



HAL
open science

O.R. and music generation

Dorien Herremans, Elaine Chew

► **To cite this version:**

Dorien Herremans, Elaine Chew. O.R. and music generation. 2018, <10.1287/orms.2018.01.03/full/1/14>. <hal-03687099>

HAL Id: hal-03687099

<https://hal.science/hal-03687099v1>

Submitted on 3 Jun 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization



February 5, 2018 in [O.R.](#)

O.R. and music generation


The role of operations research in state-of-the-art automatic composition systems: challenges and opportunities.

By Dorien Herremans, Elaine Chew

SHARE: [f](#) [in](#) [twitter](#) [envelope](#)


PRINT ARTICLE: [print](#)

<https://doi.org/10.1287/orms.2018.01.03>



The digital marketplace is rapidly changing the global music industry. Image © Pavel Ignatov | 123rf.com

Operations research (O.R.) is a diverse field centered on the modeling and solving of complex problems. The robustness of O.R. techniques is evidenced in their successful applications to problems as diverse as finding optimal phone keyboard layouts [26] and modified driven equilibrium fourier transform sequence optimization to improve 3-D brain imaging scans [7]. Apart from addressing problems of industrial utility and health, operations research has also been applied to problems farther afield in art and design such as architectural layouts [8] and the finding of musical solutions that satisfy compositional constraints [28, 10]. It is the last category of problems that forms the focus of this article – the role of operations research in state-of-the-art automatic composition systems, the challenges and opportunities. In particular, we highlight the contributions

of O.R. modeling in the context of the MorpheuS project [30, 15]. For a more complete overview of music generation systems, the reader is referred to Herremans et al. (2018) [17]. 

The Digital Music Economy

The digital marketplace is rapidly changing the global music industry. In 2014, the industry's digital revenue increased by 6.9 percent, reaching a total of \$6.85 billion [9]. Vast amounts of music-related apps have surfaced for smartphones. These easy-to-use and seemingly simple apps are powered by clever algorithms, often inspired by techniques common to O.R.; cases in point include Shazam and Spotify. In this article, we aim to highlight an up-and-coming area that has only recently gained popularity: automatic music generation.

Automatic music generation serves many purposes. The objective is not to replace composers – the algorithms still rely heavily on human-conceived designs and are nowhere near producing high art music – but serve the purpose of creating changeable music in scalable ways for more mundane recurring scenarios. Imagine a world where small independent video makers can access large amounts of copyright-free and bespoke music for use as background music to their videos, or where a user can generate a novel soundtrack to accompany their vacation photo slideshow for online sharing. What if a computer game could have different and new music that follows the suspense each time the player enters the game? It is scenarios like these that have provided researchers with the motivation for developing methods for automatic music generation.

The idea of automatically composing music existed long before the first computers were built. In the words of Ada Lovelace, widely regarded to be the world's first programmer, as she wrote about the general-purpose computing machine:

"[The Engine's] operating mechanism might act upon other things besides numbers [...] Supposing, for instance, that the fundamental relations of pitched sounds in the signs of harmony and of musical composition were susceptible of such expressions and adaptations, the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent."

– Lovelace, 1843 [18]

Today, several techniques exist for generating music that sounds reasonable at the local scale. The big challenge that the field faces is that of creating music that possesses long-term structure and large-scale coherence. We believe that looking at music generation as an optimization problem offers a viable framework for solving this problem.

Optimization-Based Approaches to Music Generation

Composing music can be construed as a constraint satisfaction or combinatorial optimization problem. Composers such as Stravinsky have likened the composition process to one of following self-imposed constraints:

"My freedom consists in my moving about within the narrow frame that I have assigned myself for each one of my undertakings. ... The more constraints one imposes, the more one frees one's self of

the chains that shackle the spirit.”
 – **Stravinsky (as quoted in Strauss [27])**



The task is then to find the right combination of notes that best satisfies these constraints. But how does one judge if a piece of music is good, or if one solution is better than another? The late music theorist Edward Cone has written that “a good composition manifests its structural principles” and that most music criticism “sets about discovering the standards implied by a given work and testing how well it lives up to them” [6]. This theory thus offers a way to assess the quality of a composition, but evaluation – how we construct an objective function that effectively measures the quality of music – remains a central challenge in music generation.

