This is where the orchestra comes in. The performance of established music is perhaps the genre's defining trait, evident in the programming

13 Marcus du Sautoy, *The Creativity Code: Art and Innovation in the Age of AI* (Cambridge, M.A.: Fourth Estate 2020).

14 Gaëtan Hadjeres, François Pachet, & Frank Nielsen, 'DeepBach: a Steerable Model for Bach Chorales Generation' in *Proceedings of the 34th International Conference on Machine Learning* (*2017*); Raymond Whorley, & Robin Laney, 'Generating Subjects for Pieces in the Style of Bach's Two-Part Inventions' in *Proceedings of the 2020 Joint Conference on AI Music Creativity* (2021); Alexander Fang, Alisa Liu, Prem Seetharaman, & Bryan Pardo *Bach or Mock? A Grading Function for Chorales in the Style of J. S. Bach* (2020), ArXiv:2006.13329 [Cs.SD].

15 Jason Goodyer, *How an artificial intelligence finished Beethoven's last symphony*, (Science Focus, 2021), https://www.sciencefocus.com/news/ai-beethovens-symphony/

16 Robert Langkjær-Bain, 'Five ways data is transforming music' in *Significance* 15:1 (2018), 20–23; AIVA, *AIVA - The AI composing emotional soundtrack music* (AIVA), https://www.aiva.ai/.

of the vast majority of symphony orchestras.[17] Why do we do this? In my view, it is because we believe that ideas from the past can have something to say in the present – something beyond merely being a benchmark by which to judge technical progress. We never try to merely recreate. Conversely, modern composers often use references to older music, or different genres of music, to make exciting and fascinating musical arguments (e.g. Michael Gordon's *Rewriting Beethoven's Seventh Symphony* or Sky Macklay's *Many Many Cadences*).

Silicon's first movement, *Mind,* is scored for a classical orchestra with a few additional instruments. There are no electronics: the relationship between orchestra and AI is realised solely through the live performers. It is the longest movement of *Silicon* at around fifteen minutes. Through it, I explore Future and Past in two specific ways:

By using an AI designed to imitate composers to push orchestral music into new places.

By creating a piece that might make sense to a theoretical, far-future AI mind.

The AI I chose to use for this piece is called MuseNet.[18] MuseNet, developed by OpenAI, produces MIDI data (i.e. sheet music) and has been trained on a very large dataset of music throughout history. This allows it to imitate a wide range of styles, including my own (though I did not use this functionality while composing *Silicon*). I have developed familiarity with its interface and quirks through using it in several projects over the last few years.

MuseNet works by providing the trained model with some music (a 'prompt') and giving it some basic instructions of how to continue writing the piece (a 'response'). For this movement, I instructed MuseNet to continue it in the style of Mozart. This suited the classical orchestra, and relates to the use of the sonata form in this movement, which will be discussed later. I would provide MuseNet with prompts in short score and sift through the many possible responses it generated.

17 Donne, *Women In Music, 2019-2020 - Donne, Women in Music Research* (Donne 2020), https://donne-uk.org/2019-2020/; Mark Gotham, 'Coherence in Concert Programming: A View from the U.K.', in *IRASM* 45:2 (2014), 293–309.

18 Christine Payne, 'MuseNet' in *OpenAI* (2019).

In contrast to an approach aiming to show how well the algorithm learned, where I might have chosen to keep responses I thought sounded like Mozart, I instead elected to retain the responses which I found unusual, uncanny, obsessive and – in a word – bizarre. I was fascinated by these responses, because it was clear that MuseNet understood them to be just as stylistically accurate as the responses that really did sound like Mozart. It felt like there was a strange, almost alien, methodology at work under the surface, which to me was far more interesting than a computer showing it can learn rules we already know concerning harmony, counterpoint and voice-leading. Perhaps the learning machine had discovered a deeper 'style' that underpins this music, which does not necessarily relate to the surface-level features of the music.

