



HAL
open science

Personalizing AI for Co-Creative Music Composition from Melody to Structure

Ken Déguernel, Mathieu Giraud, Richard Groult, Sebastien Gulluni

► **To cite this version:**

Ken Déguernel, Mathieu Giraud, Richard Groult, Sebastien Gulluni. Personalizing AI for Co-Creative Music Composition from Melody to Structure. Sound and Music Computing (SMC 2022), 2022, Saint-Étienne, France. pp.314-321, 10.5281/zenodo.6573287 . hal-03618015

HAL Id: hal-03618015

<https://hal.science/hal-03618015v1>

Submitted on 7 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.

Personalizing AI for Co-Creative Music Composition from Melody to Structure

Ken Déguernel¹ Mathieu Giraud¹ Richard Groult^{2,3} Sébastien Gulluni

¹Univ. Lille, CNRS, Centrale Lille, UMR 9189 CRIStAL, F-59000 Lille

²Normandie Univ., UNIROUEN, LITIS, F-76000 Rouen

³Univ. Picardie Jules-Verne, MIS, F-80000 Amiens

contact@algonus.fr

ABSTRACT

Co-creativity is a unique artistic situation where human interact with computer, and raises challenges on interaction, steerability, and personalization. We present a new co-creative music composition approach that we used for our participation in the “AI Song Contest 2021”, an international music contest involving artificial intelligence (AI). We *personalize* the artificial creativity methods to adapt to the needs and expectations of a composer. Interactions between the composer and different AI methods occurred throughout the whole composition process, for the generation of melodies, chord progressions, global structure, and textural variations, both through *data sharing* for machine learning based AI and through *knowledge sharing* for rule-based AI. We describe these AI methods and how the composer interacted with them: The personalization of the AI methods enabled the composer to explore new musical territories while keeping their original style, with AI music generation which “*sounds like it has been generated for him*”. The song “The last moment before you fly” ranked 3rd place in this contest, the jury underlining the “personal feel” of the song. We discuss here how these methods open the way to new co-creative approaches, both using AI and personalization.

1. INTRODUCTION

1.1 Music Composition and Co-creativity

Music assisted composition developed throughout the years, employing the new scientific advances: In the 1950s, automated compositions mainly used first-order Markov chains. Fred and Carolyn Attneave analyzed and generated Western-style melodies described as “convincing” [1], whereas Pinkerton generated melodies from the analysis of 39 children’s songs [2]. Probabilistic methods were expanded through several studies [3,4]. Methods based on n -grams were used to predict (using entropy measures) and generate melodies in the style of Gregorian music or in the style of Bach [5]. Models based on linguistics were also used, using sets of rules and/or grammars to describe composition processes [6]. Evolutionary computation methods for music generation have been developed for

the generation of melodies, drum tracks, or for live-coding applications [7]. Neural networks have been trained on Bach and traditional European melodies to extract stylistic regularities and to generate new melodies [8]. Today, most machine learning methods use neural networks and deep learning techniques are used extensively for music generation [9–11].

Herremans et al. did a thorough review of music generation techniques [12]. Most of them focus on the generation of harmony (sequences of chords) and of melody (sequences of notes). One of the key challenges in music generation is to be able to generate structured music with long-term correlations [12]. Several strategies have been employed in order to integrate some long term structure in melody generation, for instance by using a combination of short- and long-term models [13], structural constraints [14, 15], or, with Markov models and their extensions, by modeling a prior hierarchical knowledge [16, 17]. Another strategy is to generate elements of a *narrative* providing a feel of long-term structure. For instance, Hyperscore [18] provides a graphical interface for computer-assisted composition to edit and to visualize musical structures from low- to high-level features (melodic shape, harmonic tension, etc.), and Morpheus [19] uses a tension profile, designed by a user or computed from a template, as a constraint for music generation.

