



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE
DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Evolutionary Music: Statistical Learning and Novelty for Automatic Improvisation

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Barbaresi M., Roli A. (2022). Evolutionary Music: Statistical Learning and Novelty for Automatic Improvisation. Cham : Springer [10.1007/978-3-031-23929-8_17].

Availability:

This version is available at: <https://hdl.handle.net/11585/927180> since: 2023-07-05

Published:

DOI: http://doi.org/10.1007/978-3-031-23929-8_17

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Barbaresi, M., Roli, A. (2022). Evolutionary Music: Statistical Learning and Novelty for Automatic Improvisation. In: Schneider, J.J., Weyland, M.S., Flumini, D., Fuchslin, R.M. (eds) Artificial Life and Evolutionary Computation. WIVACE 2021. Communications in Computer and Information Science, vol 1722. Springer, Cham.

The final published version is available online at: https://doi.org/10.1007/978-3-031-23929-8_17

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Evolutionary Music: Statistical Learning and Novelty for automatic improvisation

Mattia Barbaresi[†][0000–0003–2684–5311] and Andrea Roli^[0000–0001–9891–5441]

Dipartimento di Informatica - Scienza e Ingegneria,
Alma Mater Studiorum Università di Bologna,
Campus of Cesena, Via dell'Università 50, Cesena, Italy
{mattia.barbaresi, andrea.roli}@unibo.it
† Corresponding author: mattia.barbaresi@unibo.it

Abstract. In this work we combine aspects of implicit learning with novelty search in an evolutionary algorithm with the aim to automatically generate melodies with improvisational flavour. Using Markov chains, the technique we present combines implicit statistical knowledge, extracted from musical corpora, with an adaptive novelty search mechanism. The algorithm is described along with the main design choices. Preliminary results are shown in two different musical contexts: Irish music and counterpoint compositions.

Keywords: Evolutionary Art · Computational Creativity · Statistical Learning · Music · Novelty

1 Introduction

Computational Creativity (CC) is a renewed and vivid field of AI research that aims at understanding human creativity while trying to produce “machine creativity”, in which the computer appears to be creative, to some degree [7]. Leaning on such general definitions, approaches to CC often have a twofold perspective: *(i)* developing systems that generate “creative” artefacts and *(ii)* take this opportunity to investigate the cognitive aspects of such processes on a computational basis. Following such practice, this work builds on the statistical approach to implicit human learning and cognitive development, aiming at engineering more human-like and creative computational procedures.

Implicit Statistical Learning (ISL) refers to the general, implicit and ubiquitous ability of the brain to encode temporal and sequential phenomena and more generally, to grasp the regularities in the environment, in an implicit and unconscious way. This approach results from the recent attempt to unify two research venues in psychology and cognitive science, namely Implicit Learning (IL) and Statistical Learning (SL) [26, 9]. Implicit Learning refers more in general to mechanisms and knowledge, in the brain, that are unconscious. Statistical Learning, on the other hand, was initially introduced for language acquisition,

and it is now invoked in various domains of psychology and neuroscience to account for the human ability to detect and use statistical regularities present in the environment [30]. Additionally, many animal species are sensitive to distributional statistics, which suggests that learning from distributional statistics is a domain-general ability rather than a language-specific one [2]. More remarkably, it has been suggested that improvisational musical creativity is mainly formed by implicit knowledge. The brain models music—and other sequential phenomena such as language or movements—as a hierarchy of dynamical systems encoding probability distributions and complexity [14]. SL also plays a role in the production of sequences (e.g. notes or actions); from a psychological perspective, transitional probabilities distributions (TPs) sampled from music may refer to the characteristics of a composer’s implicit (statistical) knowledge: a high-probability transition may be one that a composer is more likely to predict and choose, compared to a low-probability transition corresponding more to an unusual variation [12].