Selecting which to retain was a creative act in this work, in which I was rejecting the idea of 'best' and instead comparing each generation on their own merits according to my own subjective criteria. In general, I was drawn to responses that introduced a fundamental change in some parameter, for example substantially altering the harmonic rate, introducing completely novel melodic material, or repeating material in an unexpected way. After choosing a response, I might compose the next few bars, or sometimes simply give that response as the prompt for the next iteration of MuseNet to work with. In this way, the work goes down a unique rabbit-hole that I couldn't have planned on my own. This is one of the joys of working with AI – it can introduce an element of spontaneity to work that you can't account for, like a voice from another world. Analysing the responses can also be quite illuminating, and in these responses I was particularly taken by MuseNet's approach to repeated material, which is very polarised. Often material is stated once and never repeated or developed. Equally often, material will be obsessed over for an extremely long period. Sometimes it does repeat phrases in a way we would associate with classical music (such as composing a four-bar phrase and repeating it with an ending in the dominant rather than the tonic). This shows the AI did, or could, learn how repeats are used stylistically, but that it made the decision *not* to do this. I found these approaches to repeated material very inspiring and developed them throughout *Silicon*.

In my second area of exploration, I investigated what music of the future might sound like if an AI wrote music intended for other computers, not humans. Even when informed by real AI research, this remains very speculative and there are many areas this might affect. However, I was particularly taken by the way that AI algorithms consider time. The notion of general algorithmic time is one that already has strong parallels with music. Rohrhuber's argument that 'algorithmic methods suggest a break with the idea of time as an immediate grounding' is based upon the idea that an algorithm defines its own time through unfolding its pre-programmed computations, or steps, in a specific order.[19] Time 'begins' with the first step and 'ends' when the algorithm is complete. It does not necessarily matter how long, in actual time, these steps take. This perspective chimes with some orchestral music. We can imagine a sonata form or the four movements of a traditional symphony as a type of algorithmic time, where it is more informative to understand the relationship of the internal sections that unfold in a specific order (e.g. exposition – development – recapitulation, or sonata – rondo – finale), than it is to count how many seconds have passed in actual time. That is not to say that this is the only way of viewing musical time, which is a complex phenomenon – or even that, on its own, algorithmic time is a sufficient method of understanding musical time. It was to me, however, a new perspective on how algorithms, as a concept, can relate to classical music. *Silicon Mind* is in a warped sonata form, partly to explore this connection.

Musical time has more specific connections with AI algorithms. AI learns from audio data like WAV files, or from symbolic data like MIDI files. Whether audio or MIDI, for AI training purposes this data can be transformed into an image, such as a spectrograph or a MIDI roll.[20] Images do not exist in time, they are static. It's only when we tell the machine-learning algorithm to play that image from left to right that the dimension of time suddenly originates. But if a machine was creating music for itself, in a theoretical future where machines exist that enjoy listening to music for its own sake, musical time probably wouldn't need

19 Julian Rohrhuber, 'Algorithmic Music and the Philosophy of Time', in *The Oxford Handbook of Algorithmic Music*, ed. by A. McLean & R. Dean (Oxford: Oxford University Press 2018), pp. 17–40.

20 Carykh, *AI Evolves to Compose 3 Hours of Jazz!* [video], YouTube, 5 July 2017, https://youtu.be/nA3YOFUCn4U

to work in the way we experience it. The image-music could be enjoyed all at once, top-to-bottom, right-to-left or the traditional start-to-end.

To enact this in *Silicon Mind*, I created several axes of reflection across the piece. On either side of these axes, we hear the same music both forwards and backwards. This is not only retrograding rhythms and pitches, but also the timbre, decay and attack of the sound. If we imagine reading a spectrograph backwards (right-to-left) the entire sound is in reverse. Realising this with only the physical instruments of the orchestra presented an enjoyable challenge: reversing the sound of the vibraphone, for example, requires the percussionist to first bow the note to produce a sustained note, before striking and damping the note with a mallet. In this way we hear decay followed by attack followed by silence. Returning to the sonata form idea, I composed the two subjects as a contrasting pair: where one is heard forwards, the other is heard backwards, synchronising at the end of the movement.

MuseNet can be instructed to create any number of responses to a prompt, which will all be created simultaneously, and each will be different. This also fed into the approach to musical time during *Silicon Mind*. Several times while composing, I added one MuseNet response before rewinding back to the start of that phrase to use another, creating a sonification of constant progress through many iterations of the same task – the core tenet of machine-learning AI.