Another approach to music generation is to consider the AI models not as an independent agent, but rather as part of an interactive practice between human and machine, creating new dynamics of interactions through a *co-creative* process [20,21]. Co-improvisation systems such as DYCI2 [22] are based on real-time interactions between human musicians and a machine: Each performer (human or machine) listens to the music produced by the other and responds appropriately, bringing on the musical discourse in a novel way each time. Co-creativity can be used also in the context of a composition process either live through human-machine interactive systems [23, 24] or offline, in a *generate-then-select* fashion where the AI system generates a set of musical fragments, one of which will eventually be selected and possibly re-shaped by the composer to meet a specific musical need. For example, Ghisi proposes a method where a user chooses among fragments generated by a LSTM neural network [25], and Daren Banarsë composes folk music from melodies generated with a recurrent neural network [26]. Seeing the computer/AI system as a *colleague*, or even as a *co-creative colleague* who sug-

Copyright: © 2022 Ken Déguernel et al. This is an open-access article distributed under the terms of the [Creative Commons Attribution 3.0 Unported License](https://creativecommons.org/licenses/by/3.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

gest ideas, was studied by Lubart [27], Kantosalo and Jordanous [28], and Micchi et al. [29].

1.2 Motivations and Contributions

In this article, we present the co-creative music composition process that served as the artistic ground of our collaborative submission to the AI Song Contest 2021. Our team consisted of a Music Information Retrieval (MIR) researcher team (KD, MG, RG), part of which already participated in the AI Song Contest 2020 [29], with diverse knowledge and different AI methods, from music analysis to music generation, and a professional composer (SG) who is also a sound designer.

Co-creativity raises challenges in interaction, steerability, and personalization [30]. Our main goal was here both to use existing tools from the field of automatic music generation, and to develop new techniques to further experiment with a *personalized co-creative approach* by maintaining a constant interaction between composer and AI methods. For this composition, we focused on:

- an exploration of a co-creative expression, through a collaboration **involving people with distinct roles**. The collaboration is between a composer and a MIR team equipped with different knowledge and AI.
- a composition that goes beyond a simple “*generate-then-select*” approach. We wanted to **involve the composer in the creative process right at the beginning of each generative process**. The first compositional gesture for each component is given by the composer through knowledge sharing for knowledge-based AI, and through the creation of a training corpus using the own data of the composer (as well as their influence) for machine-learning AI.
- the **use of AI methods throughout the whole pre-production compositional process**, on multiple levels of the compositional process, including the “big picture” of the song. That is to generate chords and melodies, but also methods to explore co-creativity for generating long-term structures and variations of compositional processes (e.g. melodic or rhythmic variations, timbral curves, etc.).

In Section 2, we present the methods and the different models we used and developed for the different needs and interactions that the MIR team wanted with the composer. In Section 3, we present the results, open up a discussion about our thoughts (especially the ones of the composer), as well as the contests’ jury opinions, about the end results and the composition process.

2. METHOD

2.1 Approach

Generating structured music, especially involving long-term correlations between elements, is a key challenge [12, 31–34]. We decided to tackle this problem by focusing on a composition process with structural *sections*. First,

harmonic progressions were generated by the AI models and selected by the composer for several structural sections. Second, melodies were generated upon the selected chord progressions for each section. Then, song structures were generated, organizing the order and repetitions of the different sections. And finally, *arrangement layouts* were generated on the structure with propositions of variations of timbre, rhythm, and melody.

We combined three types of “intelligence”:

- the *human composer*, with their own compositional gesture, and music knowledge,
- *data-based AI* using machine learning for the generation of chords, and melodic contents,
- *rule-based AI* with some generative algorithms, for the generation of structures and variation templates.

Although the MIR team chose and developed the AI methods used, the interactions between the AI models and the composer were personalized through communication between them. Communication between the composer and the AI models was necessary at the initial step of each process (Figure 1). For the machine learning oriented AI, after initial knowledge sharing, that communication was *through data sharing*. For the generative algorithms, communication between the composer and the AI models was done *through knowledge sharing*, iterated and refined.

The methods presented here were developed during a 1-month timeframe, with a continuous discussion process among all the members of the team. The following paragraphs detail this composition process, focusing on personalization, for chords (Section 2.2), melody (2.3), structure (2.4), and finally variations (2.5). The music production is briefly explained in Section 2.6 but did not involve AI personalization.

2.2 Chords

Our first step was to create chord progressions for several sections of the composition. To personalize the generation, the first compositional gesture was done by the composer who selected a corpus of data mixing some of his own compositions as well as some of his influences (Ambient Electronica).

