

Motifs, Phrases, and Beyond: The Modelling of Structure in Symbolic Music Generation

Keshav Bhandari¹[0009–0008–2045–6319] and Simon Colton²[0000–0002–4887–6947]

¹ k.bhandari@qmul.ac.uk

² s.colton@qmul.ac.uk

School of Electronic Engineering and Computer Science
Queen Mary University of London

Abstract. Modelling musical structure is vital yet challenging for artificial intelligence systems that generate symbolic music compositions. This literature review dissects the evolution of techniques for incorporating coherent structure, from symbolic approaches to foundational and transformative deep learning methods that harness the power of computation and data across a wide variety of training paradigms. In the later stages, we review an emerging technique which we refer to as “sub-task decomposition” that involves decomposing music generation into separate high-level structural planning and content creation stages. Such systems incorporate some form of musical knowledge or neuro-symbolic methods by extracting melodic skeletons or structural templates to guide the generation. Progress is evident in capturing motifs and repetitions across all three eras reviewed, yet modelling the nuanced development of themes across extended compositions in the style of human composers remains difficult. We outline several key future directions to realize the synergistic benefits of combining approaches from all eras examined.

Keywords: Generative Music · Deep Learning · Computerized Music

1 Introduction

Musical structure embodies the thoughtful organization of fundamental elements such as pitch, harmony, rhythm, and timbre in a composition [50]. Structure gives a sense of coherence, unity, and direction to the piece. On a basic level, notes are organized into motives and phrases. These phrases are then combined into higher-level sections such as verses, choruses and melodic segments. The repetition, variation, and development of musical ideas and themes manifest in the music and create relationships between these sections to form an overarching structure. Understanding and perceiving this musical structure is key to appreciating and enjoying music for listeners [24].

Musical structure is undoubtedly a crucial aspect of people’s cognitive process that influences their perception of music [39,59,60]. Surprisingly, research involving infants as young as four and a half years old suggests that musical experience might not be a prerequisite for perceiving musical phrase structures

[37]. Studies focusing on rhythm have delved into perceptual elements like timing and tempo [28] and syncopated patterns [22]. For human listeners, perceiving musical structure can be thought of as a pattern recognition task [65]. As we listen, we segment the music into motives, phrases, and sections. We identify similarities, repetitions, variations, and points of departure over time. Through this process, we form a mental representation of the music’s structure. However, human memory’s limitations become evident when individuals struggle to recall details; instead, they remember specific impactful moments [50] and distinct fragments rather than entire compositions, highlighting the selective nature of musical memory [48,65].

Given this understanding, songwriters strategically incorporate catchy repetitive phrases to enhance memorability and emotional impact for listeners [79,43]. However, Burns [6] argues that repetition gains meaning through its relationship with change. Both endless repetition and constant variation could lead to monotony. Thus, repetition and change are opposite possibilities from moment to moment in music. The tension between them can be a source of meaning and emotion. For example, in pop music, the repetition of verse and chorus sections helps emphasize music ideas, while the contrast between verse and chorus can create more emotional intensity [44]. Similarly, within Western tonal music, artists leverage repetition of harmonic progressions (sequences of chords) to guide listeners through a journey that creates dramatic narratives, conveying a sense of conflict that demands a resolution [39,62].

While structure and predictability are key for many applications involving composition, cognitive science and musical therapy [69], some contexts like video game and social media music require adaptive, event-driven approaches [2]. The music in video games often evolves in response to the player’s progress, creating a unique journey for each player. In this scenario, long-term structural plans around musical form might be impractical, as the musical narrative needs to be flexible and responsive to the unfolding events in the game. Similarly, social media music quickly sets moods with catchy hooks rather than intricate structures [23]. Additionally, avant-garde music is another area that deliberately challenges conventional musical structures, requiring artists to explore abstract and unpredictable forms that push listener expectations.