Fake and Real (*Silicon Body*)

One of the biggest concerns raised by AI concerns authenticity. In recent years we have become familiar with AI's capacity for creating believable fakes. This technology is used to automatically generate stories that resemble human-written news and by social media giants to encourage engagement, with the dissemination and promotion of fake news stories a known by-product.[21] It is now a regular occurrence to see AI algorithms used to create fake videos showing public figures in unfavourable light and it has also been used in movies to allow deceased actors to appear

21 Patrick Wang, Rafael Angarita, & Ilaria Renna, 'Is this the Era of Misinformation yet: Combining Social Bots and Fake News to Deceive the Masses', in *The Web Conference 2018 - Companion of the World Wide Web Conference, WWW 2018* (2018), 1557–1561.

in new releases (e.g. Peter Cushing and Carrie Fisher in *Rogue One: A Star Wars Story*) or to de-age live ones.[22] This kind of technology is often colloquially called deepfake, but the technical term for this field is Style Transfer.

We are now becoming used to questioning the provenance of believable-looking sources in a way that we were not even at the turn of this century. I believe this question of authenticity will be one that defines society over the next generation, and even if the identification of human-made or AI-made content were resolved. Real is not the same as authentic, and this is especially clear in the creative fields. An AI might generate *real* music in the style of Mozart (as discussed), but this music might not feel *authentic* to all listeners. Authenticity is a much more subjective question than truth or untruth. Here, the orchestra, and classical music more generally, can offer a perspective.

Classical musicians are familiar with questions of authenticity. Discussions and disagreements emerging around, for example, performing Bach on the modern piano,[23] using vibrato in eighteenth-century symphonies,[24] or casting singers of colour to sing operatic roles representing minority groups can be viewed, at least partially, as questions of authenticity.[25] Though this is not exclusive to this musical genre alone, it remains true that classical musicians are plugged in to whether music feels authentic, in addition to what it sounds or looks like on the surface. A music that deliberately grapples with fake and real is already in a good place to survive the onslaught of a deepfake society.

This leads to the questions that are at the foundation of the second orchestral movement, *Silicon Body.* What exactly is fake music? And does fake or inauthentic music become any more authentic when performed by an orchestra, by real people? Perhaps most importantly, I wanted

22 Johnny Botha, & Heloise Pieterse, 'Fake News and Deepfakes: A Dangerous Threat for 21st Century Information Security', in *ICCWS 2020 15th International Conference on Cyber Warfare and Security* ed. by B. K. Payne & H. Wu (2020), pp. 57–66; Alexi Sargeant, *The Undeath of Cinema* (The New Atlantis, 2017), https://www.thenewatlantis.com/publications/the-undeath-of-cinema

23 Aron Edidin, 'Playing Bach His Way: Historical Authenticity, Personal Authenticity, and the Performance of Classical Music', *Journal of Aesthetic Education* 32:4 (1998), 79.

24 Roger Norrington, 'The sound orchestras make', *Early Music* 32:1 (2004), 2–6.

25 Naomi André, Karen M. Bryan, & Eric Saylor in *Blackness in Opera* (Urbana: University of Illinois Press, 2012).

to hear what this deepfake technology actually sounds like. I wanted to embed an instrument that uses AI deepfake technology within the orchestra, to be played by an orchestral musician, as a kind of model for how orchestras might be constituted in the age of AI.

One research paper that particularly interested me showcasing deepfake technology is called 'Everybody Dance Now'.[26] It demonstrates taking a video of a dancer (Source), an image of a second person (Target), and the use of AI to make the Target appear to move like the Source. To do this, it strips the Source video down to a basic set of moving points and lines, abstractly representing the human body. With this distilled from the Source, the AI then rebuilds the video, this time with the Target fleshing out the skeletal nodes. I found it fascinating the way that computer vision 'sees' people fundamentally differently to how we see people, and also perhaps a little unnerving. An answer to what fake music might sound like lay, for me, in the relationship between the surface – the Target – and the hidden layers – the Source.

Silicon Body has a Source, a layer of music I composed that sits underneath the whole piece. It's a skeletal musical framework, made up of mathematical patterns of pitches and rhythms moving in cycles. This musical Source is not performed by a regular orchestral instrument, but instead by a digital instrument called DDSP (Differential Digital Signal Processing)[27] developed by Google Magenta. DDSP is a Style Transfer instrument that works in a similar way to the earlier dancer example, except that the Source and Target are audio-based rather than video-based. We can play any sound into DDSP and instruct it to transfer that sound's harmonic content into any other timbre using AI. This can be very convincing, but it can also be rather uncanny: a sound that is nearly right, but not quite. It is also possible to push the instrument outside of its intended comfort zone to create exciting new timbres through AI. Embedding new AI instruments within the ensemble was creatively interesting and engaging, and there is much more work to be done in this field.