Based on these assumptions, this work aims at combining implicit-knowledge mechanisms with novelty search in a genetic algorithm to emulate an agent’s (i.e. the musician who is composing impromptu) effort to produce novel and interesting sequences of actions (musical pieces) which have to be, at the same time, both familiar (concerning the knowledge initially provided) and novel. We assess our technique in two musical contexts that are characterized by a high degree of improvisation: Irish music and counterpoint.

2 Related Work

The applications of Markov chains in music have a long history dating back to the 1950s [27]: for detailed reviews on AI methods in music, or other examples and techniques, see [8, 23, 19, 22, 16]. Similarly, evolutionary computation has been used for generating music since long [20, 5] and there are currently several systems that generate music by means of an evolutionary technique [4, 24, 22].

From our perspective, however, music generation is just a case study: we focus on modeling a general (context-independent) method for generating sequences (not limited to music) based on implicit mechanisms. In addition, the search towards creativity represents a different approach compared to the more common optimization practice, as the objective function tries to capture several, somehow subjective, features of the piece of art produced.

However, some of the latest and most comparable approaches to this work are perhaps those in [25, 15, 21]. *Continuator* is an interactive music performance system that accepts partial input from a human musician and continues in the same style as the input [25]. The system utilizes various Markov models to learn from the user input. It tries to continue from the most-informed Markov model (higher order), and if a match is not found with the user input, the system continues with the less-informed ones. In [15] the authors describe a method of generating harmonic progressions using case-based analysis of existing material that employs a variable-order Markov model. They propose a method for a hu-

man composer to specify high-level control structures, influencing the generative algorithm based on Markov transitions. In [21] the authors propose to capture phrasing structural information in musical pieces using a weighted variation of a first-order Markov chain model. They also devise an evolutionary procedure that tunes these weights for tackling the composer identification task between two composers. Another work is that of *GenJam* [5]. It uses a genetic algorithm to generate jazz improvisations, but it requires a human to judge the quality of evolved melodies. Finally, *GEDMAS* is a generative music system that composes entire Electronic Dance Music (EDM) compositions. It uses first-order Markov chains to generate chord progressions, melodies and rhythms [1]. In a previous work [3] we conceived a genetic algorithm (GA) which, starting from a given inspiring repertoire and a set of unitary moves, generates symbolic sequences of movements (i.e. choreographies) exploiting similarity with the repertoire combined with the novelty search approach [35].

3 Materials and Methods

According to blind-variation and selective-retention principles of creativity, creative ideas must be generated without full prior knowledge of their utility values [32, 33]. Herein, genetic algorithms offer a natural setting for the blinded-divergent and convergent mechanisms involved in the creative process, terms of diversification and intensification [6]. The evolutionary approach is in itself an exploratory process: the combination of two individuals from the population pool is a combinational process, but the use of a fitness function guides the exploration toward promising areas of the conceptual space, which is bounded and defined by the genetic encoding of the individuals. Losing the fitness function, or having one that is unable to effectively guide the exploration, reverts the mechanism to pure combinational creativity, where elements of the conceptual space are joined and mutated hoping to find interesting unexplored combinations. In this work, we combine an adaptive genetic algorithm with Markov chains—built up from a corpus of music excerpts. The algorithm evolves the parameters (weights) of a constructive procedure that acts on the model and produces new pieces of music that are intended to be novel variations upon familiar music. In addition, the model is used also for evaluating the similarity (the objective function) of generated sequence to the starting knowledge. In this work we build upon the novelty approach presented in [3] and we make it adaptive so as to make it independent on specific ranges of the functions involved in the algorithm. Algorithm 1 shows the main loop and Algorithm 2 shows such a procedure.