Quantifying Music Quality




Figure 1: Multiple embedded helices in the spiral array representing pitch classes, major and minor chords (triads), and major and minor keys (Figure reproduced from Chew 2014, p.59 [4]).

There are three main ways to evaluate a musical piece. The first is via *human evaluation*. Arguably the best way to find out if a piece sounds good is to test it in front of a live audience. Unfortunately, this technique does not scale well and would be impossible to embed in a quantitative algorithm as it would cause huge delay. This problem is often referred to as the human fitness bottleneck [2].

Next, we could turn to music theory to obtain a set of rules that defines what is permissible in a given style. Music generation algorithms that rely on rules focus primarily on counterpoint, a style of music popular in the Renaissance period that has multiple harmonically interdependent voices [10, 23]. Counterpoint is one of the most formally defined musical styles; for most other genres of music, no such well-behaved set of rules exist.

A third, more robust technique is to use *machine learning* to construct the objective function. Since the 1950s, Markov models have been used to capture statistical properties of musical pieces and genres [24, 3]. While these models can capture many aspects and features of music, their traditional sampling methods (e.g., random walk, Gibbs sampling) do not provide a method for generating structure and thus the resulting pieces often lack long-term coherence. More recently there has been interest in using deep neural networks for composition [16,17], yet the problem of constraining long-term structure remains.

Herremans et al. [12] developed a method to integrate Markov models in the objective function of a metaheuristic. The approach allows for the resulting music to be evaluated based on a machine-learned model, yet supports the use of hard constraints to fix a larger structure. In our Morpheus system [13, 14], this approach is further expanded to tackle complex polyphonic music and automatically detected recurring patterns. We also developed a way to combine machine learning and music theory to construct a new type of objective function based on musical tension.