26 Caroline Chan, Shiry Ginosar, Tinghui Zhou, & Alexei A. Efros, 'Everybody Dance Now' in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2019), 5933–5942.

27 Jesse Engel, Lamtharn Hantrakul, Chenjie Gu, & Adam Roberts, *DDSP: Differentiable Digital Signal Processing* (2020), https://doi.org/10.48550/arxiv.2001.04643

On top of this Source are superimposed three Target styles of music that are performed by the orchestral instruments. Inspired by 'Everybody Dance Now', these three styles are based on different types of dance music: big-band jazz, electronic and folk. Continuing the idea of reference to traditional orchestral music, this also makes *Silicon*'s second movement a kind of dance movement, to follow the first warped sonata form. While each of these dance styles sounds completely different on the surface, they are each controlled by the Source instrument. As the piece progresses, the three Targets are rotated faster and faster, until they reach a breaking point and the Source is revealed finally on its own. The music is meant to sound fun, uncanny and sinister, reflecting the many uses of deepfake technology.

System and Secret (*Silicon Soul*)

Both case-study algorithms discussed so far, MuseNet and DDSP, are partly interesting due to their imperfections. Both do not always do what might have originally been intended, and in not doing so they reveal interesting artefacts, processes and obsessions of AI. For many artists, these artefacts are one of the main reasons to use AI in the creative process, and the uncanny or strange is actively sought out.

Let us now imagine the inverse as a thought experiment, which will lead to the third duality: System and Secret. Our thought-experiment algorithm has no artefacts, and it can achieve whatever musical task we set it. It can analyse any amount of data, unrestricted by hardware limitations, and can produce new data (i.e. music) trivially quickly. It can produce sound indistinguishable from human musicians in any genre, historical period or ensemble. It can even produce entirely new music by combining existing music in novel ways or identifying gaps in its dataset that have never been exploited. It is, in a word, perfect. But is it music?

Across society there is a trend to understand any problem as inherently solvable by data. Here I draw upon the work of Federico Campagna, especially his book *Technic and Magic*,[28] but also more

28 Federico Campagna, *Technic and Magic* (London: Bloomsbury Academic, 2018), https://doi.org/10.5040/9781350044005

widely upon recent public policy decisions in the UK and abroad. The implication, in the field of music-related AI, is that with enough data and computer resources, an AI does not only produce something that *sounds* like music, but that it *is* music, that is, that there is nothing that constitutes music that exists outside of data. Given a sufficiently sophisticated system, creating meaningful music is trivial.

But would people accept this music, or do we require some kind of secret ingredient in order to feel a genuine connection with art? We do not know the answer to this question because AI has not yet reached the fluency of our thought experiment, but it is reasonable to imagine that it will some day. And if we take the view that there is more to music than computer data can communicate, what is that secret? Does it exist inherently within the music, or can this secret be imagined or imposed by the audience? Will AI research, in its dogged pursuit of a systematic understanding of the world, help us understand what the secret of music is?

This question is already under active consideration from a wide range of artists and scholars. Campagna argues that embracing a worldview he terms 'magic', informed by elements of spiritualism, mysticism and religion, can help alleviate the difficulties, both personal and social, inherent in a worldview reliant on data. Similarly, the authors of the *Atlas of Anomalous AI* explicitly state their aim to 're-mythologise AI in a way that reveals the roots of the technological project in spiritual practices, institutions and frameworks'.[29] Composer and academic George E. Lewis, when explaining why he is interested in teaching computers how to improvise music, describes a view of improvisation as 'something essential, fundamental to the human spirit', before going on to assert that attempting to teach computers to improvise 'can teach us how to live in a world marked by agency, indeterminacy, analysis of conditions, and the apparent ineffability of choice'.[30] The question of what the secret is, and how computers and data might help us understand it, has as many answers as answerers.

29 Ben Vickers, & K. Allado-McDowell, 'Introduction', in *Atlas of Anomalous AI*, ed. by B. Vickers & K. Allado-McDowell (London: Ignota Books 2021).

30 George E. Lewis, 'Why Do We Want Our Computers to Improvise?', in *The Oxford Handbook of Algorithmic Music*, ed. by A. McLean & R. Dean (Oxford: Oxford University Press 2018), pp. 123–130.

I set out to provide one response, if not an answer, to this difficult question by examining it through the lens of orchestral music. I wondered why audiences still go to see the orchestra today. As the COVID-19 pandemic has shown, it is perfectly possible to livestream performances to tune into from home, and there are even sample libraries – frequently used for video game, movie and TV soundtracks – that allow us to emulate the orchestral sound without needing any humans at all. What is its secret that compels people to physically come and watch humans make these sounds live?