3.1 Markov model: chains and score

Markov chains allow us to grasp the statistical structure of sequential phenomena (i.e. music, movements) but also statistical learning and knowledge in humans [13]. It has been observed that transitional probabilities sampled from

Algorithm 1: Pseudo code for GA

```

eval ← mono
for #iters do
  eval(pop)
  offspring, elite ← pop
  offspring crossover and mutation
  pop ← offspring + elite
  eval ← select_objective(pop)
  archive ← archive_assessment(elite)

```

Algorithm 2: Pseudo code for *select_objective()* function

```

eval =  $\begin{cases} \textit{mono} & // \textit{ the Markov Score} \\ \textit{biobjective} & // \textit{ Pareto}(\textit{MarkovScore}(), \textit{Novelty}()) \end{cases}$ 
prevFit = bestFit
bestFit = selBest(pop)
if eval == mono then
  if prevFit ≈ bestFit then
    counter ← counter - 1
    if counter == 0 then
      reset(counter)
      lastAvg = avg(pop)
      eval ← biobjective
    else
      restart(counter)
  else
    bestAvg = avg(pop)
    if lastAvg ≈ bestAvg then
      eval ← mono

```

music (based on Markov models) may also refer to the characteristics of a musician’s statistical knowledge and captures temporal individual preferences in improvisation [11]. We consider here sequences of symbols from a finite alphabet, which can represent e.g. melodies. To model this implicit knowledge, we computed the Markov chains with memory m (or Markov chain of order m) up to the $m = 5$ order¹ starting from a set of musical pieces. For each order m , transitional probabilities are computed for each excerpt as frequency ratios: $P(y|x_m) = \frac{\#x_m y}{\#x_m}$, given a symbol y and a (sub-)sequence x_m (the past) of length m , where $x_m \rightarrow y$ is the inspected transition. As we increase the context size,

¹ In the data we used, orders higher than 5 are not “expressive” because of the data limits and structure: at some point, higher orders contain about the same information held in the previous ones.

the probability of the alphabet becomes more and more skewed, which results in lower entropy. The longer the context, the better its predictive value.

3.2 Objective function

The objective function is intended to capture the familiarity (or the membership, the similarity) of a sequence with respect to the Markov model resulted from the (inspirational) musical corpus. So given a sequence $X = x_0x_1\dots x_n$ the Markov score is defined as the product of TPs of symbols in the sequence

$$score(X) = P(x_0) \times P(x_1|x_0) \times P(x_2|x_0x_1) \times \dots \times P(x_n|x_{n-m}\dots x_{n-1}) \quad (1)$$

where $x_{n-m}\dots x_{n-1}$ is the (past) sequence of length m up to the $(n-1)_{th}$ symbol. For a given chain, it might happen that a transition (past \rightarrow symbol) does not exist. If that actual past does not match a transition in that current order, we shorten the past ($x_{n-m}\dots x_{n-1}$ becomes $x_{n-m+1}\dots x_{n-1}$) and move down a order (i.e. a chain), looking at shorter contextual information to guide the generation. Finally, we apply the negative logarithm to the Markov score and turn the GA objective into a minimization problem:

$$minimize_X : \quad -\log(score(X)) \quad (2)$$

3.3 Encoding

The GA manipulate the parameters of a randomized constructive procedure that acts on the Markov model. The genotype is an array of 6 decimal elements—that sum up to 1—representing the weights to assign to each computed chain in the Markov model. Namely a weight i for each i_{th} -order Markov chain (i.e., a categorical distribution for the chain choice), as in Table 1. Every positional

Table 1. Example of an individual

w_0	w_1	w_2	w_3	w_4	w_5
0.0	0.3	0.05	0.4	0.2	0.05

value of the array weights the probability of the corresponding order in the model when generating a sequence. Notably, these arrays weigh a Monte Carlo process that selects, for each symbol to be emitted in the generation, the order—i.e. the Markov chain—to look at when looking for the transitions to produce it. The phenotype indeed is represented by all sequences generated by the model with that given array of weights (the individual of the GA).