From a computational perspective, automatically analyzing and generating musical structure poses challenges for music research. Unlike other art forms like stories which have explicit plot structures, the “language” of music structure is more abstract and relies heavily on repetition, variation, and development of themes. This makes explicitly modelling musical structure difficult. While well-structured music generation with AI spans both raw audio and symbolic (sheet) music domains, this literature review specifically focuses on research pertaining to musical structures in the symbolic domain. While there has been a plethora of surveys on music generation systems, recent surveys such as [5,7,12,35,42,61] have covered works that use deep learning techniques. Among these, [7] touches upon early deep learning methods built upon recurrent neural networks (RNNs), and generative adversarial networks (GANs), while placing greater emphasis on

symbolic methods. Similarly, [5] also surveys early deep learning models until Music Transformer [33]. Furthermore, [42] and [61] have a very narrow focus on multi-track music using GANs and three specific deep learning models respectively. While [35] is comprehensive in their selection of papers by area of application, our literature review differs in two ways. First, we offer a deeper analysis on structure by exclusively focusing on studies that aim to improve musical structure in their methodology. Our selection of papers predominantly spans between 2007 to 2023. Early papers spanning this time frame can be seen in section 2 which provides a foundational overview by covering symbolic methods before delving into later advancements using deep learning in section 3. Second, we examine an emerging “subtask decomposition” approach in section 4 that breaks down the music generative process into smaller steps incorporating neuro-symbolic methods or musical knowledge. Reflecting on this trend, we address future directions for modelling long term musical structure in section 5.

2 Symbolic Methods

In early work on computerized music generation with long-term structure, [16] proposes a reinforcement learning approach on top of a Variable-Length Markov model for musical style imitation and improvisation with long term dependencies. Multiple agents model different musical attributes (e.g. pitch, rhythm, harmony) to capture the anticipatory foundations of musical expectation in a collaborative yet competitive framework. The chosen “behavior” agent for each episode takes actions and updates the others’ policies accordingly. Agents leverage *Factor Oracles*, which compactly gather repetitive sequence factors, enabling efficient access to long musical contexts. Results modeling Bach inventions demonstrates the model’s sensitivity to phrase boundaries, which clearly segment the score into formal sections.

The authors of [53] propose a constraint programming approach for controlled generation of Markov sequences, with applications to music. They formulate the Markovian property as a cost function represented by stacks of Elementary Markov Constraints and generate sequences by exploring the space of all sequences satisfying both the constraints and optimizing the Markovian cost. A key advantage of this approach is the ability to specify arbitrary structural constraints beyond just Markovian properties and find globally optimal solutions through search rather than greedy methods. These constraints can encode musical structures like repetition, motifs, endings, etc. Different Markovian cost functions can also bias generation towards more structured sequences. In addition, the model performs chunk-wise generation which allows it to maintain structure over long sequences (full songs) rather than just short clips. The chunks can be stitched to maintain longer-term dependencies.

Another notable example of a Markov approach for generating music with long-term repetitive and phrasal structure is the Racchmaninof-Jun2015 algorithm [13]. In this study, the authors develop a Markov model for inheriting long-term repetitive structure from template pieces using the SIACT pattern discov-

ery algorithm from [15]. The beat-relative-MIDI state space representation used in their paper is an improvement over a previous model called Racchmaninof-Oct2010 [14] that uses beat-spacing state space representation in which the spacing difference with respect to the previous note could consist of tonally obscure representations. Using the extracted structural framework of a template piece, the Markov model can generate complete musical textures in a forward as well as backward generative way to form phrases that are perceived as having a beginning, middle and an end.

Herremans et al., [27] present an application of Markov models combined with Variable Neighbourhood Search (VNS) to generate structured music in the style of Bagana, a traditional lyre from Ethiopia. The structure of Bagana pieces are encoded as cyclic patterns capturing repetitive cycles and global form. These patterns are realized by the VNS optimization technique that generates candidate melodies fitting the predefined structure of a template piece. A first order Markov model learned from the Bagana corpus is used in the next step to evaluate how well the melody matches the corpus statistics and musical structure. The VNS then selects the melodies optimizing the objectives to improve over iterations. While the Markov model does not directly generate notes, it provides the statistical basis for the objective functions guiding the VNS to produce melodies adhering to the structure with stylistic consistency. The integration of optimization and statistical learning helps in structuring music in Bagana's style.

Morpheus [26] extends the VNS approach to polyphonic music, with automatic detection of more complex long term patterns in the template piece with respect to a user-specified tension profile. Morpheus comprises 3 components that are core to its framework. First, the authors use the COSIATEC and SIATECComPress [47] pattern detection algorithms to extract repeated note patterns from a template piece. Users can choose between COSIATEC, which captures each note in precisely one pattern, or SIATECCompress, which captures more relationships between different notes, resulting in overlapping patterns. This allows imposing long-term structure in the generated music. Second, these patterns are then constrained in a VNS optimization framework that assigns pitches to generate new music matching a target tension profile. And finally, tension used in the VNS optimization algorithm is quantified using the spiral array model of tonality developed in [25]. While Morpheus has received favorable feedback overall, certain limitations are evident from a structural perspective. These include limited flexibility in pattern detection and the absence of captured variations or transformations, attributed to the pattern recognition algorithms used. Additionally, Morpheus is restricted in developing motifs, primarily repeating patterns without sophisticated development, and lacks a formal quantitative evaluation of structure retention, relying solely on informal listening tests for assessment.