To generate new chord progressions, in the style of the given corpus, we used a factor oracle [35], a simple machine-learning AI which is at the core of the co-improvisation software developed at Ircam (OMax [36], SoMax [37], ImproteK [22], DYCI2 [38]). The factor oracle learns music content, here chord progressions, by connecting places in the memory which share a similar context. It can be considered as a variable order Markov model where each state has the maximum order possible in connection to the past, and can be used to generate new musical content whilst keeping the learned musical style. The main advantage of using a factor oracle for our approach is its capacity to generate music in a given style from a limited amount of data (compared to neural-network based methods). This is important for the personalization aspect of our approach as it would be difficult for a composer to




Figure 1. Workflow of our structure-based composition process and interactions between the composer and AI.




Figure 2. Factor oracle created from the chord progression D G D C G Bm C G D C D. Curved arrows at the upper half are used during the construction of the oracle, but are not used anymore. Suffix links (dashed arrows) connect each state to the past states sharing the longest common context. The two states connected by the bold suffix link share the common context G D C. A possible generation by navigating the oracle is D G D C * D . G, taking this suffix link towards the future (*), then, after the last state D, taking a suffix link towards the past. This generation was not present in the original chord progression but nevertheless share significant sub-progressions with it.

provide a substantial amount of data representative of the music they are trying to compose. Moreover, the main limitation of the generative power of the factor oracle, i.e. the lack of long-term structure, is alleviated by the fact that the global structure is composed later using another method (Section 2.4).

The composer provided a corpus of 21 chord sequences (5 from his own composition and 16 from his influences) totaling 498 chords, including 49 different chords. Chords qualities include major/minor/7th chords (major, minor and dominant), but also suspended chords and slash chords. Chords were not transposed. We trained a factor oracle on these chord sequences (Figure 2). Generation then followed the heuristics proposed in [36], i.e. the ability to jump both *forward and backward* in the memory, a *continuity factor* (3 chords must be generated before doing another jump in the memory), and a *taboo list* to avoid too many repetitions. The MIR team generated 50 new chord progressions of 16 bars each, and the composer created 5 sections, A, B, C, D1, and D2 from a selection of them (Figure 3).

The composer said:

“I started from the generated progressions that I truncated and split, selecting what pleased and surprised me. [...] Parts A and B are the beginning of the generations #37 and #38 that I truncated. The sequence from C to D1 is from the same generated chord sequences #41/#42 that I split into 2 parts (same for C and D2).



Figure 3. Chord progressions selected for each of the five sections. Red elements indicate further modifications chord made by the composer during the final touches of the composition process.

These parts sounded good in isolation but there was a missing link to go from B to C so I transposed C, D1 and D2 by a semitone to smooth the transition.”

The composer eventually expanded some chords to create a more modal feel.

2.3 Melodic content

Once the chord progressions for each section were selected, the next step was to generate melodies to go on top of them. Once again, to personalize the composition, the composer selected a corpus of 14 melodies for a total of 767 notes from his own music and inspirations. This corpus contained melodic and harmonic information.

A factor oracle was trained on this corpus, but this time the oracle was informed by interpolated probabilistic models [39] also trained on that same corpus. The factor oracle is created solely on melody, acting as a set of constraints on the probabilistic models which are representing the knowledge of the system about the melodic consistency, the chord-melody relationship and the balance between the two. The melodic consistency was modeled by a tri-gram: $P(M_t|M_{t-1}, M_{t-2})$ (where M_t is the note played at time t), and the chord-melody relationship was represented by a vertical model $P(M_t|C_t)$ (where C_t is the chord at time t). The two models were then interpolated with log-linear interpolation [40] and using




Figure 4. A probabilistic model linking melodies and chords is used to provide probability transitions towards each reachable states [39]. In our case, the factor oracle is built on melody, (μ_i corresponds to the melodic content in state i), whilst the environment is the chord progressions for each sections of the composition. At each step of the navigation, the transition probability for going for current state i to state j ($P(i \rightarrow j)$) is computed from the probabilistic models $P(M_t|M_{t-1}, M_{t-2})$ and $P(M_t|C_t)$, and then normalized according to the reachable states.

back-off smoothing [41]. Figure 4 shows a representation of one step in the navigation of a factor oracle informed by a probabilistic model. The probabilistic oracle proposed melodies for each of the chord progressions (taking into account melody and harmony).

