In the Name of Creativity: En Route to Inspiring Machines

Chun-yien Chang*
College of Humanities and Social Sciences
National Yang Ming Chiao Tung University
HsinChu City, TAIWAN
ccy@cctcc.art

Ying-ping Chen*
Department of Computer Science
National Yang Ming Chiao Tung University
HsinChu City, TAIWAN
ypchen@cs.nctu.edu.tw

Abstract

In this short paper, we reflect on the long quest for intelligence and creativity of computing machinery as well as the suitability for prevailing machine learning techniques to be used in creative tasks. We believe that modularization and multi-layered structures are among essential ingredients constituting creative minds and may greatly benefit machines on handling creative tasks. For proof of concept, we select musical composition, particularly, Species Counterpoint, as the task, adopt a recently proposed computational framework designed for investigating creativity and the creative process, and present an implementation capable of producing the outcomes that exhibit the desired effect.

Introduction and Reflections

Computing machinery has been fascinating to human beings for quite a long time. Accompanied with the introduction to the idea and design of a programmable, general-purpose, mechanical computer, well known as Babbage's Analytical Engine (Menabrea 1843), almost two centuries ago, expectations and speculations on the potentials, especially in aspects of intelligence and creativity, of such a machine had been boldly made by Lovelace (1843), "..., the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent." A century later, Turing (1950) asked the question, "Can machines think?" to address the intelligence aspect of machines and to argue that machines may eventually exhibit intelligent behavior as playing well in the imitation game. Now, we wish to ask the question, "Can machines create?" Under the current circumstances, we are unable to directly, appropriately answer this question. Instead, in this article, within the scope of musical composition, we would like to make a discussion on the apparent lack of certain essential components, capabilities, and properties that enable or permit machines to create in the present prevailing techniques. Moreover, we provide our preliminary implementation as a viable technical construction with its generated results indicating that the existence of some these essential ingredients brings machines one step further closer to being able to create.

The most prominent, prevailing computational techniques in the related fields of artificial intelligence are undoubtedly the methods in the family of deep learning and artificial neural networks, and in the area of music generation, there have already been enormous studies and results (Briot, Hadjeres, and Pachet 2020). We do not intend to diminish the importance or undermine the practical value of those works, but if methods of this category are adopted, from the viewpoint of creators, or more specifically, music composers, in terms of the present form of the methodology, there are certain limitations in the aspects of formality, capability, and efficacy (Pearl and Mackenzie 2018). The technical framework of deep learning requires a huge amount of data, i.e., existing music pieces in this case, to train models which no matter will be used as classifiers or generators. Insufficient data will be unable to render useful or meaningful outcomes, let alone models that can *create* in the common sense. In history, there are only a few productive music composers creating certain amount of music pieces. Simply according to this situation, in contrast to the usage and requirements of the deep learning methodology, it can be seen that creativity and the action of creating may not properly fit with how deep learning operates and functions.

As to the characteristics of deep learning, please allow us to make an arguable analogy. We know that any boolean function, as long as the truth table is given, in theory, we can directly construct its combinational circuits in a systematic way by analyzing and identifying the essential prime implicants. It is in the fundamentals of logic design. However, in practice, except for certain, usually extremely simple, circuits, most circuits are not designed in this way. Instead, their design process usually incorporates modularization, multi-layered structures, domain knowledge, and even personal experience of designers. Although the techniques of deep learning keeps evolving and advancing, employing popular deep learning techniques on music generation tasks is intrinsically similar to making attempts to piece together superficial elements, like prime implicants in the circuit case, to produce target outcome in the basic, primitive way. For music, such an approach in fact not only conceptually ignores the separate, usually totally different ideas and emotions that the music composer would like to convey and express via individual music pieces but also decontextualizes the music pieces by not considering the essence constitutes the creation such as the historical background, the cultural heritage, and even possibly the factors of instruments, including timbre, registers, and the difficulty to perform.