Genetic algorithms, a class of evolutionary computing approaches that are inspired by the principles of natural selection, have also been explored in the concept of evolving musical motifs and structures over generations. For example, in [63], the authors propose an automatic music composition system called Phrase Imitation-based Evolutionary Composition (PIEC) that generates new melodies

with structure and form by imitating phrases from a sample melody composed by people. The framework uses a genetic algorithm along with intraphrase and interphrase rearrangement to mimic the ascending or descending melodic progression and note distribution of phrases in the sample melody. Similarly, in [1], the authors use a genetic algorithm to generate variations on a musical piece. To ensure that the generated variations have the underlying melodic skeleton of the original piece, the paper derives a melodic similarity fitness function from the spiral array model proposed in [9].

3 Deep Learning Methods

Symbolic methods, though useful in their versatility, are limited in their development across diverse musical datasets because they rely heavily on specific templates and encode strict musical rules, potentially restricting the generated music. In contrast, deep learning breaks free from these constraints. By learning from extensive datasets, deep learning models understand complex patterns and functions, and exploit locality through the use inductive biases. This enables them to contribute to improved quality in applications such as prompt based conditional generation, unconditional generation with no inputs or assistance, musical infilling as well as accompaniment generation [45].

3.1 Foundational Deep Learning Techniques

Lookback RNN and Attention RNN [66] by Google Magenta were two notable approaches early in the deep learning era that modeled long term structure in the generated melodies. Lookback RNN encoded repeating patterns from 1 or 2 bars ago in the input to allow the RNN model to identify repeating events better in the generated melody. Attention RNN used an earlier version of the self-attention mechanism from [3] to generate melodies with longer dependencies without storing information in the LSTM cell’s state. However, while these improvements are noticeable over brief periods, they may not necessarily impact the higher-level structural aspects of the piece.

Building on the Lookback RNN model, the authors of StructureNet [46] propose an enhancement at the generation stage of an RNN based melody generation model that induces repetition and structure in the generated melodies. StructureNet is trained on sequences of structural elements like repeats extracted from a dataset of melodies. When generating new melodies, it runs alongside the melody RNN model, biasing it towards notes that would form repeats similar to those seen in the training data. To evaluate if StructureNet improves structure, the authors compared statistics related to repeats as well as pitch and rhythm distributions between melodies generated with and without StructureNet. The results showed that StructureNet increased the occurrence and lengths of repeats, matching the melodies in the training data more closely, while preserving pitch and rhythm characteristics.

In another study, [8], the authors proposed using WaveNet [51], a generative model based on dilated convolutional neural networks, for generating new melodies that fit a given chord progression. They found WaveNet was able to better learn musical structure compared to its LSTM counterpart. This is because the WaveNet model encodes melodic structure more explicitly through its dilated convolutions across multiple time scales. Each layer doubles the receptive field of the previous layer, allowing it to learn dependencies over larger spans of melody. In contrast, an LSTM relies solely on its internal memory to capture structure, which is less explicit. As a result, WaveNet was able to generate coherent rhythms and melodies over many time steps that exhibited more repetitive patterns and long-range dependencies over the LSTM baseline that was reflective of compositional structure.

3.2 Transformative Deep Learning Approaches

Along the lines of model based enhancements for inducing structure in longer compositions is Music Transformer [33], which uses the transformer model equipped with the relative attention mechanism. This allows the model to learn patterns based on relative distances between events, making it easier to capture motifs and repetition on multiple timescales. However, this approach doesn't explicitly capture abstract structural elements such as how musical phrases and sections transition and develop over time. Additionally, when generating longer music pieces (over a minute), the model is not able to adhere to the initial prompt sequence and produces output with poor musicality and global structure.

