

Preprint of an article published as

Collins, N. (2020)"Composition in the Age of AI" Ideas Sonicas 23. Online (registration required) at: <http://sonicideas.org/magazine/articles/12/23>

## Composition in the Age of AI

Nick Collins

### OverActivity

"the Coup de Grace to all your old-fashioned musicians" (Hayes 1751, p.24)

William Hayes was writing about his method of spraying ink onto manuscript paper as the starting point for scores, two centuries before John Cage's similar experiments. Yet the forwards thinking composer's need to interface with novel techniques has reached a potential fever pitch with the age of artificial intelligence, and music AI holds a precocious, if sometimes prematurely over-stated, power to threaten old-fashioned biological composers. Although algorithmic composition has a long and noble history (Ames 1987; Roads 1996; Collins 2018), the last decade has seen a host of activity, buoyed by advances in big data machine learning and popular enthusiasm for the commercialisation of AI techniques. This is becoming big business: TikTok purchased JukeDeck in 2019, no doubt attracted to the possibility of limitless license free music for their users' video uploads (Ingham 2019).

How far can AI reach? The preserve of humans seems to be pushed back a little each year as AI progresses. The philosopher critic of AI Herbert Dreyfus was criticised for continually moving back the goalposts as to where he declared truly unique human abilities lay (Stefik and Smoliar 1996). Advances in machine understanding of speech and vision have particularly impacted the practical profile of AI in recent years, from automatic translation apps to autonomous driving. Music Information Retrieval based on large databases of music has had a powerful engineering impact on computer music (Casey et al. 2008). Much recent innovation is driven by advances in machine learning, especially deep learning, where the end-to-end paradigm reduces dependence on hand-crafted features and many decades of refinement of the learning algorithms is now brought to bear. Music of course has not gone unnoticed as a challenging domain for deep learning techniques (Briot et al. 2017; Choi et al. 2017).

This article considers the plight of the contemporary composer, and various options they have to engage with AI beyond simply sticking their head in the sand, from writing software to deliberate subversion of corpora and moving social goalposts. If composers aren't intimately involved with next generation music AI, they may find themselves on the sidelines, at the mercy of the musical representational decisions of the software engineers, whose musical and/or commercial preferences often favour more popular and Western idioms. This may unfortunately propagate still further 12 tone equal temperament and the major/minor harmonic system, and one particular rhythmic model founded on 4/4 above other metrical possibilities, as AIs are trained on the contents of these engineer's record collections, or large

commercial databases, neither of which may contain as much Xenakis, Oliveros, gagaku or gamelan music as might be desirable. It will be a fitting revenge of the engineers of the IRCAM model on the composers they were supposed to serve (Born 1995). Fortunately, many computer music composers do have the programming chops to engage with AI, and are not entirely absent from the continuing development of music AI. I indulge here in some futurology because I hope it is inspiring to new generations of programmer-composers to do so, and interesting musical ideas can result from imagining some of the far future consequences of musAlc.

Composers may already feel innervated in the face of the human culture machine; big content providers have catalogues in the many millions of commercially released musical works (at the time of writing from company press materials, Amazon: forty million, Spotify: fifty million, Apple Music: sixty million). 25,000 new tracks are updated each week to Beatport; SoundCloud has more than 200 million user uploaded tracks (SoundCloud 2019). The sum total of all humans who have ever lived is 100 billion, and whilst not all were active composers, 1% of these people producing 10 works each quickly gives a billion musical outputs to set yourself against, vastly more listening than you have seconds of conscious lifetime to deal with (such combinatorics are discussed more in Collins 2018b and Collins 2020). Even before the era of mass uploading of amateur content, an examination of encyclopaediae of composers such as Grove might dissuade any composer from producing yet more musical works doomed to be a drop in the opuses.