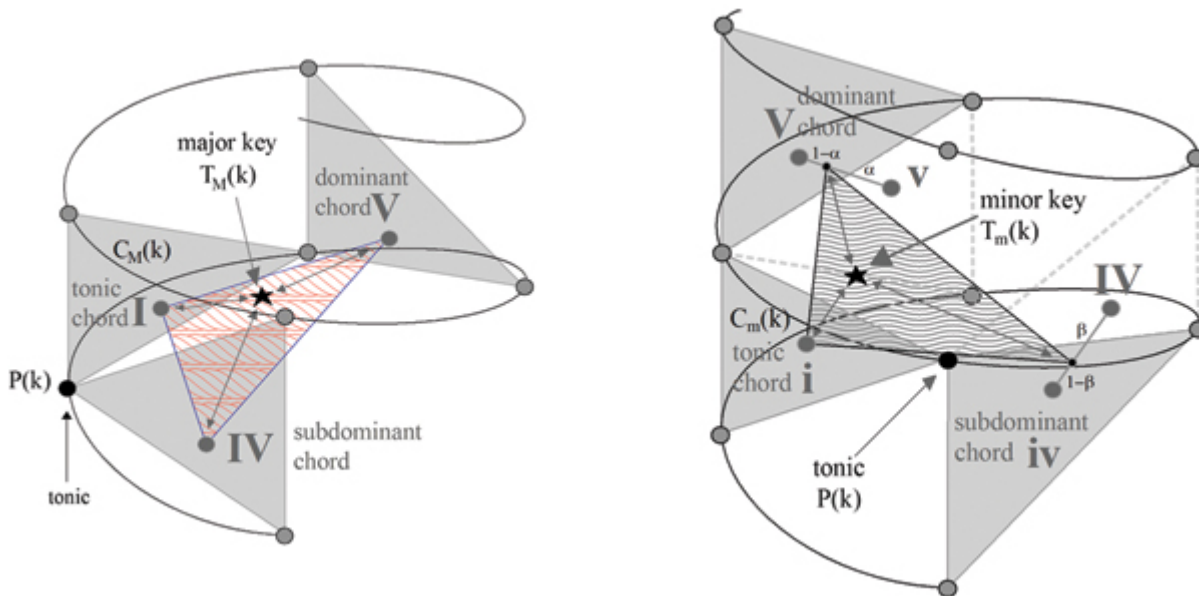


Figure 2: Spiral array pitch class, chord, and key representations: chords are weighted combination of their component pitches, and keys of their defining tonic (rooted on the first scale degree), subdominant (rooted on the fourth scale degree), and dominant (rooted on the fifth scale degree). (Figures reproduced from Chew 2014, p.54-55 [4]).

Representing Music




Figure 3. Elaine Chew performing with the MuSA_RT app [31] at the New Resonances Concert of the 2012 International Computer Music Modeling and Retrieval Conference.

In order to work with music in a quantitative way, we need a mathematical representation. Typically, music consists of a number of notes, each with properties such as pitch (often expressed as a MIDI pitch), duration, start time and beat number. In the MorpheuS project, the music generation problem is re-cast as one of finding new pitches that satisfy certain structural constraints given a rhythm template extracted from an existing piece. Our goal is thus to morph a pre-existing piece into a new composition.

Almost all the music that we hear is tonal, i.e., adhering to a system of hierarchical pitch relations. The tonal properties of a piece of music and how they evolve over time shape the listener's perception of musical stability and tension. The tension profile of a piece is one of the main defining characteristics that imbues upon the piece long-term structural, which we use as a structural constraint in the MorpheuS project. This tension profile is calculated based on the spiral array model for tonality [4].

The spiral array consists of multiple embedded helices (see Figure 1), representing tonal entities at three hierarchical levels: pitch classes, chords and keys. The pitch class helix is effectively the line of fifths (most simply explained as pitches having a simple 2:3 frequency ratio) wrapped around a cylinder so that pitch classes align vertically every four quarter turns. Chords and keys are each generated as convex combinations of their defining elements as shown in Figure 2.

A real-time implementation of the spiral array and its tonal analysis algorithms has been used in live performances [5] to visualize tonal relationships in a music piece (see Figure 3).

Tension Ribbons




Figure 4: Joplin's German Sixth chord mapped to the spiral array; the chord in the key of Eb major consists of the notes Cb, Eb, F# and A.

Leonard Meyer [22] ascribes the emotional content of music to the composer's choreography of expectation. Expectation suspended elicits tension, and expectation fulfilled leads to resolution of that tension. The MorpheuS system aims to recreate the palpable tension that music generates; to this end, we have developed a model that captures tonal tension based on the spiral array [13].

The spiral array tension model first segments a piece into equal length subdivisions and maps the notes to clouds of points in the array. Three aspects of tonal tension are captured using these clouds: the *cloud diameter* measures the dispersion (an indicator of dissonance) of note clusters in tonal space; the *cloud momentum* measures the movement of pitch sets (change in tonal context) in the spiral array; and, *tensile strain* measures the distance between the local and global tonal context. Each attribute can be visualized as tension ribbons on the score (see Figure 5).

Figure 4, shows a well-known tense chord in the spiral array: the German Sixth. It is noticeably spread out in the tonal space. The German Sixth chord occurs on the last beat of the second bar of Scott Joplin's "Binks' Waltz" (Figure 5). When this chord occurs, there is a noticeable increase in the tension ribbons representing cloud diameter and cloud momentum in the spiral array model. The tensile strain starts to increase slightly earlier, adding to the buildup of tonal suspense leading up to the German Sixth chord.



(a) Cloud diameter



(b) Cloud momentum



(c) Tensile strain

Figure 5: Tension ribbons of the German Sixth chord in Scott Joplin's Binks' Waltz (1905).