In [4], the authors introduced Long-Short Term Universal Transformer (LSTUT), a fusion of transformers and RNNs, aiming to capture both local patterns and extensive musical relationships. This approach stemmed from a fundamental hypothesis: the intertwining of recurrence for short-term structure and attention for long-term dependencies enhances inductive biases, enabling the modelling of music spanning several minutes. However, while their methodology is promising, the evaluation is somewhat limited. Their assessment primarily relies on cross-entropy loss and manual inspection of attention heads to decipher the learned musical features. The lack of a subjective test and absence of generated samples restricts a comprehensive understanding of the model's potential.

Along the lines of feature representation based enhancements is the Pop music transformer [34] in which the authors propose REMI (revamped MIDI derived events) to represent MIDI data following the way humans read them. REMI uses position and bar events to embed an explicit beat/bar grid, tempo events to allow flexible tempo changes, and chord events to represent harmony. This equips models with a sense of rhythm and harmonic structure to better learn musical dependencies. Pop Music Transformer uses REMI with a Transformer-XL backbone. Objective evaluation shows the model generates piano compositions with more consistent rhythm and clearer downbeats compared to baselines that use the MIDI-like representation [52]. Subjective listening tests also indicate REMI results in more pleasing and coherent continuations of given musical prompts. While REMI offers a simple yet effective way to inject musical inductive biases

into sequence models through data representation, the input sequence length can be a point of concern for music of longer duration.

The Theme Transformer model [58] effectively tackles the constraints of Music Transformer by incorporating thematic materials into the generated outputs. Initially, the authors employ an unsupervised approach, combining contrastive learning with clustering techniques to extract recurring thematic patterns from 2-bar melodic segments. These themes are then integrated into a transformer encoder, and music generation occurs through a transformer decoder, which cross-attends to the encoder. The architecture of the Theme Transformer incorporates gated parallel attention and theme-aligned positional encoding. This design ensures the sustained influence of the thematic condition over extended generation periods. This is achieved by modulating between the self-attention mechanism based on the decoder’s auto-regressive inputs and the cross-attention from the encoder’s thematic outputs. Evaluations demonstrate that the Theme Transformer excels in generating music with repeated thematic elements compared to prompt-based Transformers. However, it faces challenges in creating variations and evolving the recurring themes, showcasing areas for potential improvement.

In [32], a novel R-Transformer model is introduced, aiming to tackle the under-explored problem of modelling varied motif repetitions. To facilitate this, a new music repetition dataset comprising over 500,000 labeled motif repetitions across 5 types is constructed from an existing piano music dataset. R-Transformer combines a Transformer encoder for note representation learning and a repetition-aware learner that exploits repetition characteristics based on music theory, enabling controlled generation of designated repetitions. The model is trained using both reconstruction and classification losses. Evaluations show it outperforms baseline models in generating varied high-quality repetitions of motifs and creates subjectively more enjoyable music than prior models, as rated by both musicians and non-musicians. This is attributed to explicitly modelling repetition types and their combinations, which prior generative models overlook.

3.3 Hierarchical Neural Networks

There exists a class of models that draw inspiration from the hierarchical nature of musical elements like notes, chords, bars, phrases, and movements. These models leverage this inherent hierarchy by employing specialized network structures to dissect and generate intricate melodic sequences.

MeloNet [29] is an early example of a small feed-forward based neural network architecture that uses a hierarchical approach to generate melodic variations in Baroque style given a melody input. Specifically, it employs two neural networks operating at different timescales - a “supernet” that predicts motif classes at a higher level based on the melody and previous motifs, and a “subnet” that generates the actual notes based on the motif, harmony, and previous notes.

Similarly, [70] introduces a hierarchical RNN model for monophonic melody generation, employing three Long-Short-Term-Memory (LSTM) sub-networks that operate in a sequential manner: generating bar profiles, beat profiles, and notes. The outputs of higher-level sub-networks guide the generation of finer

melody components in lower-level sub-networks. The subjective experiments conducted using this hierarchical approach demonstrate improved melodic quality and long term structure over MidiNet [74] and MusicVAE [55].

Expanding on this concept, Guo et al. [79] employed a three-tiered hierarchical RNN to simultaneously model rhythmic and pitch structures. This approach allowed the model to capture hierarchical musical relationships across varying time scales, evident in the presence of more repeated patterns indicated by a higher compression ratio metric. MusicFrameworks [17] adopted a multi-step hierarchical process, breaking down melody generation into manageable subtasks guided by music frameworks. This method facilitated the use of simpler models trained with less data, enabling the creation of long-term structures, including repetitions, within full-length melodies.