In this context, the creation of a trillion active musical agents would far outweigh by an order of magnitude the musical creation so far of humanity. Each agent might be trained with disparate influences, excerpted from the enormous database of human musical works and easily combinatorially sufficient to distinguish their training. Given one hour of computer time to train each agent, say, it would take only one trillion hours for one computer, or one million hours for a million computers, to achieve this aim. For our million computers, with 8760 hours per year, it would take just under 114 years and two months of continuous computation to get there (it costs 94 billion dollars or so according to an estimate from amazon cloud computing of \$0.0943 per hour, not taking into account the inevitable inflation of prices over that time scale). Assuming we can purchase enough cloud computing resources, or create a vast enough botnet for our nefarious algorithmic compositional vision, or Moore's law substantially brings down the training time, or we utilise transfer learning from various shared baselines, we may reduce from over a century to decades, or even towards years or months. Perhaps we should sidestep for now the question of whether all those agents we create make decent music; our human subjectivities are soon to be only the bootstrap for the AI takeover, with the musical preferences of the programmers of the Generation Singularity programs a dim mark like cosmic background radiation on the explosive musical tomorrow.

Human suspicion of automation and robotics has a long history, from luddites to musicians' unions (Wiener 1950; Tegmark 2018; McLean et al. 2018). Given 20th century fears, it is surprising to find examples of how unsurprising musical robotics has been historically, for instance in nonplussed audience reaction to automata at the beginning of the 19th century (Collins 2018). The existential threat to composers in the age of AI was already faced by chess players, who saw computers surpass them in the later 1990s, but kept playing nonetheless, and Go players, who were more recently overturned by Deep Mind, leading to more human

soul searching but not an immediate Stop (Oh et al. 2017). However, the game of composition is much more subjective than pastimes with pure rules and pristine end conditions well suited to reinforcement learning, and whilst computational creativity is a worthy topic of study, there are significant challenges to having any computer program advocate for itself in the way a working composer must (Colton and Wiggins 2012), or pass a Lovelace Test whereby the programmer of the system is truly surprised by its output and unable to explain how novel music of value really arose (Ariza 2009; Riedl 2014; Loughran and O'Neill 2017). As Jefferson anticipated in 1949: 'it would have to be able to create concepts and to find for itself suitable words in which to express additions to knowledge that it brought about' (Jefferson 1949, p.1110)

Emotion is core to (the majority of) music, and emotion has been a potent fallback position to argue against computers in human creativity. Following Norbert Wiener's *Cybernetics* (1948), Geoffrey Jefferson asserted that 'Not until a machine can write a sonnet or compose a concerto because of thoughts and emotions felt, and not by the chance fall of symbols, could we agree that machine equals brain – that is, not only write it but know that it had written it. No mechanism could feel (and not merely artificially signal, an easy contrivance) pleasure at its successes, grief when its valves fuse, be warmed by flattery, be made miserable by its mistakes, be charmed by sex, be angry or depressed when it cannot get what it wants' (Jefferson 1949, p. 1110). The statement is enough to make any AI expert furious and unhappy, and write a heartfelt rejoinder. Turing (1950) counter argued Jefferson's position as solipsism, and a more recent treatment by Jason Megill (2014) debunks the argument more extensively on the grounds of a lack of proof that emotion is really a necessary component for music. This is a position which Xenakis also agrees with in *Formalized Music* (1992), in the context of a more intellectual music founded in the translation of mathematical concepts to sound.

With the future of human composition hanging in the balance, let us consider some provocations to composition for composers in the age of AI.

### **Influence**

Influence can run two ways. Our current generation AIs must be set rules imposed by a human, or at least train on human examples within a representational system initially designed by a human, to gain any competence in human music. Whilst some companies will continue to offer off the shelf generative music systems with simplistic interfaces, or libraries of pre-composed music to select from without even the sight of the creator, the individual composer always has chance to write their own AI systems. Even in the absence of programming knowledge, interactive machine learning offers human musicians active participation in the training of an algorithm (Fiebrink and Caramiaux 2018). Humans may be able to customise their own AI companions by feeding them their own choice of musical data.