Patterns in Music

Repetition forms another important aspect of music structure. Recurrent themes and motives are reinforced with each repetition and stay in the listener's mind through the listening experience. A study by Margulis [19] revealed that listeners rated unfamiliar music that had repetitions as more enjoyable, interesting and artistic than the original non-repeated version. With repetition, patterns begin to form; Plato alludes to the importance of patterns in music:

"I would teach children music, physics and philosophy; but most importantly music, for the patterns in music and all the arts are the keys to learning."

– **Plato as cited by Plato & Bloom [25]**

Thus, in MorpheuS, we aim to generate music with patterns having the same length and that recur at the same places as the original template piece. In order to generate music with recurring patterns, we need to first find them. Compression algorithms such as SIATECCompress [20] and COSIATEC [21] can be used to find maximal translatable patterns, i.e., the largest subset of consecutive notes that recur either verbatim or transposed by a constant interval.

Using the elements discussed above, we can define music generation as an optimization problem. MorpheuS takes as an input a template piano piece. From the exemplar piece, the rhythmic structure

is extracted and retained as a template. The algorithm aims to find new pitches for each note in the original piece, and thus morph the piece into a new one. Recurring patterns in the template piece are detected using COSIATEC. These patterns form hard constraints during the optimization so as to enforce perceptible musical motifs and to ensure long-term coherence. The objective of the algorithm is to generate music with a specific tension profile; the user can choose to have the algorithm match the tension profile of the template piece, or follow a user-given tension profile.

To solve the multi-objective tension profile matching problem, we have implemented an efficient variable neighborhood search (VNS) metaheuristic based on the algorithm developed by Herremans [11] for generating counterpoint music. This original VNS has been extended – with a move that changes multiple notes in a time slice – in order to handle more complex polyphonic music.

Conclusions

Systems like MorpheuS can now compose music with limited global structure, which may well be a starting point for scalable systems for automatic music generation. The output is promising, but still falls short of human creations. Take, for example, the issue of playability. The MopheuS-Haydn piece contains many awkward leaps and hand crossings; although these are not impossible, they do not often appear in human compositions. In the literature, apart from some algorithms for rudimentary fingering [1, 29], quantifying playability (which is far more than fingering) remains an important challenge.

Secondly, little is known about how the interactions of local and global patterns generate coherence. Could there be a way to generate viable patterns without relying on a template?

Third, the shaping of music is closely tied to body movement, and feelings such as stress and lilt are part and parcel of musical communication. There are currently no known methods for, say, generating a lilting feeling or incorporating a danceability factor, which is not automatic even when the time signature is three-quarters (as in a waltz).

Finally, aspects such as having the time to breathe and to respire are essential components of human and thus musical communication. How does a composer modulate, beyond tonal tension, other factors for moments of release? As we consider and tackle these challenges, we find ourselves faced with the deeper question of what it means to be human.

References & Notes

1. Balliauw, M., Herremans, D., Palhazi Cuervo, D., & Sörensen, K., 2017, "A variable neighborhood search algorithm to generate piano fingerings for polyphonic sheet music," *International Transactions in Operational Research*, Vol. 24, No. 3, pp. 509-535.
2. Biles, J. A., 2001, "Autonomous GenJam: eliminating the fitness bottleneck by eliminating fitness," in *Proceedings of the 2001 Genetic and Evolutionary Computation Conference Workshop Program*, San Francisco.
3. Brooks Jr, F. P., Hopkins Jr, A. L., Neumann, P. G., & Wright, W. V., 1957, "An experiment in musical composition," *IRE Transactions on Electronic Computers*, Vol., pp. 175-182.