“I used the oracle outputs as alternative suggestions for melodic direction. So I used the note sequences from the oracle and adapted them rhythmically to fit the song.”

Melodic contents have been used for the main melody but also to create background accompaniments for each section. In the end, not all selected melodies (either for the main melody or for the accompaniments) were actually used for each section in the composition: some sections working better with a single melody, or accompaniment parts appearing later creating some variational changes.

2.4 Structure

The third step of our composition process was to create a global structure for the piece, that is to decide how to combine the different sections with their harmonic and melodic elements. Candidate structures were generated by a first-order Markov model with additional constraints such as “do not output C before B already appeared”. The constraints are created and the transition probabilities of the model are adapted to better emulate the style and intention of the composer in order to personalize the results. They are decided through knowledge sharing during conversations with the composer, with a refinement process: some structures are generated, the composer then assessed the

quality of the generation, identify things that should occur more or less (or even not at all), new constraints are added and the probability of transitions are refined. For instance, the composer wanted to ensure that the initial A was repeated at least twice to “anchor the key” of the piece.

This process was repeated several times until most of the generated structures were satisfactory. The composer then selected one of them (Figure 5, top):

“I first selected structure 23 (that included D2 at the end), but, after three days of work, it did not work. I then tried to work with structure 24 – with another end – but finally, I decided to make it even more simple (looping on C), keeping a 100 BPM to fit the 4-minute song constraint. As with traditional composition, with AI, the eraser is the best tool!”

2.5 Variations

Many experiments in AI (co-)creativity generate melody and/or chords. However, the arrangement plan is a key feature of a song. To add some co-creativity in the arrangement of the piece, the composer proposed a simple taxonomy of variations on rhythm, melody, and timbre/texture from his own conceptualization of his composition process (Figure 6).

A generative procedure based on a random selection of variations, upon generative curves, generates where and how variations are to be applied. For each musical dimension identified in the composer’s taxonomy (rhythm, melody, timbre/texture), a curve is drawn out of a set of hand-crafted curves (decided through information sharing during discussions with the composer). The curve defines how many variations are being randomly selected out of the taxonomy of variations for each section of the piece. These curves were hand-crafted to have more variations towards the end of the song (Figure 7). Moreover, a dark/light curve was also generated by a similar procedure.

Twenty variations profiles were generated on structure 24s and the composer chose the variation profile 24s-02, and selected some of the variations inside this profile (Figure 5, bottom):

“I used the variations proposed by the AI to motivate me to give movement to the track while leaving me the possibility to improve the proposed solutions. Altogether, the structure and variation suggestion tools thus helped me to think outside the box, as for example, the timber/texture curve was not what I was thinking for.”

The composer chose how to implement each variation (which instruments and/or musical texture apply the different variations).

2.6 Final Touches, Music Arrangement, and Production

The composer created and produced the song given the harmony, the melody of each section, and the structure with

```

structure 23:   AAA B C A C  D1 B C D2 A
structure 24:   AA B C A C  D1 B C D1
==> final structure 24s:   A B C A CC D1 B CC

## Structure with variations 24s-02
(Timb-0 A)
Timb-0 A      (T-subst)
Timb-1 B
Timb-1 C
Timb-2 A      (*R-dense*,M-trans)           [Sixteenth pulse]
Timb-2 CC     (*M-irreg*,T-subst)           [Strings, C second time]
Timb-0 D1     (T-far)
Timb-0 B      (*M-ornem*,M-trans)           [Ornamented piano patterns]
Timb-1 C      (*R-dense*,M-ornem,M-trans)    [Sixteenth pulse]
Timb-2 (C)    (R-light,M-irreg,M-trans,*T-far*) [Fade-out]
    
```

Figure 5. Final structure (top) with variations (bottom). The system proposes a timbre for each section from the structure as well as a series of potential variations for them. The starred variations were the one selected by the composer and the right side shows how this variation is implemented in the composition.

Rhythm	R-dense	Densification
	R-light	Lightening
	R-frag	Conversion, fragmentation
Melody	M-trans	Diatonic transposition
	M-irreg	Irregular copy
	M-ornem	Ornamentation, appoggiaturas
Timbre, Texture	T-far	Distancing
	T-near	Rapprochement
	T-subst	Substitution
	T0..T1..T2	Dark..Light

Figure 6. Variation taxonomy. The initial vocabulary also contained harmonic variations, but the composer eventually decided that the chords had already enough such variation.