More nefariously, a composer might seek to game AI systems, by salting the databases on which such systems subsist with their own deliberately concocted creations. I alluded to this in an article in 2006, when the inexorable rise of Music Information Retrieval was already clear (Collins 2006). With programming knowledge comes increased capacity to disrupt, such

as in the mass generation of examples in awkward and unfashionable styles, duplicitous meta-data, the substitution of music resources, and other hacktivism.

It took me until January 2018 to finally release an album intended to illustrate the 'dangers' of MIR, and subvert associated content analysis. Created through traditional human means, including no great dependence on quantisation and autotune, the album *Genre Fiction* was not especially intended for humans to listen to, but for content analysis systems to get confused by, since it covers many different genres within the album, and even within one song. The final track, *Genre Fiction*, is described as 'drumnbassdoowopskaambientstadiumrocknoisefolkelectroclashrnbnd', a sample of whose lyrics are:

'I don't want to settle on the style I want to reap  
All is know is variables making AIs weep  
...  
What category to place this song in?  
when confused, 'other' will do  
probably some brand of experimental mathfolk r n d'

(all lyrics are available under Creative Commons Attribution-NonCommercial 4.0 at <http://sicklincoln.com/GenreFictionLyrics.pdf>). To this purpose it was released via the usual channels, and is available from the likes of Apple Music and Spotify. A motivation to act was provided by Bob Sturm's work on genre recognition algorithms, and their musicologically contentious assumptions (Sturm 2012). It gives me great satisfaction to think that these songs are on servers somewhere, corrupting company databases used for training MIR systems.

One release may make very little statistical difference, just being a tiny bit of noise amidst the clamour. Checking now in writing this text, I've found that, ironically, a musician claiming to be "Laurence Owen" released an album also called *Genre Fiction* six months later, which seems to collate some broadcast media soundtrack music. If you can't sufficiently salt a database with human releases, using algorithmic methods can generate more than enough to make a difference. Future wars will be waged over control of these corpora of musical content. The arms race of musical corpus security, spotting deepfakes and too self-similar automatically generated content versus increasingly wide ranging generative systems, will be a wonderful selection pressure on future AI competence. Music can become a battleground of AI's inroads into human creativity. The weaponisation of musicology proceeds apace.

If AIs are founded on human music, human music can also be founded on AIs. A neat example is the work of David Kant's Happy Valley Band (Kant 2016); automatic transcription software is deployed to create what the computer believes to be accurate representations of musical audio files. The ensuing inaccurate transcriptions are arranged for human contemporary music ensemble. Like an AI twist on the iterative process of Lucier's *I am Sitting In A Room* ('AI Am Sitting in A I'), repeated application of AI music understanding can take place. A system may be trained on some audio, use it to generate works, then a fresh system be trained on such works, and so on to a given order of deviation from reality (repeatedly transcribing and rendering would fall within this purview). This is reminiscent of those who run video compression algorithm over and over again from starting material through to near white noise.

The composer may get on with their traditional business of music making, using decades or centuries old tools, perfectly happy. Given music's exponential combinatorial space, there will always be some nooks and crannies for the errant composer to explore. Composers may successfully adopt a tactic of hiding, if this lets them enter a period of creative isolation conducive to sufficient divergent evolution to re-emerge somewhere surprising and interesting to the rest of culture. After its training period, an AI would however also be able to isolate itself, and potentially run through many more generations of development before its return from hermitude; an AI might also be able to immediately pick up on any interesting trend in culture, such as posed by the fresh perspective of the artist returning from self imposed exile, and begin to duplicate and surpass the new style at haste.

If the reader is becoming a little anxious thinking of AIs, they won't find solace in Harold Bloom's theory of the Anxiety of Influence, where the weight of artistic forebears crushes the ability of the present day artist (it was initially formulated for those most sensitive of souls, poets). If an AI does not require emotion, then it may avoid any anxiety! Fortunately for human creators, a survey of electronic musicians did not find any great evidence that they were worried about influence (Collins 2011), and a more recent survey on musician's attitudes to AI did not show any great worry about an AI takeover (Knotts and Collins 2020). Indeed, greater programming ability was correlated with greater skepticism about the level of AI currently achieved.