Proceedings of the 12th International Conference on Computational Creativity (ICCC '21)
ISBN: 978-989-54160-3-5

Moreover, pre-existing knowledge, pre-determined settings, and personal preference or experience are extremely difficult to inject into the use of deep learning methods if at all possible. While the models obtained from deep learning can be presented in detail in the form of many parameters, usually millions, and easily duplicated for replicating the results, the operation as a whole fundamentally forms a black box. Thus, a successful, practically applicable artificial neural network model can provide little information for gaining insights or triggering inspiration. In recent years, while the research directions such as interpretable machine learning and explainable artificial intelligence have emerged (Linardatos, Papastefanopoulos, and Kotsiantis 2021), the advancement is currently quite limited.

Therefore, in this article, we wish to respond to the call made by Turing for making machines intelligent, or in our case, capable of creating. We would like to take a small step towards making machines able or seemingly able to create in the common sense. In order to integrate the computational framework with the concept of modularization and the multi-layered structures of the creating process, we adopt our recently proposed meta-framework, ants on multiple graphs, AntsOMG (Chang and Chen 2020) and one of its showcase, the composition of organum motets (Chang and Chen 2021), in the hope that the essential ingredients in the creating process, especially in music composition, can be observed. Based on the design and properties of AntsOMG, we expect the presented implementation to possess certain characteristics, such as accessibility, scalability, and explainability. Moreover, at the level of technical details, since the implementation is multi-layered and modularized, the "components" of the produced model for composing music can even be separated and swapped with ease. Hence, transfer learning, utilizing pre-existing knowledge, incorporating pre-determined settings, integrating human experience, and the like can be achieved. By conducting research along this line, hopefully injecting creativity into machines may someday be accomplished.

Related Work

While the goal of this study is to investigate creation behavior and mechanisms, in the hope that a small step towards enabling machines to autonomously create may be accomplished, automated music composition is closely related to this article since a particular music genre, organum motets, is adopted as the study subject. Hence, selected studies available in the literature in the realm of computational intelligence related to music generation and automated composition are included in this section for reference. More comprehensive surveys and reviews can be found by (Loughran and O'Neill 2020; Carnovalini and Rodá 2020; Liu and Ting 2017; Herremans, Chuan, and Chew 2017; Lopez-Rincon, Starostenko, and Martín 2018; Briot, Hadjeres, and Pachet 2020; Gifford et al. 2018; Fernández and Vico 2013).

Evolutionary algorithms are population-based, stochastic optimization methodologies relatively easy to be used to handle a variety of tasks of very different nature. Because of their flexibility and versatility, they have been utilized for generating music decades ago. Genetic algorithms, one of the major evolutionary algorithms, have been used in the task of computer-assisted music composition (Horner and Goldberg 1991; Jacob 1995; Marques et al. 2000). More complicated music constructs are also considered by researchers, including chord progression (Kikuchi and Osana 2014), measures and phrases (Ting and Wu 2017), harmonization (Donnelly and Sheppard 2011), and music theory (Liu and Ting 2012).

For other branches of evolutionary algorithms, genetic programming has been applied to evolve music generation (Phon-Amnuaisuk, Law, and Kuan 2009) and to compose the 16th-century counterpoint (Polito, Daida, and Bersano-Begey 1997). Ant colony optimization has also been adopted to generate music (Guéret, Monmarché, and Slimane 2004) and to create Baroque harmonies (Geis and Middendorf 2008).

Proof of Concept and Outcomes

Aiming at enabling machines to create, as aforementioned, we select classical music as scope, in particular, Species Counterpoint. We will firstly give some background regarding Species Counterpoint as it is closely related to what we would like machines to have, what role such a mechanism plays in the creative process, and what effect can be observed when machines have such a capability.

During the period of conventional music theory, approximately from Johann Joseph Fux's famous textbook *Gradus ad Parnassum* (Fux 1725) to its decline after Franz Schubert (Mann 1994), Species Counterpoint has been an essential way to pursue the ability of composition. Even the great composer Ludwig van Beethoven also learned Species Counterpoint from Joseph Haydn and left many manuscripts of his studies (Mann 1970). It is also a well-organized and constructive way to approach the composition of polyphonic music. Fux's title *Gradus ad Parnassum* (Steps to Mount Parnassus) revealed the step-by-step essence to construct the required knowledge for the sixteenth-century counterpoint.