For me personally, it is in understanding an orchestral performance not primarily as an act of creating sound, but rather as an act of community shared between musicians and audience. For the third movement of *Silicon, Silicon Soul*, I wanted to experiment with including AI inside such a framework. This means not just providing notes to be interpreted, or brought to life, by human musicians – like *Silicon Mind* – nor providing a fixed, uncanny perspective against the human members of the orchestra – like *Silicon Body* – but understanding AI as an integral part of the orchestra's communal act.

For this movement, I used an AI called PRiSM-SampleRNN. PRiSM-SampleRNN is an audio-based AI that learns to create new sounds from a dataset of existing sound. It produces the raw audio – both 'notes' and timbre, making it different to MuseNet and DDSP which both require human performers somewhere in the realisation. The version of PRiSM-SampleRNN I used was released by PRiSM in 2020.[31] One method I explored for integrating AI into the community of the orchestra was to make the AI personal to *that* orchestra. I trained PRiSM-SampleRNN on recordings of the BBC Philharmonic exclusively. This means that whatever it learned about music, it has learned from analysing that orchestra alone. In performance, the results of this training will be heard alongside the BBC Philharmonic – like an apprentice performing alongside a master. In this instance AI is used as a tool to increase the personalisation and site-specific nature of a piece, rather than as a tool to make general rules about music. It is in service of defining what the nature of *this* ensemble is, and the audience is challenged to make their

31 Christopher Melen, 'PRiSM SampleRNN', *RNCM PRiSM* (2019), https://www.rncm.ac.uk/research/research-centres-rncm/prism/prism-collaborations/prism-samplernn/

own decisions about the differences in sound between the physical orchestra in front of them and its AI doppelganger.

There was also literal integration of the AI sound with the orchestral sounds. Unlike the use of DDSP in *Silicon Body*, where the distinction between orchestra and AI is visually and audibly apparent, I wanted the overall sound of this movement to be an indivisible unit where AI and human sounds seamlessly merged. To achieve this, I requested that the PRiSM-SampleRNN audio be dispersed amongst many audio monitors, instead of the more standard stereo pair, and that these monitors be hidden amongst and underneath the orchestral players. To complement this, the compositional dialogue between orchestra and AI was considered on a spectrum throughout the movement, where at one end the two are totally overlapping in musical material, while at the other they make separate musical arguments. This brings to mind questions of human-computer interactivity and interdependency, an area I look forward to other artists and myself exploring with even more sophisticated tools in the future as technology develops.

Conclusions and the Future

Having set out how I have recently used AI in my work, I will briefly discuss the overall effect on my compositional craft and outlook that using this technology has had. These are beyond the general benefits, which should not be understated, of learning to code and of incorporating complex electronics into my music.

My music has always been driven by an interest in structure and musical time. Working with AI has driven this interest in new directions. Most prominent is the role of spontaneity within the compositional process. When working with algorithms such as MuseNet or PRiSM-SampleRNN, it is not possible to plan exactly the form of a work in advance, because it is not possible to control those algorithms' generations. I have therefore adopted a balance between spontaneity and structure in my work, which has (in my opinion) been for the better. Even on a small, moment-by-moment scale, there is very little control over specific details (for example, the timbre of a PRiSM-SampleRNN generation). Rather than forcing this AI-generated material into a pre-set idea, I have found myself listening to what this material *is*, and how I might either heighten or obscure that essence. In turn, this has informed my approach to working with human performers too. Working with AI has made me consider the role of the human performer a great deal. AI does not (currently) generate particularly idiomatic parts for performers, which has caused me to consider the roles of playable and awkward music, particularly regarding the idea of authenticity in classical music.

Perhaps the greatest effect AI has had on my compositional process is in its treatment of the basic elements of music. When I sit down and analyse its approach to, say, repeated material (as discussed), tuning, rhythm, time or harmony, I find that it is subtly yet fundamentally different to approaches I am familiar with. As a composer, it is easy to get caught up in fine details of musical elements. These remain interesting, but working with AI has opened my eyes to the benefits of searching out a paradigm shift. Much of my recent non-AI compositional work, for example, has been focussed on exploring new forms of tuning, which I became interested in through analysing AI-material.