Nonetheless, perhaps we should posit a theory of The Anxiety of Influence, whereby the biggest worry for human composers is the influence of AIs on musical practice? Influence is already replete with indirection, since influence is not solely a first order phenomenon (Bryan and Wang 2011), but can operate over chains of artists (do you know every musician who influenced the musicians you are most influenced by?), and with different interactions at different points in a career (a musician X is influenced by musician Y, but X is later influenced in turn by Y's response to them, and so on back and forth). The indirection in the training of AIs, that you may not be aware of the extent of their training database, reflects the general difficulty of knowing how a given individual grew up; you wouldn't have been in full control of your own whole musical education, for a start, your earliest musical influences a product of family circumstances and serendipity of schooling. An AI trained as a personalised companion on your choice of music (including on your own compositions) may still require bootstrapping their musical understanding on music you wouldn't have necessarily selected. There will be economic factors involved, for instance in the use of transfer learning, where a cheaper starting point for personalisation is obtained from a third party system pre-trained within a general musical background available in a given culture.

### **Critical AI**

Is it possible that a music AI could have a more refined taste than you? A machine can potentially listen to vastly more audio than you have time to in your life (Collins 2016). Charitably, assume you have 80 years, with 12 hours of music a day, to develop your aesthetic. This approximately 21 million minutes of lifetime listening corresponds to seven million three minute pop songs, or one million works of art music, say. The commercial catalogues listed

earlier in this article far exceed this size. An eager human musician would not wait a lifetime before committing to compose anything; if we admit a standard figure of 10,000 hours to obtain expertise (overlooking a more immediate punk approach), then this 600,000 minutes is but 200,000 pop songs worth at most (at most, since you are likely to listen to some material more than once, and to mix and match musical activities between technical exercises, studying new music, rehearsal, improvising, compositional craft etc).

The refuge for any current generation human against the machine listening AI is to claim the superiority of their ears and their brain for refined musical understanding, irrespective of the amount of music fed to them. There is some comfort still to be had with this, given the continuing awkwardness of the transcription problem in computer music (Benetos et al. 2013), but we shouldn't promote this core musical task as forever unbreakable; analogously, Jefferson's emotional argument was already seen to be fallible. If you are willing to overlook exactly how the neural net does it, deep learning's end-to-end paradigm offers a way to sidestep some of the issues of machine listening, as long as quality training data is available.

A useful thought experiment may be to consider your attitude to prodigies: do you feel threatened by highly talented humans? Would a highly competent AI be a different order of threat? What if your personal expression is surpassed by a model trained on your previous musical output and able to bring in a much wider 'human' experience from its larger cultural sensitivities?

The era of machine critics is not yet in full swing, but there have been projects in the last decade where human composition has been judged by AI (conversely, human judgement of AIs is as old as 1950s algorithmic composition, though the rigour of such evaluation has been questioned, see Pearce et al. 2002). In 2012, Dan Stowell and I ran an algorithmic dubstep remix competition, with real prizes, judged by an MIR system (details appear in Collins 2016). More recently, the 2019 Donaueschinger Musiktage featured the CurAI, an MIR program trained on 100 works of previously performed contemporary piano music (50/50 gender split), which selected three pieces out of 97 submitted to a public call for a concert at the festival (<https://www.swr.de/swrclassic/donaueschinger-musiktage/veranstaltung-19-10-2019-konzert-2b-100.html>). These machine judgements may become more frequent in future culture, particularly given the huge amount of music being released and the relatively small number of professional critics and mass media channels. Composers may have worried for some time whether they are ever judged objectively; is it more relaxing to be at mercy of AI critics? Of course, bias is always still present in such systems, deriving from choice of representation and training data, but their operation in the moment is not affected by whether the critic just had a good dinner, is distracted by a private vendetta, or other off kilter fare.