Species Counterpoint is not only a step-by-step guide, it also exhibits the building blocks essential to a classic type of creative behavior:

- Monophony: The structural rules for both cantus firmus and counterpoint melodies;
- First Species: The backbone of all contrapuntal interactions consists of note-against-note consonant intervals;
- Second Species: Introducing the usage of dissonance (passing tones) and the concept of strong-weak beats;
- Third Species: Introducing more types of dissonances including neighboring tones and cambiatas, and more differentiated contrapuntal layers and rhythmic structure;
- Fourth Species: Introducing the most important type of dissonance at that time (suspensions), and the translocation of strong beat by shifting the notes;
- Fifth Species: Integration of all previous Species to achieve a fully organic unfolding of melodic beauty in the context of counterpoint;

  ... and more.

Figure 1: An example of black-box AI. Google (2019) launched their first AI-powered Doodle *Celebrating Johann Sebastian Bach* on March 21, 2019. We use Bach's own chorale melody fragment from "Ach wie flüchtig, ach wie nichtig" (Bach 2008) to get the results of harmonization twice: The first result contains a bass line that is very awkward for human voice and also a strange chord in the beginning of the second measure; the second result comprises the uncompleted inner voices which are not even correctly notated. It demonstrates the problems of black-box AI: the machine still lacks any basic ideas of Bach Chorales even after receiving the so-called training.

This definite building process does not limit the development of composers, like Haydn or Beethoven, but rather serves as the foundation of their exploration of distinct personal styles. This makes us reflect on whether it is sufficient to solely regard the black-box machine creativity as the mainstream methodology while exploring computational creativity, an even broader realm of artificial intelligence. What are the inner foundations the composers have built while learning the tedious counterpoint rules, so that they not only gain the ability to write the counterpoint exercises like a machine but also develop their own styles upon it?

In order to explore this topic, we began to construct the building blocks of Species Counterpoint and to make attempts to "inject" this longlast heritage of music creativity into machines in a white-box way—different from black-box AI—to explore the different possibilities of machine creativity. (See Figure 1)

To demonstrate the different possibilities of machine composition, we formulated a musical form with elements from First Species Counterpoint, medieval organums, and Renaissance motets and implemented a computational framework to enable machines to compose a full length musical piece with its own compositional inclination.

First, we implemented the algorithm to generate the plainchant based on the rules of the following melodic intervals (Jeppesen 1992):

- Ascending and descending: major and minor second, major and minor third, perfect fourth, perfect fifth, and perfect octave.

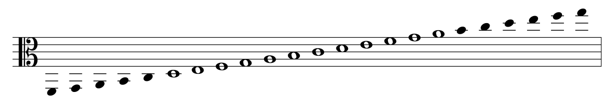

Figure 2: The gamut. White notes stand for the range of cantus firmus, and black notes stand for extended range of upper/lower counterpoint parts.

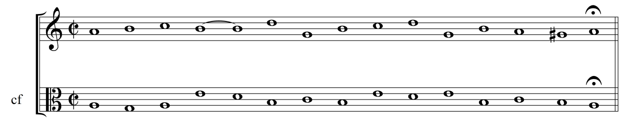

Figure 3: An example of a generated phrase of First Species Counterpoint.

- Ascending only: minor sixth.

For Renaissance vocal counterpoint, a melody is not only a series of numbers but also a medium suitable for human voice singing and conveying religious feelings. Therefore, we constructed the algorithm operating with multiple graphs: a *gamut graph* derived from the pitch set (Figure 2) and the aforementioned melodic interval rules as well as a *meta graph* representing the thinking process of composing melodies to regulate the output from the *gamut graph*.

Next, based on this algorithm, we implemented the most important rules of harmonic intervals from First Species Counterpoint: (Jeppesen 1992)

- Only consonant combinations may be used (the fourth is considered a dissonance).
- One must begin and end with a perfect consonance (octave, fifth, and so on). However, if the counterpoint lies in the lower part, only the octave or unison may be used at both beginning and the ending.