If I have described a possible future relationship between the orchestra and AI, then what are the concrete steps that might lead us along this path? Some areas for further development, based on this chapter's discussion, will be highlighted. In terms of making AI algorithms more generally useful for musicians themselves, the most obvious point is that of user interface. Many AI algorithms are interfaced with through Python, a coding language that few musicians are comfortable with. Creating a graphical interface that requires no coding, as exemplified by Google Magenta, will be the most important step to encourage more musicians to experiment with this technology. As further AI instruments are developed, the orchestra and other performing ensembles might take an active role in facilitating experimentation with these new technologies. While I used DDSP for its timbral and metaphorical advantages in *Silicon Body*, I think that a major benefit that AI instruments will bring to classical music is accessibility. AI has been shown to be proficient at learning its users' gestures and actions (e.g. your phone recognises your face and no-one else's), and this opens up the real possibility of utilising AI to create adaptive instruments for people whose physical impairments might preclude the practice of traditional instruments.

The composition of *Silicon Mind* highlighted the fact that current AI generates music from start-to-finish. A more useful compositional tool might be an AI that can generate music back-to-front, or perhaps bridge two materials that a composer provides. This has been achieved on a small scale through algorithms such as CocoNet,[32] where the user provides a melody which is harmonised by AI in the style of Bach. Many composers do not write their pieces from start to finish, so future AI intended to be helpful for composers should take this into account.

It is now the winter of 2022, and the work on *Silicon* for orchestral and artificial intelligence is finished. This chapter has used it as a backdrop and anchor to explore three relatively abstract dualities shared between AI and classical music, and to transform these dualities into something audible. It also showed three different ways AI can be employed as a technology within the orchestral texture itself, playing the role of composer (MuseNet), instrument (DDSP) and performer (PRiSM-SampleRNN). I hope that through experimentation like this, the orchestra can remain at the forefront of instrumental and compositional exploration, as it did in the days of Berlioz and Wagner.

Three dualities, and three technological use-cases of AI within the orchestra, barely scratch the surface of the potential relationship here, and my biggest hope is that this chapter has caused readers to consider other ways AI might be a useful concern for the orchestra and classical music more generally. My experience and interests using AI within the orchestra are biased towards the way that I operate as a composer, and so I hope to see other composers' substantially varied takes on this technology in the coming years. This chapter has been quite speculative, and I am sure that while some concerns described here might end up as dead ends, there will also be unforeseen questions emerging from this technology.

My own work has been supported by several institutions, which has placed me in the fortunate position of being able to seek out the positive elements of AI that can help me attain my artistic goals. As stated at the beginning of this chapter, however, this technology can be used in many ways, not all of them positive. AI has profound implications for the future of labour. In classical music, this question could be posed as: who might be replaced by AI? Already film composers utilise high-quality

32 Cheng-Zhi Anna Huang, Tim Cooijmans, Adam Roberts, Aaron Courville, & Douglas Eck, 'Counterpoint by Convolution' in *Proceedings of the 18th International Society for Music Information Retrieval Conference, ISMIR 2017* (2019), 211–218.

sample libraries which have, by and large, replaced live musicians for much of their work. Will composers be replaced by AI algorithms next? What kinds of music will audiences accept as being written by AI, and which will they demand should be written by humans? Will it be mandatory to announce whether music is AI- or human-composed, to avoid hoodwinking the public? These are not yet questions that have been tested, but may well feature in future discourse.

There is also the question of access. Currently, AI is prohibitively difficult to use both in terms of expertise and resources. Until more AI tools are open-source, or at least made affordable to use, it is difficult to see this technology becoming widespread. This would preclude its use in most educational contexts, an area that I believe AI has much to offer (i.e. assisting a student composer in mixing their music, or recommending several different paths for the music to take from a given moment).

Audiences and programmers of classical music might find both positives and negatives stemming from this technology. On the one hand, it might provide novel and exciting programmes (for example, the 2019 Cheltenham Science Festival used AI to suggest the titles of talks, which were then programmed). On the other, considering where AI-generated music might fit into the classical music programming model might require a radical rethink of how seasons and individual events are put together. This is not to mention the power of AI as a marketing tool: institutions that are slow to incorporate AI may be left behind compared to rivals using advanced AI to attract audiences.

Ultimately, I am sure that AI is here to stay, like the printing press, computer and internet before it. It already touches upon much of our society, whether we know it or not, and its expansion is only likely to continue in the future. It is partly through the work of artists that we might further understand its place in our world.