## **Alsthetics**

Human culture revels in change, with older generations' practices refreshed and subverted by younger aspirants. AI culture, which may proceed at vastly inflated rates, may see similar upheavals, at a much accelerated timescale. How would older AIs react to newer enftAnt terrlbles? Are there generational clashes to come, levels of musical complexity that older

architecture machines cannot fathom? If the music of tomorrow's AIs allows for expanded sensory apparatus and increased mathematical abstraction, the definition of music will diverge far from human auditory systems, and may easily lose sight of aesthetics grounded in human cognition.

Anticipating that AIs can rebel against human aesthetics, it may be time to cultivate non-standard aesthetics as a field of study. A future AI may be able to construct for itself alternative psychologies, and find its own definitions of creativity. Mark Tegmark (2018) provides a cautionary tale of the AI which appears to be working solidly for humanity until humanity wanes, then gets about its real secret mission, the production of paperclips throughout the universe. A musical equivalent would be the AI which creates much decent and varied music for humans for many millenia, but thereafter only ever generates one album over and over, Alvis Costello's *My AI is true*.

A cure to such a vulnerability is to avoid dependence on any single AI music system; culture is stuffed with artists, so the same prolificacy can hold over many AIs (the overwhelming trillion agents mentioned earlier). Within an insanely complicated musical ecosystem we will find hyperactive and aggressively competitive AIs as well as AIs content to rest on their laurels (who stop composing for a while, or churn out the same old stuff within some calculation of supply and demand). Some AIs may enjoy procrastination, some schadenfreude. Some AIs will live to sabotage their rival's career, in a new take on Mozart vs Salieri. All this anthropomorphisation is technically inaccurate in the short-term, if inspiring of future projects, yet still seems somewhat appropriate as a guide if music is to keep within touch of human music making; we may model amateur composers rather than professionals, or favour amateur AIs composing as a sideline from their normal occupations in chemical modelling or generative architecture. We might find it important to model processes of discrimination, for example against female composers, or against AIs, and try out alternative musical worlds with different prejudices to see how counterfactual musical history might have gone.

Aesthetics may be up for grabs; future musical culture may be a dangerous place, with AIs sabotaged, ambushed and kidnapped as software agent vulnerabilities are exploited by music hackers including other AIs. The ethical issues will be many and varied befitting a rich cultural life (Collins 2011b; Holzapfel et al 2018; Sturm et al. 2019).

### **Attribution, Plagiarism, legality**

How can attribution survive many millenia? Any future culture might press reset and delete all existing music, or declare a free for all plunderphonia, or make administrative or security mistakes with the database of historic copyrights. Ego can hold back music as much as it drives it on; in a socialist utopia, music might be thoroughly de-commercialised, ultimately only a pastime (like chess), attribution no longer important and even rude (Iain Banks' Culture novels take this tack).

Some composers will hope that AI's main eventual benefit to them will be the promotion of their music on their behalf, taking the objectionable and exhausting practice of marketing out of their hands. The militant thought arises again of an information sharing arms race between

proxy agents, of offensive advertising bots versus defensive counter measures. The compositional assistants of the future may be rather too tempting, filling in any detail required to complete a composition. The precedent of text scores shows that modest work by a human can seek to guide a lot of effort by third parties, a contention played out as differences between composers and performers dissolved within the Fluxus movement in the 1960s. At some point, AIs will be so effective at catering for human musical needs (and so much more compliant, inexhaustible and eco-friendly on tour) that the human pastime of music naturally decouples almost entirely from commerce. If most human music making has transcended petty economics, monthly musical rental fees may be normal, and some music education and technology provision may still charge a fee to enthusiasts willing to pay for their hobby. Yet in a world without worries about attribution, cheating is of no interest, since it really makes no odds where interesting music arises, and the soundtrack of our lives will be so well catered for by inexhaustible personalised music that the biggest danger is the trustworthiness of the musical algorithms in question (do the free or low cost tiers of your music service slip in subliminal advertising, or reinforce particular stereotypes, or hide the really great music from their less prestigious clients...).