We expanded the gamut to allocate the upper and lower counterpoint parts (Figure 2) and transformed it with the First Species harmonic intervals rule set and the assigned cantus firmus into a new graph for counterpoint composition. Figure 3 is an example of a musical excerpt from one of the generated compositions.

In this article, we intend to demonstrate the "Neo-classical" fun that we discovered over the process of the injection of creative knowledge. In the early twentieth century when many creative methods flourished, Neo-classicism is among the most remarkable ones, which is somewhat sarcastic in the interwar atmosphere. Russian composer Sergei Prokofiev's comment "Bach on the wrong notes" (Tierney 1977) on his fellow countryman Igor Stravinsky (See Figure 4), vividly and arguably indicated the impression of Neo-classical works on people. When Prokofiev composed his "Classical" Symphony, he used an interesting metaphor: "It seemed to me that has Haydn lived to our day he would have retained his own style while accepting something of the new at the same time. That was the kind of symphony I wanted to write: a symphony in the classical style." (Prokofiev and Shlifstein 2000)

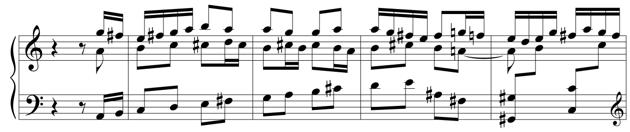

Figure 4: Music excerpt from Movement I of Stravinsky's Concerto for Piano and Wind Instruments (Piano solo part, from measure 49). Here it exhibits a taste of Bach's three-voice works, except that it is "on the wrong notes".

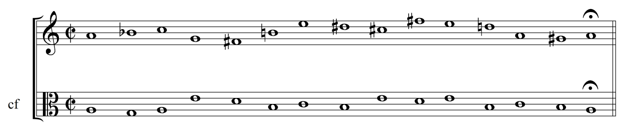

Figure 5: An example of a phrase of an organum motet generated from the replaced *gamut graph*, based on the cantus firmus in Figure 3.

Taking one's intelligent and emotional experiences as a foundation, while constantly pursuing the innovation of creative elements, is the lifelong effort of the so-called "serious" music composers. However, just take the example of the combination of pitch sets, when a human composer has gone through the process from being a student, all the way to the point when a revolution becomes apparent, this person has accumulated a vast amount of experience on composition. In other words, a strong and inevitable tendency has been formed in the years over one's development, it is never easy to switch like an instantly replaced component. This may be one of the reasons why Neo-classicism has seemed attractive and challenging to many twentieth-century composers. Most of them have more or less ventured the route of Neo-classicism, regardless of what their own mature styles may be.

For machines, the switching of algorithmic components is undoubtedly one of their strengths, as long as they are constructed or expressed in the form of a white box. Take the first species counterpoint as an example, under the identical harmonic and melodic interval rules, swapping the gamut graph and the corresponding score functions renders interesting results, as shown in Figure 5, of a very different style.

Over the presented process, we adopted a computational framework, called AntsOMG, specifically developed for investigating creativity and the creative process. AntsOMG inherits certain characteristics of ant colony optimization, ACO (Dorigo and Gambardella 1997), letting machines develop a set of style models through a large quantity of practices, and we slightly loosen some of the rules to amplify this tendency to allow ample space for development. Our implementation is not limited only to produce counterpoint fragments but to create a complete musical piece called organum motet, which comprises several segments of plainchants and counterpoint based on them. Thus the aforementioned swapping is also applicable to the generation of the complete organum motets—with a stylistic twist. More examples of both types of organum motets are provided in the supplemental material.