4. Chew, E., 2014, "Mathematical and Computational Modeling of Tonality: Theory and Applications," International Book Series on Operations Research and Management Science, Vol. 204, New York, N.Y.: Springer.
5. Chew, E., & Francois, A. R., 2003, "MuSA. RT: music on the spiral array. real-time," in *Proceedings of the eleventh ACM international conference on Multimedia*, pp. 448-449, ACM.
6. Cone, E., 1960, "Analysis Today," *The Musical Quarterly*, Vol 46, No. 2, pp. 172-188. Retrieved from <http://www.jstor.org/stable/740369>.
7. Deichmann, R., Good, C. D., Josephs, O., Ashburner, J., & Turner, R., 2000, "Optimization of 3-D MP-RAGE sequences for structural brain imaging," *Neuroimage*, Vol. 12, No. 1, pp. 112-127.
8. Flack, R. W., & Ross, B. J., 2011, "Evolution of architectural floor plans" (pp. 313-322), Springer Berlin Heidelberg.
9. IFPI, 2015, "Digital Music Report 2015 - Charting the path to sustainable growth," available online at: <http://www.ifpi.org/downloads/Digital-Music-Report-2015.pdf>
10. Herremans, D., & Sørensen, K., 2013, "Composing fifth species counterpoint music with a variable neighborhood search algorithm," *Expert Systems with Applications*, Vol. 40, No. 16, pp. 6,427-6,437.
11. Herremans D., 2014, "Compose=Compute - Computer Generation and Classification of Music Through Operations Research Methods," Ph.D. thesis, University of Antwerp.
12. Herremans, D., Weisser, S., Sørensen, K., & Conklin, D., 2015, "Generating structured music for bagana using quality metrics based on Markov models," *Expert Systems with Applications*, Vol. 42, No. 21, pp. 7,424-7,435.
13. Herremans D., Chew E., 2016, "Tension ribbons: Quantifying and visualising tonal tension," Second International Conference on Technologies for Music Notation and Representation (TENOR), Cambridge, U.K.
14. Herremans D., Chew E., 2016b, "Morpheus: automatic music generation with recurrent pattern constraints and tension profiles," *Proceedings of IEEE Tencon*, Singapore. DOI: 10.1109/TENCON.2016.7848007
15. Herremans D., Chew, E., 2017, "Morpheus: generating structured music with constrained patterns and tension," *IEEE Transactions on Affective Computing*, pp. 99, DOI: 10.1109/TAFFC.2017.2737984.
16. Herremans D., Chuan C.H., 2017, "Modelling Musical Context with Word2Vec," *Proceedings of the first workshop on deep learning for music joint with IJCNN*, Anchorage, Alaska, DOI: 10.13140/RG.2.2.22227.99364/1.
17. Herremans D., Chuan C.H., Chew, E., 2018, "A Functional Taxonomy of Music Generation Systems," *ACM Computing Surveys*. Vol. 50, No. 5, pp. 69.
18. Lovelace, A. A., 1843, "Notes by AAL" [Augusta Ada Lovelace], "Taylor's Scientific Memoirs," pp. 666-731.
19. Margulis, E. H., 2013, "Aesthetic responses to repetition in unfamiliar music," *Empirical Studies of the Arts*, Vol. 31, No. 1, pp. 45-57.
20. Meredith, D., Wiggins, G., & Lemstrom, K., 2002, U.S. Patent Application No. 10/478,458.
21. Meredith, D., 2015, "Music analysis and point-set compression. *Journal of New Music Research*, Vol. 44, No. 3, pp. 245-270.
22. Meyer, L., 1956, "Emotion and Meaning in Music," The University of Chicago Press.