3.4 Adaptive Novelty Search

To steer the generation towards novel productions, we followed the novelty search described in [3, 35]. This algorithm is applied to compensate for a lack of diversity concerning the best individuals already found. It consists of a bi-objective optimization activated when the main objective (e.g. the fitness) stagnates. The novelty is computed as the mean L_2 norm between a genotype (weight vector) and an archive of past genotypes. Using the Markov model, there is no explicit boundary for the objective function since the Markov score depends on the length of the sequence being evaluated. To tackle this problem, we conceived an adaptive mechanism for the activation of novelty based on the results obtained in previous iterations. At each iteration, the algorithm stores the best result obtained. If such value does not change for a given number of iterations, somehow the algorithm is stuck at a minimum. In such a case, the algorithm starts to look at the novelty of individuals. When novelty search has moved the score away of a certain amount from the last best value found, it is turned off and the regular evolution with the Markov score is resumed. See Equation 3.

$$\begin{cases} \text{if } bestFit \simeq prevFit, \text{ for } k \text{ times,} & \text{switch to bi-objective} \\ \text{if } |lastAvg - bestAvg| \simeq stdevLast, & \text{switch to mono} \end{cases} \quad (3)$$

As well as for the objective function, we applied the negative logarithm to novelty too. Thus the bi-objective optimization is intended to minimize both the main objective and the novelty of individuals.

$$\underset{X}{\text{minimize}} : -\log(\text{novelty}(X)) \quad (4)$$

We remark that biasing towards novelty does not mean just adding randomness, but rather diversifying with respect to the best solutions found.

Archive assessment For the assessment of the archive we followed the approach used in [35] except for one aspect; we did not consider a threshold for the individual in order to be added to the archive. Instead we considered, at each iteration of the genetic algorithm, the individuals of the elite, from the elitism process. For these individuals we calculate the dissimilarity as in the mentioned work.

4 Results

The most suitable musical contexts in which our technique can be applied are those in which improvisation plays an important role; but we also need structure to some degree, in such a way that the implicit (soft) constraints imposed by the style can be detected. This way, the music resulting from our method has some amount of novelty, yet still in the style of the examples provided. For our experiments we chose two notable musical contexts: traditional Irish tunes and Orlande de Lassus' *Bicinia* [18]. In this section we first introduce these cases and subsequently we present a selection of the typical results achieved by our technique.

4.1 Irish songs and Bicinia

Traditional Irish music is strongly characterized by its melodies: most old tunes are just melodic (see e.g. [34]) or they are the result of an improvisation upon a given *ground*, i.e. a bass line providing also a harmonic base (see e.g. [28]). In any case, the melodic part of a traditional Irish music is currently the most important component and melodies are usually played with variations, improvising upon a given melodic structure. A large corpus of traditional Irish airs is available in abc notation,² which makes it possible to extract melodies as sequences of symbols, each representing both note and duration. A typical traditional Irish air is shown in Fig. 1. From these airs we extracted all the ones in the key of G and assigned one symbol to each $\langle \text{note}, \text{duration} \rangle$ pair.



Fig. 1. Score of a well known traditional air titled “The south wind”.

The second musical context we have chosen is that of two voices counterpoint, which is one of the simplest and oldest forms of polyphony [29]. In origin, a voice was superimposed to a given one, called *cantus firmus*, in improvisational settings. This original impromptu spirit was subsequently substituted by a more elaborated compositional approach, leading to marvelous multi-voices counterpoints, such as the ones composed by Gesualdo da Venosa. The main technical characteristic of these pieces of music can be summarized in a small set of rules involving the intervals, i.e. the distance in semitones, between the upper and the lower voice. For example, the distance between C and F (above C) is 5 semitones. Obviously, these are not all hard constraints, but some are rather preferences, and they have also been subject to change in time according to different musical aesthetics. Examples of such rules, typical of XVI century counterpoint, are:

- no parallel fifths or octaves are allowed (e.g. C-G cannot move to D-A)
- fifths and octaves should be intercalated by *imperfect consonances*, i.e. thirds and sixths (e.g. an allowed sequence is C-G, C-A, D-A)
- *dissonances*, i.e. all intervals except for unisons, octaves and fifths, should *be prepared* and then *resolve* to a consonant interval by descending (e.g. D-B, C-B, C-A)