Figure 7. Variation curves. The left curve makes variations increase along the song. The right curve further pauses variations before the half of the song. Such curves were randomly selected, then stretched, to generate the variations on Figure 5.

variations presented above. Although some AI was also used during the music production part (after the composition), it did not involve AI personalization. The composer, being also sound designer and music producer, took great care in selecting/crafting instruments and textures in a minimalist “Ambient Electronica” style, combining synthesizer and drums with classical instruments like a cello and a string section. The main melody is played by a cello, while chords are presented with arpeggios played by a synthesizer. The densification and ornamentation are achieved by adding more complex accompaniment on strings or in the background, or by adding another layer of synthesizer playing the harmonic arpeggios with a new rhythm.

The mastering was done with the software LANDR that proposes online mastering services via an AI algorithm. Since the spectral shape of the song is not very common, AI based mastering fails to propose an adapted mastering




Figure 8. Interactive annotation of the song on Dezzrann (dezzrann.net). The top lines use the sections and variations features described on Figure 5. Selecting some labels, such as here “About the oracle”, provides more information.

equalization so the composer decided to apply a slight mastering equalization before LANDR [42].

3. RESULTS AND DISCUSSION

In this section, we focus on how AI impacted the composition process for the different musical dimensions and how the personalized aspect of the AI mattered.

3.1 Song Availability and Reception

The song “The Last Moment Before You Fly” is available online on various platforms. We also annotated the song, both for its global structure as well as for specific points describing some of the decisions taken during the compositional processes, on the Dezzrann platform¹ (Figure 8).

To remain as transparent as possible regarding the impact of the AI on the song, we released the code that we used as well as some outputs of the generative process that has been used for the composition, under GNU GPLv.3 and Creative Commons CC-BY-SA 4.0 respectively.² Note that, as requested by the composer, we are not releasing the training corpora for the chords and melodic content, as it contains copyrighted material and is considered to be

¹ Links available from www.algomus.fr/before-you-fly

² www.algomus.fr/code

part of the composer's identity.

The song was awarded 3rd place at the AI Song Contest 2021 out of 38 entries. Mark Simos, a member of the jury of this contest, songwriter, said in the jury session:

“What was really unique about that project is the way that the composer got a chance to train the data on him specific individual artistic work, so he felt like he were actually having a kind of conversation with him own sort of stylistic corpus, which gave it a really kind of personal feel.”

3.2 Low-tech AI and Co-creativity

Indeed, we tried to follow the composer's creative process as much as possible for the AI generation of the different musical elements, and also to respond to him needs and expectation from his first time creating a piece with an AI system:

“My goal for this project was to use AI algorithms to find new musical ideas and provoke inspiration while staying within a predefined style. We employed a co-creative approach in which the AI methods and the human fuel each other's creativity in an iterative process.”

Our models were purposefully “low-tech AI” (low need of computing power/data), to ease the exchange and enhance the place of the composer in the generative process (being able to generate musical elements with a tailored-made but a limited amount of data). This way, we experiment more on co-creativity while involving AI in parts of the composition process that are traditionally human-made (arrangement). Moreover, this choice also complements the *minimalist vibes* of the composition, that the composer wanted.

For all layers of the composition, the composer was the one doing the first compositional gestures through knowledge and/or data sharing, enabling a true personalization of the AI models to him needs, and was also doing the final compositional gestures and had the final say.

3.3 A Modular Co-creative Experience

We started by creating blocks of chord progressions and melodies in order to decide the atmosphere and direction for the piece. The blocks were then arranged in a generated structure. This generation was experimental, and the probability curves were hand-crafted. We would definitely like to train more flexible statistical models on actual data – but this would require having such structural data that is nonexistent for now. Altogether, chords, melody, structure, and variations were perceived as stimulation by the composer. Even though some elements for them felt something unusual, he found the generation inspiring and enabled them to explore ideas he would not have had without him interaction with the AI models:

“The first use of the AI tools was for inspiration and to begin the creation process. It allowed me to expand my

musical vocabulary and proposed some chord changes and melodic directions I would not have thought of without such assistance. When I was selecting the chord progressions there were a lot of options and it was difficult to choose. So, I decided to choose the progressions that moved me the most. I found that emotion that I wanted to capture and then I tried to make it grow through instrumentation and production. [...] I don't think I would have come up with these ideas without using these tools and I can say that this song is an artifact from my interaction with these algorithms.”