References

AIVA, *AIVA—The AI composing emotional soundtrack music* (AIVA), https://www.aiva.ai/

André, N. A., Bryan, K. M., & Saylor, E., *Blackness in Opera* (Urbana: University of Illinois Press, 2012).

Botha, Johnny, & Pieterse, Heloise, 'Fake News and Deepfakes: A Dangerous Threat for 21st Century Information Security', in *Proceedings of the 15th International Conference on Cyber Warfare and Security, ICCWS 2020*, ed. by B. K. Payne & H. Wu (2020), pp. 57–66.

Campagna, Federico, *Technic and Magic* (London: Bloomsbury Academic, 2018), https://doi.org/10.5040/9781350044005

Carykh, *AI Evolves to Compose 3 Hours of Jazz!*, YouTube, 5 July 2017, https://youtu.be/nA3YOFUCn4U" https://youtu.be/nA3YOFUCn4U

Chan, C., Ginosar, S., Zhou, T., & Efros, A. A., 'Everybody Dance Now', *Proceedings of the IEEE/CVF International Conference on Computer Vision* (*ICCV*) (2019), 5933–5942.

Cowls, J., Tsamados, A., Taddeo, M., & Floridi, L., 'The AI gambit: leveraging artificial intelligence to combat climate change—opportunities, challenges, and recommendations', *AI and Society* 1 (2021), 1–25, https://doi.org/10.1007/s00146-021-01294-x

Cox, Geoff, & Riis, Morton, '(Micro)Politics of Algorithmic Music', in *The Oxford Handbook of Algorithmic Music*, ed. by A. McLean & R. Dean (Oxford: Oxford University Press 2018), pp. 603–626, https://doi.org/10.1093/oxfordhb/9780190226992.001.0001

Dastin, Jeffrey, 'Amazon Scraps Secret AI Recruiting Tool that Showed Bias against Women', in *Ethics of Data and Analytics*, ed. by K. Martin (Boca Raton: Auerbach Publications, 2022), pp. 296–299.

Donne, Women In Music, 2019-2020—Donne, Women in Music Research (Donne 2020), https://donne-uk.org/2019-2020/

du Sautoy, Marcus, *The Creativity Code: Art and Innovation in the Age of AI* (Cambridge, M.A.: Fourth Estate, 2020).

Edidin, Aron, 'Playing Bach His Way: Historical Authenticity, Personal Authenticity, and the Performance of Classical Music', *Journal of Aesthetic Education* 32:4 (1998), 79, https://doi.org/10.2307/3333387

Engel, J., Hantrakul, L., Gu, C., & Roberts, A., *DDSP: Differentiable Digital Signal Processing* (2020), https://doi.org/10.48550/arxiv.2001.04643

Fang, A., Liu, A., Seetharaman, P., & Pardo, B., *Bach or Mock? A Grading Function for Chorales in the Style of J. S. Bach* (2020), ArXiv:2006.13329 [Cs.SD].

Fiebrink, Rebecca, & Cook, Perry R., 'The Wekinator: a system for real-time, interactive machine learning in music', *The 11th International Society for Music Information Retrieval Conference* (2010), http://code.google.com/p/wekinator/

Gomez-Uribe, Carlos A., & Hunt, Neil, 'The Netflix Recommender System', *ACM Transactions on Management Information Systems* (*TMIS*) 6:4 (2015).

Goodyer, Jason, 'How an artificial intelligence finished Beethoven's last symphony', *Science Focus*, 14 October 2021, https://www.sciencefocus.com/news/ai-beethovens-symphony/

Gotham, Mark, 'Coherence in Concert Programming: A View from the U.K.', *IRASM* 45:2 (2014), 293–309.

Hadjeres, G. Pachet, F., Nielsen, F., 'DeepBach: a Steerable Model for Bach Chorales Generation', *Proceedings of the 34th International Conference on Machine Learning* (2017). https://proceedings.mlr.press/v70/kim17a.html.