² <http://www.norbeck.nu/abc/>

The image shows a musical score excerpt for two voices: Soprano and Tenor. The Soprano part is on a treble clef staff with a key signature of one flat (B-flat) and a common time signature (C). The Tenor part is on a bass clef staff with the same key signature and time signature. The lyrics are written below the notes. The Soprano part has a measure with a circled '5' above it. The lyrics for the Soprano are: "Te de - pre - ca - mur, te de - pre - ca -". The lyrics for the Tenor are: "Te de - pre - ca - - - - - - - - - mur lar - gi - us, te".

Fig. 2. Excerpt of bicinium no. 1, “Te deprecamus”, by de Lassus. Score extracted from https://imslp.org/wiki/Category:Lassus,_Orlande_de.

In our tests, we have taken all the twelve two-voices counterpoint compositions, called *bicinia*, by Orlande de Lassus, which are available as MIDI files.³ In Fig. 2 we show an excerpt of a bicinium by de Lassus. This second context was chosen to assess to what extent our method is able of identifying recurrent patterns and rules typical of a music genre. In this case, we have encoded the twelve MIDI bicinia as sequences of intervals (i.e. distances in semitones between upper and lower voice). As the two voices have in general different durations, we have sampled the music at steps of duration $1/32$ and taken the intervals in semitones, deleting repetitions. This provides the repertoire on which the Markov models are computed. A typical result from our system is a sequence of integer numbers representing intervals in semitones which can be used as a guideline for composing the upper voice upon a given *cantus firmus*.

4.2 Experimental settings

Differently from usual optimization contexts, in our case a good performance does not correspond to the one that leads to the overall best objective function values, but rather to a good balance between similarity (Markov score) and novelty. Therefore, we tuned the parameters of the algorithm trying to attain an effective interplay between score and novelty. The results we present have been obtained with a population of 100 individuals, uniform crossover with probability 0.5, Gaussian mutation ($\mu = 0$, $\sigma = 0.3$) with both chromosome and gene probability equal to 0.35, and 200 generations. The novelty is activated after 5 idle generations (the best score s_{pop} in the current population is stored, along with the standard deviation of the populations scores σ_{pop}) and deactivated when the difference between the score of current best individual and s_{pop} is greater than $\sigma_{pop}/3$. The plot of score and novelty of a typical run is shown in Fig. 3.

We can observe that the score oscillates: whenever the algorithm stagnates, novelty is activated so as to increase diversification. When this latter is high enough, only the Markov score is kept as objective function. In a sense, we can describe the dynamics of the algorithms as a biased exploration of local minima, as typically done by Iterated Local Search techniques [6].

³ <http://icking-music-archive.org/ByComposer1/Lasso.php>.




Fig. 3. Plot of score and novelty of a typical run. Both the functions are to be minimized and novelty is activated, adaptively, only when diversification is needed. The number of individuals in the archive, involved in the calculation of novelty [3], is also plotted.

4.3 Musical results

Due to limited space we can just provide a few examples of the musical results obtained. The generation of melodies inspired to traditional Irish airs has been evaluated by sampling some weight vectors from the final populations and using them to generate actual music. By analyzing the results both through visual inspection and by listening to them, we observed that the music generated is similar to the repertoire provided but with variations and recombinations of patterns. A couple of excerpts are shown in Fig. 4, where we can observe variations of typical Irish melodic and rhythmic patterns: the characteristic *run* (i.e. a fast sequence of notes, typically in a scale) in bars 4, 5 of the first example and the syncopated and composite rhythm in the second one.



Fig. 4. Two typical excerpts of automatically created Irish music.