The MIR team was impressed by how fast the composer did appropriate the generative models with a clear insight on their inner working, which we think was possible by the “low-tech” nature of our methods. This facilitated a lot more in-depth discussion about the composition process, and helped with knowledge-sharing, in order to turn musical ideas into concrete computer model improvements and changes of model parameters.

Regarding the personalization of the AI, the methods using data for training presented here have a smaller generative power than some neural-network based methods from the state-of-the-art, such as [31, 43] which use much more data. However, the fact that factor oracle based method work on small corpora has an important impact on co-creativity. Creating a corpus large enough for neural network based methods would be much more difficult and time-consuming for a composer, and the overall content of the corpus would have to be more generic. Working on a smaller corpus is much more manageable for a composer and enables them to pinpoint more clearly the specific influences and style they want for their composition, which impacts directly their involvement in the human-AI interaction:

“Unlike the usual methods that use large generic corpora, the methods we employed used only my own data and influences to stay within a well-defined Ambient Electronica style. [...] Creating a corpus to train an AI with my own composition and influences makes it feel very personal. It was the first bricks of my own world that I would like to carry on contributing to by adding more stuff, piece by piece for future projects.”

With more time, more data would have been used, and the exploration of AI generation would have been longer by exploring more directions. The composer said:

“The result sounds less generic than some of the things I've heard from deep models. [...] It sounds like it has been generated for me.”

As said above, this aspect was also noted by members of the jury of the contest who were not part of the process.

3.4 Towards Autonomous Co-creation?

The method presented here for the generation of structure and variations shares similarities with MorpheuS [19]. In

both cases, profile curves are designed by the composer in order to guide the elements of the overall narrative of the composition. In our methodology, however, the composer define and describe their own vocabulary to describe their own approach to variation when composing. Therefore, the AI system is able to propose variations, communicating with the composer with their own taxonomy. On the one hand, this provides the composer with a better understanding of what the AI models are proposing. And on the other hand, this enables the composer to be more specific on which variational elements they want to consider in their composition. The autonomy is here shared between the machine and the human, who keeps such an autonomy “in the intentional sense” [44].

But could composers work by themselves, without any connection with a computer music team? A current limitation for the knowledge-based method is indeed the lack of a user interface, meaning that the MIR team was necessary to do the transition to interpret the needs of the composer and to implement them. Future work could include implementing such a user interface, in which the composer could describe their vocabulary, as well as a set of constraint and variation curves.

Here the MIR team and the composer worked in close interactions and the workflow evolved somewhat organically throughout the 1-month timeframe. More generally, we believe that this experience can be generalized and this workflow could be used again for a completely different project or with another composer. The different models used are agnostic of the data or of the knowledge, letting the composer express their intention and communicate with the AI models, through different data, and through new expressive vocabularies. However, we also felt that the interaction between humans, here the composer and the MIR team, was also central to the project, giving the composer a new perspective on their work while giving the MIR team new musical insights. As such, co-creation with AI also fosters creativity through new human relations.