While retaining most of the building blocks and the style development mechanism, we have changed the creative phenomenon produced by machines merely by swapping a single graph, which may be intriguing in the context of transfer learning because of the nature of modularization and its multi-layered design. As the present form of AntsOMG, the weights/parameters on the developed model, unlike those on artificial neural networks, are quite interpretable, and consequently the behavior of the whole algorithmic construct may be explainable. It is a simple display of the benefits of injecting white-box knowledge into machines. In the meanwhile, it opens up many possibilities of machine creativity through the unlimited combinations or swapping of the building blocks, such as the pre-existing or newly generated rule sets, creative materials, and even the outcomes from the interactions between them. The diversity of computational methods and their combinations, along with the current achievement of machine learning, will largely enrich the field of computational creativity, and even human creators may further benefit from being inspired by the plentiful spectrum created from the landscape of machine strengthened creativity.

Supplemental material of this paper, including scores, music, and source code information, is available and can be accessed online via `https://e.cctcc.art/iccc21` or `https://github.com/nclab/iccc-21`.

Conclusive Notes

In this article, we presented our ideas and thoughts on what machines must have in order to possess the capability of creating. Human creators learn the domain knowledge, gain their personal experience, and accordingly cultivate their own styles during the time of their existence. Machines are intrinsically different from humans and do not need to act like humans. However, before machines are able to independently, without any human help or dictation, exhibit a sufficient level of intelligence, investigating how creative minds operate and gaining insights into the creative process are necessary to make progress towards inspiring machines.

Acknowledgments

The authors would like to thank Taiwanese renowned cellist Dr. Sheng-wen Liu for her professional advices and Mr. Chih-fang Chiang for his assistance with the experiment data processing and music rendering.

The work was supported in part by the Ministry of Science and Technology of Taiwan under Grant MOST 110-2221-E-A49-076.

References

Bach, J. S. 2008. *389 Choralgesänge für vierstimmigen gemischten Chor*. Wiesbaden: Breitkopf & Härtel. Edited by B. F. Richter.

Briot, J.-P.; Hadjeres, G.; and Pachet, F.-D. 2020. *Deep Learning Techniques for Music Generation*. Springer.

Carnovalini, F., and Rodá, A. 2020. Computational creativity and music generation systems: An introduction to the state of the art. *Frontiers in Artificial Intelligence* 3(14).

Chang, C.-Y., and Chen, Y.-P. 2020. AntsOMG: A framework aiming to automate creativity and intelligent behavior with a showcase on cantus firmus composition and style development. *Electronics* 9(8). article 1212.

Chang, C.-y., and Chen, Y.-p. 2021. Contrapuntal composition and autonomous style development of organum motets by using AntsOMG. In *Proceedings of 2021 IEEE Congress on Evolutionary Computation (CEC 2021)*, 2023–2030.

Donnelly, P., and Sheppard, J. 2011. Evolving four-part harmony using genetic algorithms. In *Applications of Evolutionary Computation (EvoApplications 2011)*, 273–282.

Dorigo, M., and Gambardella, L. M. 1997. Ant colony system: A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation* 1(1):53–66.

Fernández, J. D., and Vico, F. 2013. AI methods in algorithmic composition: A comprehensive survey. *Journal of Artificial Intelligence Research* 48(1):513–582.

Fux, J. J. 1725. *Gradus ad Parnassum, sive manuductio ad compositionem musicæ regularem: Metyhodo novâ, ac certâ, nondum antè tam exacto ordine in lucem edita*. Wien: Johann Peter van Ghelen.

Geis, M., and Middendorf, M. 2008. Creating melodies and Baroque harmonies with ant colony optimization. *International Journal of Intelligent Computing and Cybernetics* 1(2):213–238.

Gifford, T.; Knotts, S.; McCormack, J.; Kalonaris, S.; Yee-King, M.; and d'Inverno, M. 2018. Computational systems for music improvisation. *Digital Creativity* 29(1):19–36.

Google. 2019. https://www.google.com/doodles/celebrating-johann-sebastian-bach. Accessed on June 17, 2021.

Guéret, C.; Monmarché, N.; and Slimane, M. 2004. Ants can play music. In *Proceedings of International Workshop on Ant Colony Optimization and Swarm Intelligence (ANTS)*, 310–317.

Herremans, D.; Chuan, C.-H.; and Chew, E. 2017. A functional taxonomy of music generation systems. *ACM Computing Surveys* 50(5). Article No. 69.