The second test case concerns de Lassus' *Bicinia*. The main result attained is that it was able to discover the basic rules that characterize two voices counter-

point. In particular the rules extracted that have more strength are: *incipit* with a perfect consonance (unison, octave or fifth), no consecutive octaves or fifths, and dissonant intervals followed by consonant ones—both perfect and imperfect. In Fig. 5 we show an excerpt of the counterpoint produced by applying one of the sequences generated by our algorithm to a given cantus firmus (Chanson CXXVI from manuscript Bibl. Nat. Fr 12744 published by G. Paris). As the algorithm returns a sequence of intervals, it could be used as a tool that assists composers by suggesting feasible note choices, once one of the two voices (*cantus firmus*) is given.




Fig. 5. An example of a two voices counterpoint. The lower voice is the cantus firmus, while the upper voice has been generated by applying a sequence of intervals generated by our algorithm.

In conclusion, for both the contexts the algorithm was able to identify the core regularities and to elaborate around them. The calibration of the parameters is important to achieve a good balance between the tendency of just recombining the patterns learned and the exploration of new possibilities. However, the choice of parameter values does not seem critical, because the combined use of a stack of Markov models of varying orders and novelty search makes it easier to achieve this trade-off.

5 Conclusion and Future Work

The algorithm we have presented has proven to be able to generate novel, yet somehow familiar melodies. An interesting perspective is that of creating non-homogeneous repertoires, maybe just including music that the user likes, without any genre restriction. This way, our system can produce music that merges some of the peculiar features that meet user’s tastes.

Future work is focused on quantitatively assessing the properties of the generated sequences by means of information theory measures, such as block entropies [31] and complexity measures like set-based complexity [17]. Some of these measures can also be introduced in the generative process, so as to limit human evaluation as much as possible. In addition, some metrics can also be used to assess the distance between sequences or to cluster them [10].

As the proposed technique is general and can be applied whenever the goal is to produce sequences of actions, we plan to explore multimodal automatic

generation by combining Markov models from two different contexts, e.g. music and text.

References

1. Anderson, C., Eigenfeldt, A., Pasquier, P.: The generative electronic dance music algorithmic system (gedmas). In: Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (2013)
2. Armstrong, B.C., Frost, R., Christiansen, M.H.: The long road of statistical learning research: Past, present and future (2017)
3. Barbaresi, M., Bernagozzi, S., Roli, A.: Robot choreographies: Artificial evolution between novelty and similarity. In: Finzi, A., Castellini, A., Buoncompagni, L., Anzalone, S. (eds.) Proceedings of the 7th Italian Workshop on Artificial Intelligence and Robotics (AIRO@AIxIA2020). pp. 17–21 (2020)
4. Biles, J.: Improvizing with genetic algorithms: Genjam. In: Evolutionary Computer Music, pp. 137–169. Springer (2007)
5. Biles, J., et al.: Genjam: A genetic algorithm for generating jazz solos. In: ICMC. vol. 94, pp. 131–137 (1994)
6. Blum, C., Roli, A.: Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM computing surveys* **35**(3), 268–308 (2003)
7. Boden, M.A.: Creativity in a nutshell. *Think* **5**(15), 83–96 (2007)
8. Carnovalini, F., Rodà, A.: Computational creativity and music generation systems: An introduction to the state of the art. *Frontiers in Artificial Intelligence* **3**, 14 (2020)
9. Christiansen, M.H.: Implicit statistical learning: A tale of two literatures. *Topics in Cognitive Science* **11**(3), 468–481 (2019)
10. Cilibrasi, R., Vitányi, P.: Clustering by compression. *IEEE Transactions on Information theory* **51**(4), 1523–1545 (2005)
11. Daikoku, T.: Musical creativity and depth of implicit knowledge: Spectral and temporal individualities in improvisation. *Frontiers in Computational Neuroscience* **12**, 89 (2018)
12. Daikoku, T.: Neurophysiological markers of statistical learning in music and language: hierarchy, entropy and uncertainty. *Brain sciences* **8**(6), 114 (2018)
13. Daikoku, T.: Time-course variation of statistics embedded in music: Corpus study on implicit learning and knowledge. *PLoS One* **13**(5), e0196493 (2018)
14. Daikoku, T.: Depth and the uncertainty of statistical knowledge on musical creativity fluctuate over a composer’s lifetime. *Frontiers in Computational Neuroscience* **13**, 27 (2019)
15. Eigenfeldt, A., Pasquier, P.: Realtime generation of harmonic progressions using controlled markov selection. In: Proceedings of ICCX-X-Computational Creativity Conference. pp. 16–25 (2010)
16. Fernández, J.D., Vico, F.: Ai methods in algorithmic composition: A comprehensive survey. *Journal of Artificial Intelligence Research* **48**, 513–582 (2013)
17. Galas, D., Nykter, M., Carter, G., Price, N., Shmulevich, I.: Biological information as set-based complexity. *IEEE Transactions on Information Theory* **56**(2), 667–677 (2010)
18. Haar, J.: *Lassus, orlande de* (2021). <https://doi.org/10/g5nj>, Grove Music Online, accessed on 12 Nov. 2021