4. REFERENCES

- [1] C. Ariza, “An open design for computer-aided algorithmic music composition : athenaCL,” Ph.D. dissertation, New York University, 2005.
- [2] R. C. Pinkerton, “Information theory and melody,” *Scientific American*, vol. 194, no. 2, pp. 77–86, 1956.
- [3] M. A. Hall and L. Smith, “A computer model of blues music and its evaluation,” *Journal of the Acoustical Society of America*, vol. 100, no. 2, pp. 1163–1167, 1996.
- [4] D. Conklin and J. Cleary, “Modelling and generating music using multiple viewpoints,” in *1st workshop on AI and Music*, 1988, pp. 125–137.
- [5] D. Conklin, “Music generation from statistical models,” in *AISB Symposium on Artificial Intelligence and Creativity in the Arts and Sciences*, 2003, pp. 30–35.
- [6] F. Lerdahl and R. Jackendoff, *A Generative Theory of Tonal Music*. MIT Press, 1983.
- [7] R. Loughran and M. O’Neill, “Evolutionary music: applying evolutionary computation to the art of creating music,” *Genetic Programming and Evolvable Machines*, vol. 21, no. 1, pp. 55–85, 2020.
- [8] M. C. Mozer, “Neural network music composition by prediction : exploring the benefits of psychoacoustic constraints and multi-scale processing,” *Connection Science*, vol. 6, no. 2–3, pp. 247–280, 1994.
- [9] G. Bickerman, S. Bosley, P. Swire, and R. M. Keller, “Learning to create jazz melodies using deep belief nets,” 2010, pp. 228–236.
- [10] L.-C. Yang, S.-Y. Chou, and Y.-H. Yang, “MidiNet: A convolutional generative adversarial network for symbolic-domain music generation,” in *International Society for Music Information Retrieval Conference (ISMIR 2017)*, 2017, pp. 324–331. [Online]. Available: <https://doi.org/10.5281/zenodo.1415990>
- [11] S. Ji, J. Luo, and X. Yang, “A comprehensive survey on deep music generation: Multi-level representations, algorithms, evaluations, and future directions,” *arXiv:2011.06801*, 2020.
- [12] D. Herremans, C.-H. Chuan, and E. Chew, “A functional taxonomy of music generation systems,” *ACM Computing Surveys*, vol. 50, no. 5, pp. 1–30, 2017.
- [13] M. Pearce, M. Ruiz, S. Kapasi, G. Wiggins, and J. Bhattacharya, “Unsupervised statistical learning underpins computational, behavioural and neural manifestations of musical expectations,” *NeuroImage*, vol. 50, no. 1, pp. 302–313, 2010.
- [14] A. Donzé, R. Valle, I. Akkaya, S. Libkind, S. A. Seshia, and D. Wessel, “Machine improvisation with formal specifications,” in *International Computer Music Conference (ICMC 2014)*, 2014, pp. 1277–1284.
- [15] T. Tanaka, B. Bemman, and D. Meredith, “Constraint programming formulation of the problem of generating miltan babbitt’s all-partition arrays,” in *22nd International Conference on Principles and Practice of Constraint Programming*, 2016.
- [16] K. Déguernel, E. Vincent, J. Nika, G. Assayag, and K. Smaïli, “Learning of hierarchical temporal structures for guided improvisation,” *Computer Music Journal*, vol. 43, no. 2, 2019.
- [17] N. Carvalho and G. Bernardes, “SyVMO: Synchronous variable markov oracle for modeling and predicting multi-part musical structures,” in *International Conference on Computational Intelligence in Music, Sound, Art and Design (EvoMUSART 2021)*, 2021.
- [18] M. Farbood, H. Kaufman, and K. Jennings, “Composing with hyperscore: An intuitive interface for visualizing musical structure,” in *International Computer Music Conference (ICMC 2007)*, 2007.