Hao, Karen, 'He got Facebook hooked on AI. Now he can't fix its misinformation addiction', *MIT Technology Review*, 11 March 2021, https://www.technologyreview.com/2021/03/11/1020600/facebook-responsible-ai-misinformation/

Huang, C. Z. A., Cooijmans, T., Roberts, A., Courville, A., & Eck, D., 'Counterpoint by Convolution', *Proceedings of the 18th International Society for Music Information Retrieval Conference, ISMIR 2017* (2019), 211–218, https://doi.org/10.48550/arXiv.1903.07227

Karyuatry, Laksnoria, 'Grammarly as a Tool to Improve Students' Writing Quality: Free Online-Proofreader across the Boundaries', *JSSH* (*Jurnal Sains Sosial Dan Humaniora*) 2:1 (2018), 83–89, http://dx.doi.org/10.30595/jssh.v2i1.2297

Langkjær-Bain, Robert, 'Five ways data is transforming music', *Significance* 15:1 (2018), 20–23, https://doi.org/10.1111/j.1740-9713.2018.01106.x

Lau, Johann, 'Google Maps 101: How AI helps predict traffic and determine routes', *Google: The Keyword*, 3 September 2020, https://blog.google/products/maps/

google-maps-101-how-ai-helps-predict-traffic-and-determine-routes/

Lee, L.-H., Braud, T., Zhou, P., Wang, L., Xu, D., Lin, Z., Kumar, A., Bermejo, C., & Hui, P., *All One Needs to Know about Metaverse: A Complete Survey on Technological Singularity, Virtual Ecosystem, and Research Agenda*, arXiv:2110.05352 [cs.CY] (2021), https://doi.org/10.48550/arxiv.2110.05352

Leslie, David, 'Understanding Bias in Facial Recognition Technologies', Zenodo (2020), https://doi.org/10.5281/zenodo.4050457

Lewis, George E., 'Why Do We Want Our Computers to Improvise?', in *The Oxford Handbook of Algorithmic Music*, ed. by A. McLean & R. Dean (Oxford: Oxford University Press, 2018), pp. 123–130.

McLean, Alex, & Dean, Roger T., *The Oxford Handbook of Algorithmic Music* (Oxford: Oxford University Press, 2018).

Melen, Christopher, 'PRiSM SampleRNN', *RNCM PRiSM* (2019), https://www.rncm.ac.uk/research/research-centres-rncm/prism/prism-collaborations/prism-samplernn/

Mogaji, Emmanuel, Olaleye, Sunday, & Ukpabi, Dandison, 'Using AI to Personalise Emotionally Appealing Advertisement', in *Digital and Social Media Marketing*, ed. by N. P. Rana, E. L. Slade, G. P. Sahu, H. Kizgin, N. Singh, B. Dey, A. Gutierrez, & Y. K. Dwivedi (Cham: Springer, 2020), pp. 137-150, https://doi.org/10.1007/978-3-030-24374-6_10

Morton, Timothy, *Hyperobjects: Philosophy and Ecology after the End of the World* (Minneapolis: University of Minnesota Press, 2013).

Norrington, Roger, 'The sound orchestras make', *Early Music* 32:1(2004), 2–6, https://doi.org/10.1093/earlyj/32.1.2

Payne, Christine, 'MuseNet', *OpenAI* (2019), https://openai.com/research/musenet

Rohrhuber, Julian, 'Algorithmic Music and the Philosophy of Time', in *The Oxford Handbook of Algorithmic Music*, ed. by A. McLean & R. Dean (Oxford University Press 2018), pp. 17–40, https://doi.org/10.1093/oxfordhb/9780190226992.001.0001

Sargeant, Alexi, *The Undeath of Cinema* (The New Atlantis, 2017), https://www.thenewatlantis.com/publications/the-undeath-of-cinema

Sejnowski, Terry, *The Deep Learning Revolution* (Cambridge, M.A.: MIT Press, 2018).

Tatalovic, Mico, 'AI writing bots are about to revolutionise science journalism: we must shape how this is done', *Journal of Science*

Communication 17:1 (2018), https://doi.org/10.22323/2.17010501

Vickers, Ben, & Allado-McDowell, K., 'Introduction', in *Atlas of Anomalous AI*, ed. by B. Vickers & K. Allado-McDowell (London: Ignota Books, 2021), pp. 9-29.

Wang, P., Angarita, R., & Renna, I., 'Is this the Era of Misinformation yet: Combining Social Bots and Fake News to Deceive the Masses', *The Web Conference 2018—Companion of the World Wide Web Conference, WWW 2018* (2018), 1557–1561, http://dx.doi.org/10.1145/3184558.3191610

Whorley, Raymond, & Laney, Robin, 'Generating Subjects for Pieces in the Style of Bach's Two-Part Inventions', *Proceedings of the 2020 Joint Conference on AI Music Creativity*, KTH Royal Institute of Technology2021), https://doi.org/10.30746/978-91-519-5560-5.