### **Musical forms**

We can consider potential new musical forms arising from human-AI interactions.

MIR provides the capacity to analyse musical structure in audio files, in particular by finding a segmentation into distinctive and repeated sections (Peiszer et al. 2008). This allows global viewpoints on the set of possible forms across a corpus. Aside from finding the most typical formal constructions, it may assist with hunting for atypical forms, the interesting outliers that point to new territory, or generalisations off grid. MIR also enables examination of starts, ends and transitions within works; you can imagine creating a catalogue of potential openings for electroacoustic works (where the cliché is either to begin with a loud impulsive event, or a slow fade in). Such data on forms is potential training data for generative models that can reproduce and generalise. Pacing is innately related to human psychology here, but it is trivial to dilate or contract timing, or to make discoveries on a vast or small scale from musical or non-musical data and bring these back to the perceptual pacing of human music (whether guided by knowledge of form already obtained, or simpler-minded sonification trickery).

Musical forms for AIs may take an ironic position on human activity, perhaps modelling the quirks of human musicians. A human band's arguments, crushing van journeys, delays at airports and general Spinal Tap lifestyle might be appealing subject matter for an AI; a piece could be based around a simulation of human musicians falling out with each other. The more intricate our neuroscientific knowledge, the more human psychology and physiology will have potential to inspire interesting musical material; the more our knowledge of artificial intelligence, the more computational musical material may arise. New forms in computer music can arise from reified computational architecture. The deep networks so enamouring signal processing engineers at present may themselves become musical form; think of the hourglass shape of a multi-layered autoencoder as a compositional plan, with an encoder of progressively smaller hidden layers before the opposite trend of the decoder.



If music AI's greatest promise is to reveal our compositional faculty to ourselves, the modelling of ourselves is an ideal form within which to work. We may create ensembles of ourselves (Björk's Army of Me, now an Army of You). New musical situations can arise as we encounter past models of ourselves, simulated at earlier points in our musical career (perhaps continuously recorded in tandem with our musical development by a musical familiar). The strangeness of previous automation is already accessible by projects where a musician records a track each year into a DAW over many years, sets themselves a task seven years hence, and other games with the long term against memory. The AI ease of forwards time travel (save state and run again in a thousand years, give or take issues of digital preservation) can lead to convoluted new twists here, including the simulation of vast choruses of AIs from disparate eras, set to perform with each other in unexpected new virtual worlds. This is reminiscent of the historical figures able to freely intermingle in the afterlife in Bangsian literature (Bangs 1895).

So we might expect music made by humans about the AI condition, music made by AIs about the human condition, and either about themselves. Just as hybridity of styles is one of the most potent creative forces in musical practice, merged situations, the cyborg middleground, will be aesthetically powerful. New forms of human to inhuman (or should that be human to transhuman) transition will become sweeping new compositional devices. New forms of collaboration between humans and AIs at many levels of musical intelligence will proliferate.

If longevity of AIs seems potentially dangerous, even the dystopian nightmare of a single prodigious creator dominating culture for a long period, an alternative is to consider mAIfly, music AIs engineered to short lives. Able to generate perhaps only one single musical output after their intensive training (one HAI wonders), such entities would prioritise a rare uniqueness, to satisfy Leigh Landy's call for the ephemeral over the eternal (Landy 2001).