23. Phon-Amnuaisuk, S., & Wiggins, G., 1999 (April), "The four-part harmonisation problem: a comparison between genetic algorithms and a rule-based system," in *Proceedings of the AISB'99 Symposium on Musical Creativity*, pp. 28-34.
24. Pinkerton, R. C., 1956, "Information theory and melody," *Scientific American*.
25. Plato, Bloom, A. D., 1968, "The republic of Plato," Basic Books, Book III, pp 398-403.
26. Sørensen, K., 2007, "Multi-objective optimization of mobile phone keymaps for typing messages using a word list," *European Journal of Operational Research*, Vol. 179, No.3, pp. 838-846.
27. Strauss J. N., 2003, "Stravinsky the Serialist" in I. Cross (ed.), "The Cambridge Companion to Stravinsky," Cambridge, U.K: Cambridge University Press.
28. Truchet, C., Assayag, G., 2011, "Constraint Programming in Music," London, U.K.: Wiley-ISTE
29. Yonebayashi, Y., Kameoka, H., & Sagayama, S., 2007 (January), "Automatic Decision of Piano Fingering Based on a Hidden Markov Models," in *IJCAI*, pp. 2,915-2,921.
30. cordis.europa.eu/project/rcn/195685_en.html
31. itunes.apple.com/us/app/musa-rt/id506866959?mt=12



Dorien Herremans

Dorien Herremans is an assistant professor at Singapore University of Technology and Design and has a joint-appointment at the Institute of High Performance Computing, A*STAR. Before that she was a fellow at the Centre for Digital Music at Queen Mary University of London. She received her Ph.D. in applied economics on the topic of "Computer Generation and Classification of Music through Operations Research Methods." She graduated as a commercial engineer in management information systems at the University of Antwerp in 2005.



Elaine Chew

Elaine Chew is professor of Digital Media at Queen Mary University of London's School of Electronic Engineering and Computer Science, where she is affiliated with the Centre for Digital Music and serves as its director of Music Initiatives and co-lead of the Cognition, Creativity and Expression research theme. Her scientific research centers on the design of mathematical and computational tools to model, analyze, visualize and manipulate music structures.

SHARE: [f](#) [in](#) [t](#) [✉](#)

Keywords: [O.R.](#); [Music](#)

Recommended

Segmental Pattern Discovery in Music

Darrell Conklin, Christina Anagnostopoulou

1 AUGUST 2006 | [INFORMS JOURNAL ON COMPUTING, VOL. 18, NO. 3](#)

Tonal Description of Polyphonic Audio for Music Content Processing

Emilia Gómez

1 AUGUST 2006 | [INFORMS JOURNAL ON COMPUTING, VOL. 18, NO. 3](#)

Recurrent Neural Networks for Music Computation

Judy A. Franklin

1 AUGUST 2006 | [INFORMS JOURNAL ON COMPUTING, VOL. 18, NO. 3](#)

From the Guest Editor—Special Cluster on Computation in Music

Elaine Chew



1 AUGUST 2006 | INFORMS JOURNAL ON COMPUTING, VOL. 18, NO. 3

Math & Music: The Perfect Match

Elaine Chew

ORMS TODAY

PUBLISHED ONLINE: 2 JUNE 2008

Sign Up for INFORMS Publications Updates and News

Sign Up

SUBSCRIBE

CONTACT

ADVERTISE



The Institute for Operations Research and the Management Sciences

5521 Research Park Drive, Suite 200
Catonsville, MD 21228 USA

phone 1 443-757-3500

phone 2 800-4INFORMS (800-446-3676)

fax 443-757-3515

email informs@informs.org

Get the Latest Updates

Sign Up

[Discover INFORMS](#)

[Explore OR & Analytics](#)

[Get Involved](#)

[Impact](#)

[Join Us](#)

[Recognizing Excellence](#)

[Professional Development](#)

[Resource Center](#)

[Meetings & Conferences](#)

[Publications](#)

[About INFORMS](#)

[Communities](#)

[PubsOnLine](#)

[2022 International Conference](#)

[Certified Analytics Professional](#)

[Career Center](#)

[INFORMS Connect](#)

Copyright 2022 INFORMS. All Rights Reserved

[INFORMS Code of Conduct](#) | [Terms of Use](#) | [Privacy](#) | [Contact INFORMS](#) | [Sitemap](#)

Follow INFORMS on:  [Twitter](#)  [Facebook](#)  [LinkedIn](#)