19. Herremans, D., Chuan, C.H., Chew, E.: A functional taxonomy of music generation systems. *ACM Computing Surveys (CSUR)* **50**(5), 1–30 (2017)
20. Horner, A., Goldberg, D.: *Genetic algorithms and computer-assisted music composition*, vol. 51. Ann Arbor, MI: Michigan Publishing, University of Michigan Library (1991)
21. Kaliakatsos-Papakostas, M.A., Epitropakis, M.G., Vrahatis, M.N.: Weighted markov chain model for musical composer identification. In: *European conference on the applications of evolutionary computation*. pp. 334–343. Springer (2011)
22. Liu, C.H., Ting, C.K.: Computational intelligence in music composition: A survey. *IEEE Transactions on Emerging Topics in Computational Intelligence* **1**(1), 2–15 (2016)
23. Mor, B., Garhwal, S., Kumar, A.: A systematic literature review on computational musicology. *Archives of Computational Methods in Engineering* **27**(3) (2020)
24. Muñoz, E., Cadenas, J., Ong, Y., Acampora, G.: Memetic music composition. *IEEE Transactions on Evolutionary Computation* **20**(1), 1–15 (2014)
25. Pachet, F.: Interacting with a musical learning system: The continuator. In: *International Conference on Music and Artificial Intelligence*. pp. 119–132. Springer (2002)
26. Perruchet, P., Pacton, S.: Implicit learning and statistical learning: One phenomenon, two approaches. *Trends in cognitive sciences* **10**(5), 233–238 (2006)
27. Pinkerton, R.C.: Information theory and melody. *Scientific American* **194**(2), 77–87 (1956)
28. Playford, J.: *The Division violin: containing a collection of divisions upon several grounds for the treble-violin*. Henry Playford, London, UK, third edn. (1688)
29. Sachs, K.J., Dahlhaus, C.: *Counterpoint* (2001). <https://doi.org/10/g5nk>, Grove Music Online, accessed on 12 Nov. 2021
30. Saffran, J.R., Kirkham, N.Z.: Infant statistical learning. *Annual review of psychology* **69**, 181–203 (2018)
31. Shannon, C.: A mathematical theory of communication. *The Bell System Technical Journal* **27**(1,2), 379–423, 623–656 (1948)
32. Simonton, D.K.: Creativity as blind variation and selective retention: Is the creative process darwinian? *Psychological Inquiry* pp. 309–328 (1999)
33. Simonton, D.K.: Defining creativity: Don't we also need to define what is not creative? *The Journal of Creative Behavior* **52**(1), 80–90 (2018)
34. Various authors: *A collection of the Most Celebrated Irish Tunes Proper for the Violin, German Flute or Hautboy*. John and William Neal, Dublin, Ireland (1724)
35. Vinhas, A., Assunção, F., Correia, J., Ekárt, A., Machado, P.: Fitness and novelty in evolutionary art. In: *International Conference on Computational Intelligence in Music, Sound, Art and Design*. pp. 225–240. Springer (2016)