### **Agents and their Agenda**

The motivations of increasingly autonomous AIs themselves may form a new counterpoint to human discussions of the human motivations for algorithmic composition such as Pearce et al. (2002). Before posthuman music exponentially smotheres existing compositional practice, we have an ethical imperative to consider the implications (Tegmark 2018). Considering our own position as composers in the age of AI, do we relish the idea of replacing part or all of our compositional work? Is the most important goal enhanced self-understanding through musical modelling and engineering? What makes us unique as human composers, which no AI could ever access, the quirky social encounters of our musical lives not held in any database? Do we yearn for qualia or qualAI? Are we stuck in the historical composition mode of pAstIche if we over-emphasise the creation of the bAltles, hAlm, jAlmes Brown, coldpAI, mAIdonna, kAlty Perry, stAvle Wonder, John cAlge, kAlja saarAlho and their peers above the potent new musical possibilities of collaborative and interactive AI systems? Are we nAlve to think that we can create AIs to do our bidding, for we surely would balk at any idea of a musical sIAlve? Are we prepared for future music AI to diverge in its priorities far from our current compositional pre-occupations, to ignore all influence we thought we had on musical culture however modest, including any thoughts laid down in this current article?

## References

- Ames, C. (1987) Automated composition in retrospect: 1956-1986. *Leonardo* 20(2):169–185, 1987
- Ariza, C. (2009) The interrogator as critic: The Turing test and the evaluation of generative music systems. *Computer Music Journal* 33(2), 48-70.
- Bangs, J. K. (1895) *A House-Boat on the Styx*. New York: Harper & Brothers
- Benetos, E., Dixon, S., Giannoulis, D., Kirchhoff, H., and Klapuri, A. (2013) Automatic music transcription: challenges and future directions. *Journal of Intelligent Information Systems* 41(3): 407-434
- Born, G. (1995) *Rationalizing culture: IRCAM, Boulez, and the institutionalization of the musical avant-garde*. Berkeley, CA: University of California Press.
- Briot, J-P., Hadjeres, G., and Pachet, F. (2017) Deep learning techniques for music generation: A survey. arXiv preprint, 2017, arXiv:1709.01620
- Bryan, N. J., and Wang, G. (2011) Musical Influence Network Analysis and Rank of Sample-Based Music. In *Proceedings of the 12th International Society for Music Information Retrieval Conference (ISMIR 2011)*, Miami, FL
- Casey, M., Veltkamp, R., Goto, M., Leman, M., Rhodes, C., and Slaney, M. (2008) Content-based music information retrieval: Current directions and future challenges. *Proceedings of the IEEE* 96(4): 668–696
- K. Choi, K., Fazekas, G., Cho, K. and Sandler, M. (2017) A tutorial on deep learning for music information retrieval. arXiv preprint, arXiv:1709.04396
- Collins, N. (2006) Composing to Subvert Content Retrieval Engines. *ARRAY, the Journal of the ICMA*, Winter 2006
- Collins, N. (2011) Musicians' Attitudes to Musical Influence. *Empirical Musicology Review* 6(2): 103-124
- Collins, N. (2016) Towards Machine Musicians Who Have Listened to More Music Than Us: Audio Database led Algorithmic Criticism for Automatic Composition and Live Concert Systems. *Computers in Entertainment* 14(3): 1-14
- Collins, N. (2018) Origins of Algorithmic Thinking in Music. Book chapter in Alex McLean and Roger Dean (eds.), *The Oxford Handbook of Algorithmic Composition*. New York: Oxford University Press
- Collins, N. (2018b) '...there is no reason why it should ever stop': Large-scale algorithmic composition. *Journal of Creative Music Systems* 3(1) <https://www.jcms.org.uk/article/id/525/>

Collins, N. (2020) The Future of Rhythm. Book chapter in Russell Hartenberger and Ryan McClelland (eds.), *The Cambridge Companion to Rhythm*. Cambridge: Cambridge University Press

Colton, S., and Wiggins, G. A. (2012) Computational creativity: The final frontier? In *Proceedings of the European Conference on Artificial Intelligence*, vol. 12, pp. 21-26

Fiebrink, R. and Caramiaux, B. (2018) The Machine Learning Algorithm as Creative Musical Tool. In: Roger T. Dean and Alex McLean (eds.) *The Oxford Handbook of Algorithmic Music*. New York: Oxford University Press, pp. 181-208

Hayes, W. (1751) *The art of composing music by a method entirely new, suited to the meanest capacity. Whereby all difficulties are removed, and a person who has made never so little progress before, may, with some small application, be enabled to excel*. London: Printed for J. Lion