Another notable example of a hierarchical transformer model is HAT (**H**armony-**A**ware **H**ierarchical **M**usic **T**ransformer) proposed in [77]. HAT leverages musical harmony to jointly model form and texture structure for pop music generation. HAT represents music as tokens with attributes like notes, chords and phrases, and uses Transformer blocks in a three tiered hierarchy to capture structure at different levels. Subjective as well as objective experiments on musical structural attributes show HAT generates pop music with improved coherence in form and stability in texture compared to prior methods. However, there is still room for performance improvement for longer chord progressions and ending sections.

3.4 Long Sequence / Efficiency Models

A class of models hypothesize generating longer sequences is key to improving musical quality and structure, a notion grounded in the observation that training data often surpasses the capabilities of standard full-attention vanilla transformer models. The Compound Word Transformer work [31] groups consecutive music tokens into “compound words” and processes them together for compactness. This allows their model to handle longer sequences to capture long term dependencies and induce structure. In a user study, the authors demonstrated improved structural outcomes during unconditional generation. Similarly, MuseNet [11] adopts the Sparse Transformer [10] architecture, expanding the context window to 4096 tokens. The aim is to generate music with extended structural coherence from a given prompt. Another notable advancement is Museformer [75], a Transformer model for long sequence symbolic music generation that captures musical structure through a fine and coarse-grained attention mechanism. It allows each token to directly attend to selected “structure-related” bars based on statistical analysis of common repetitive patterns in music, providing fine-grained attention for modelling local structures. All other non-structure-related bars are summarized via coarse-grained attention to retain necessary context while greatly reducing computation compared to full attention. This combined approach efficiently generates long, structured musical sequences by capturing repetitive patterns and variations through fine-grained attention, while still incorporating global context. However, one limitation is that the bar selection

is based on statistics and may not generalize perfectly to all music. The SymphonyNet [41] approach introduces a permutation invariant byte pair encoding (BPE) representation for symbolic symphony music generation. The BPE representation is proposed to preserve note structure through a compressed encoding scheme by preventing the overflow of long sequence multi-instrument tokens.

3.5 Pre-training Methods

Pre-training involves training a model on a large unlabeled dataset before fine-tuning it for a specific task, aiding the model in learning general representations applicable to downstream tasks. Pre-training can prove advantageous in symbolic music generation tasks due to the intricate, multi-dimensional structure of music, encompassing local and global patterns involving elements like pitch, rhythm and dynamics. Music’s high complexity makes it important to learn musical syntax and patterns from vast datasets, enhancing coherence during fine-tuning. However, compared to other domains like text, there are fewer large-scale structured symbolic music datasets available [45,57,73] for task specific problems. Utilizing unlabeled data through pre-training facilitates transfer learning [30], enabling models to leverage discovered patterns and structures when fine-tuning on limited task-specific datasets, maximizing the utility of available labeled data. A few studies [21,40,57,67,73] have explored pre-training in the symbolic music domain aimed at generation with better musical structure and overall coherence.

In the LakhNES project [21], the authors pre-trained a Transformer-XL model using language modelling on a large MIDI dataset. This was then fine-tuned on multi-instrumental music generation. The goal was to leverage the patterns learned during pre-training to generate more coherent long-form music. SongMASS [57] also used a masked language modelling approach to pre-train an encoder-decoder model. They specifically lengthened the masked spans during pre-training to capture longer-term repetitive structures like verses and choruses. This helped generate songs with improved global structure. MuseBERT [67] and MRBERT [40] adapted BERT [20], a language representation model for various downstream music generation tasks by pre-training on piano music segments. They proposed techniques to handle music’s polyphonic and multi-dimensional nature compared to text. Recently proposed MelodyGLM [73] introduces a multi-task pre-training framework to model both local patterns and long-term structure in melodies. It uses melodic n-grams and long span masking objectives tailored to music during pre-training. MelodyGLM also constructs a large-scale melody dataset to enable robust pre-training. Both objective and subjective evaluations demonstrate MelodyGLM’s ability to generate melodies with improved structure and coherence.