- [19] D. Herremans and E. Chew, "Morpheus: Automatic music generation with recurrent pattern constraints and tension profiles," in *IEEE TENCON*, 2016.
- [20] G. Assayag, "Creative symbolic interaction," in *Sound and Music Computing Conference (SMC 2014)*, 2014.
- [21] P. Esling and N. Devis, "Creativity in the era of artificial intelligence," *arXiv:2008.05959*, 2020.
- [22] J. Nika, M. Chemillier, and G. Assayag, "Improtek: Introducing scenarios into human-computer music improvisation," *Comput. Entertain.*, vol. 14, no. 2, Jan. 2017. [Online]. Available: <https://doi.org/10.1145/3022635>
- [23] J.-M. Fernández, T. Köppel, G. Lorieux, A. Vert, and P. Spiesser, "GeKiPe, a gesture-based interface for audiovisual performance," in *New Interfaces for Musical Expression Conference (NIME 2017)*, 2017, pp. 450–455.
- [24] A. Parmentier, K. Déguernel, and C. Frei, "A modular tool for automatic soundpainting query recognition and music generation in Max/MSP," in *Sound and Music Computing Conference (SMC 2021)*, 2021.
- [25] D. Ghisi, "Music across music: towards a corpus-based, interactive computer-aided composition," Ph.D. dissertation, University Paris 6, 2017.
- [26] O. Ben-Tal, M. Tobias Harris, and B. L. Sturm, "How music ai is useful: Engagements with composers, performers, and audiences," *Leonardo*, pp. 1–13, 2020.
- [27] T. Lubart, "How can computers be partners in the creative process: classification and commentary on the special issue," *International Journal of Human-Computer Studies*, vol. 63, no. 4-5, pp. 365–369, 2005.
- [28] A. Kantosalo and A. Jordanous, "Role-based perceptions of computer participants in human-computer co-creativity," in *AISB Symposium of Computational Creativity (CC@AISB 2020)*, 2020.
- [29] G. Micchi, L. Bigo, M. Giraud, R. Groult, and F. Levé, "I keep counting: An experiment in human/ai co-creative songwriting," *Transactions of the International Society for Music Information Retrieval*, vol. 4, no. 1, pp. 263–275, 2021.
- [30] C.-Z. A. Huang, H. V. Koops, E. Newton-Rex, M. Dinculescu, and C. J. Cai, "AI song contest: Human-AI co-creation in songwriting," in *International Society for Music Information Retrieval Conference (ISMIR 2020)*, 2020.
- [31] J.-P. Briot, G. Hadjeres, and F.-D. Pachet, *Deep learning techniques for music generation*. Springer, 2019, arXiv:1709.01620.
- [32] G. Medeot, S. Cherla, K. Kosta, M. McVicar, S. Abdalla, M. Selvi, E. Rex, and K. Webster, "StructureNet: Inducing structure in generated melodies," in *International Society for Music Information Retrieval Conference (ISMIR 2018)*, 2018.
- [33] Y. Zhou, W. Chu, S. Young, and X. Chen, "BandNet: A neural network-based, multi-instrument Beatles-style MIDI music composition machine," in *International Society for Music Information Retrieval Conference (ISMIR 2019)*, 2019.
- [34] P. Dhariwal, H. Jun, C. Payne, J. W. Kim, A. Radford, and I. Sutskever, "Jukebox: A generative model for music," 2020.
- [35] C. Allauzen, M. Crochemore, and M. Raffinot, "Factor oracle: A new structure for pattern matching," in *International conference on current trends in theory and practice of computer science*, 1999, pp. 295–310.
- [36] G. Assayag and G. Bloch, "Navigating the oracle: A heuristic approach," in *International Computer Music Conference (ICMC 2007)*, 2007, pp. 405–412.
- [37] J. Borg, "Somax 2: A real-time framework for human-machine improvisation," Ircam, Tech. Rep., 2019.
- [38] J. Nika, K. Déguernel, A. Chemla-Romeu-Santos, E. Vincent, and G. Assayag, "DYCI2 agents: merging the "free", "reactive", and "scenario-based" music generation paradigms," in *International Computer Music Conference (ICMC 2017)*, 2017.
- [39] K. Déguernel, E. Vincent, and G. Assayag, "Probabilistic factor oracles for multidimensional machine improvisation," *Computer Music Journal*, vol. 42, no. 2, 2018.
- [40] D. Klakow, "Log-linear interpolation of language models," in *5th International Conference on Spoken Language Processing*, 1998, pp. 1695–1698.
- [41] C. Zhai and J. Lafferty, "A study of smoothing methods for language models applied to information retrieval," *ACM Transactions of Information Systems*, vol. 22, no. 2, pp. 179–214, 2004.
- [42] J. Sterne and E. Razlogova, "Machine learning in context, or learning from landr, artificial intelligence and the platformization of music mastering," *Social Media + Society*, vol. 5, no. 2, 2019.
- [43] B. L. Sturm, O. Ben-Tal, U. Monaghan, N. Colling, D. Herremans, E. Chew, G. Hadjeres, E. Deruty, and F. Pachet, "Machine learning research that matters for music creation: A case study," *Journal of New Music Research*, vol. 48, no. 1, pp. 36–55, 2019.
- [44] J. McCormack, T. Gifford, and P. Hutchings, "Autonomy, authenticity, authorship and intention in computer generated art," in *International Conference on Computational Intelligence in Music, Sound, Art and Design (EvoMUSART 2019)*, 2019, pp. 35–50.