Horner, A., and Goldberg, D. E. 1991. Genetic algorithms and computer-assisted music composition. In *Proceedings of 1991 International Conference on Genetic Algorithm (ICGA 1991)*, 437–441.

Jacob, B. L. 1995. Composing with genetic algorithms. In *Proceedings of 1995 International Computer Music Conference (ICMC)*, 452–455.

Jeppesen, K. 1992. *Counterpoint: The Polyphonic Vocal Style of the Sixteenth Century*. Dover Books on Music. Dover Publications.

Kikuchi, M., and Osana, Y. 2014. Automatic melody generation considering chord progression by genetic algorithm. In *Proceedings of 6th World Congress on Nature and Biologically Inspired Computing (NaBIC)*, 190–195.

Linardatos, P.; Papastefanopoulos, V.; and Kotsiantis, S. 2021. Explainable AI: A review of machine learning interpretability methods. *Entropy* 23(1).

Liu, C.-H., and Ting, C.-K. 2012. Polyphonic accompaniment using genetic algorithm with music theory. In *Proceedings of 2012 IEEE Congress on Evolutionary Computation (CEC 2012)*, 1–7.

Liu, C.-H., and Ting, C.-K. 2017. Computational intelligence in music composition: A survey. *IEEE Transactions on Emerging Topics in Computational Intelligence* 1(1):2–15.

Lopez-Rincon, O.; Starostenko, O.; and Martín, G. A.-S. 2018. Algoritmic music composition based on artificial intelligence: A survey. In *Proceedings of 2018 International Conference on Electronics, Communications and Computers (CONIELECOMP)*, 187–193.

Loughran, R., and O'Neill, M. 2020. Evolutionary music: applying evolutionary computation to the art of creating music. *Genetic Programming and Evolvable Machines* 21:55–85.

Lovelace, A. 1843. Translator's notes to M. Menabrea's memoir on Babbage's Analytical Engine. *Scientific Memoirs* 3:691–731.

Mann, A. 1970. Beethoven's contrapuntal studies with Haydn. *The Musical Quarterly* 56(4):711–726.

Mann, A. 1994. *The Great Composer as Teacher and Student: Theory and Practice of Composition*. Dover books on music. Dover.

Marques, M.; Oliveira, V.; Vieira, S.; and Rosa, A. C. 2000. Music composition using genetic evolutionary algorithms. In *Proceedings of 2000 IEEE Congress on Evolutionary Computation (CEC 2000)*, 714–719.

Menabrea, L. F. 1843. Sketch of the Analytical Engine invented by Charles Babbage. *Scientific Memoirs* 3:666–690. Originally published in French in *Bibliothèque Universelle de Genève* 82, October, 1842.

Pearl, J., and Mackenzie, D. 2018. *The Book of Why: The New Science of Cause and Effect*. Basic Books.

Phon-Amnuaisuk, S.; Law, E. H. H.; and Kuan, H. C. 2009. Evolving music generation with SOM-fitness genetic programming. In *Proceedings of EvoWorkshops on Applications of Evolutionary Computing*, 557–566.

Polito, J.; Daida, J. M.; and Bersano-Begey, T. F. 1997. Musica ex machina: Composing 16th-century counterpoint with genetic programming and symbiosis. In *Proceedings of 6th International Conference on Evolutionary Programming (EP VI)*, 113–124.

Prokofiev, S., and Shlifstein, S. 2000. *S. Prokofiev: Autobiography, Articles, Reminiscences*. Honolulu, Hawaii: University Press of the Pacific. Compiled, edited and notes by L. Shlifstein; translated from the Russian by R. Prokofieva.

Tierney, N. 1977. *The Unknown Country: A Life of Igor Stravinsky*. London: Hale.

Ting, C.-K., and Wu, C.-L.and Liu, C.-H. 2017. A novel automatic composition system using evolutionary algorithm and phrase imitation. *IEEE Systems Journal* 11:1284–1295.

Turing, A. M. 1950. Computing machinery and intelligence. *Mind* 59:433–460.