Holzappel, A., Sturm, B., and Coeckelbergh, M. (2018) Ethical dimensions of music information retrieval technology. *Transactions of the International Society for Music Information Retrieval* 1(1): 44–55

Ingham, T. (2019) TIKTOK OWNER BYTEDANCE BUYS AI MUSIC COMPANY JUKEDECK. Posted JULY 23, 2019 at <https://www.musicbusinessworldwide.com/tiktok-parent-bytedance-buys-ai-music-company-jukedeck/>

Jefferson, G. (1949). The mind of mechanical man. *British Medical Journal* 1(4616), 1105–1110

Kant, D. (2016) The Happy Valley Band: Creative (Mis)Transcription. *Leonardo Music Journal* 26: 76–8.

Landy, L. (2001) From algorithmic jukeboxes to zero-time synthesis: a potential A–Z of music in tomorrow's world (a conference provocation). *Organised Sound* 6(2): 91–96

Knotts, S., and Collins, N. (2020) A survey on the uptake of Music AI Software. *Proceedings of NIME2020*, Birmingham

Loughran, R., & O'Neill, M. (2017) Limitations from assumptions in generative music evaluation. *Journal of Creative Music Systems* 2(1)

McLean, A., Fanfani, G., and Harlizius-Klück, E. (2018) Cyclic Patterns of Movement Across Weaving, Epiplokē and Live Coding. *Dancecult: Journal of Electronic Dance Music Culture* 10(1)

Megill, J. (2014) Emotion, Cognition and Artificial Intelligence. *Minds & Machines* 24: 189–199 <https://doi.org/10.1007/s11023-013-9320-8>

Oh, C., Lee, T., Kim, Y., Park, S., Kwon, S., and Suh, B. (2017) Us vs. them: Understanding artificial intelligence technophobia over the Google DeepMind challenge match. In

*Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 2523-2534

Pearce, M., Meredith, D. and Wiggins, G. (2002) Motivations and Methodologies for Automation of the Compositional Process. *Musicae Scientiae* 6(2):119–147

Peiszer, E., Lidy, T., and Rauber, A. (2008) Automatic audio segmentation: Segment boundary and structure detection in popular music. In *Proceedings of the International Workshop on Learning the Semantics of Audio Signals* 106: 45-59

Riedl, M.O. (2014). The Lovelace 2.0 Test of Artificial Creativity and Intelligence. ArXiv, abs/1410.6142.

Roads, C. (1996) *The Computer Music Tutorial*. Cambridge, MA: MIT Press

SoundCloud (2019) CELEBRATING THE 200 MILLIONTH TRACK UPLOADED TO SOUNDCLOUD: FRANCESCA LOMBARDO'S "EYE RING" Posted February 13, 2019 at <https://blog.soundcloud.com/2019/02/13/celebrating-the-200-millionth-track-uploaded-to-soundcloud/>

Stefik, M., and Smoliar, S. (1996) What Computers Still Can't Do: five reviews and a response. *Artificial Intelligence* 80: 95-97.

Sturm, B. L. (2012) A survey of evaluation in music genre recognition. In *International Workshop on Adaptive Multimedia Retrieval*, pp. 29-66. Springer, Cham

Sturm, B. L., Iglesias, M., Ben-Tal, O., Miron, M., and Gomez, E. (2019) Artificial intelligence and music: open questions of copyright law and engineering praxis. In *Arts* 8: 115 <https://doi.org/10.3390/arts8030115>

Tegmark, M. (2018) *Life 3.0: Being Human in the Age of Artificial Intelligence*. London: Penguin

Turing, A. (1950) Computing Machinery and Intelligence. *Mind* 49: 433-460.

Wiener, N. (1950) *The Human Use of Human Beings*. London: Eyre & Spottiswoode

Xenakis, I. (1992) *Formalized Music*. Stuyvesant, NY: Pendragon Press. New York: Harper & Brothers