4 Subtask Decomposition

In the context of deep learning, a new method has emerged which involves decomposing the training and generation process into smaller, more manageable

steps (typically 2 stages). First, a high-level structure or plan for the music piece is created, outlining its organization and patterns. In the second stage, the actual musical content, such as notes and chords, is generated while adhering to this plan. This separation helps in addressing the challenges of modelling long-term structures and generating details locally. The plan ensures overall coherence, while the detailed music is filled in from the bottom-up, providing a holistic musical composition. There are two notable differences between pre-training and subtask decomposition methods. The former involves the use of the first pre-trained model in the fine-tuning stage whereas the latter involves an entirely separate model in the final stage. Second, the subtask decomposition framework explicitly models musical structure through repeated patterns or an outline of a melodic skeleton in its initial stage. Additionally, while subtasks could be thought of as a hierarchy, what separates them from the hierarchical models in section 3.3 is that the latter trains a single model in a hierarchical fashion. In contrast, subtask decomposition approaches may involve multiple training objectives.

An early design of this framework can be seen in [65] in which the authors use modular planning theory to generate tree like structures at multiple levels which are referred to as sentences, double phrases and phrases. Each level affects the next through object dependencies and constraints. These planning structures are then filled with musical material using a bottom-up Markov based generation technique. Their results show that automated planning can efficiently produce repetition and novelty in melodies, resembling human compositions.

In another study, [68] the authors focus on the problem of drum pattern generation, specifically drum sequences that are rhythmically and structurally compatible with a given melodic track. The proposed model uses a two stage generation framework. In the first stage, a variational autoencoder (VAE) generative adversarial network from [38] is trained to generate a self-similarity matrix (SSM) for the drum track given the SSM of the melodic track. This drum SSM captures the structural information of the music. In the second stage, another VAE model then generates the actual MIDI drum patterns conditioned on the drum SSM and using a bar selection mechanism to encourage self-repetition. Objective and subjective tests conducted against baselines yielded promising outcomes. Objective and subjective tests against baselines showed promising results, highlighting the model’s enhanced performance due to its incorporation of global song structure through the drum SSM.

Similarly, PopMNet [71] proposes a two stage process for the generation of pop melodies with well defined long term structure. It represents melody structure as a directed acyclic graph capturing repetitive and sequential relationships between bars. In the first stage, a convolutional GAN model is trained on this graph based structural representation to generate plausible melodic structures. In the second stage, an RNN model is then trained to generate the actual melody conditioned on the structure and chords from the first stage. Human evaluations reveal that melodies generated by PopMNet receive higher ratings in terms of enjoyability, humanness, smoothness, and structure compared to previous models such as AttentionRNN [66], LookbackRNN [66], MidiNet [74], and Music

Transformer [33]. The analysis shows that PopMNet melodies exhibit repetition patterns similar to real pop songs, a feature previous models struggled to capture. However, one limitation of the model is its inability to capture complex structural patterns beyond repetitions and sequences.

Melons [80] overcomes PopMNet’s limitations by utilizing a more complex multi-edge graph representation of musical structure, incorporating 8 types of bar-level relations compared to PopMNet’s focus on repetition and rhythm sequences. Additionally, Melons adopts a Transformer architecture for both structure and melody generation, an upgrade from PopMNet’s CNN and RNN models. In cases where no edges exist between bars, an unconditional transformer model generates melodies. Conversely, when relationships exist, a conditional transformer model is employed, considering the previous 8 bars and pairwise bar relationships as context. While Melons outperforms PopMNet in structure and overall musicality, there’s still a gap between generated melodies and human-created music. About 80% of the bars in the dataset used in the study cover the proposed relations, suggesting potential room for improvement by exploring additional structural relationships between bars.

In [18], the authors present a modular approach for generating stylistic pop music by imitating a single seed song. This approach uses purely statistical and rule-based methods to capture distinctive melodic, harmonic, rhythmic, and structural attributes from the seed song. A proposed structure alignment procedure helps adhere to the stylistic section lengths and repetition patterns of the seed song. Separate generative models produce melody lines, chord progressions, bass-lines, and overall song structure. The generated songs are evaluated using both objective similarity metrics and subjective listening tests. While results show the model can produce enjoyable, novel music that listeners recognize as similar to the seed song, the authors acknowledge the potential for enhancing musical quality by combining their knowledge-based method with advanced deep learning techniques. They note that much of the model’s success stems from making good sequential choices guided by music theory and statistics. However, more sophisticated deep learning sequence models may be able to make even better choices for the next note in a generated sequence. The challenge would be retaining the benefits of the knowledge-based structure and harmony modelling while leveraging the pattern learning strengths of deep networks.

WuYun [76] proposes a hierarchical two-stage skeleton-guided melody generation architecture that incorporates musical knowledge. It first extracts a melodic skeleton of structurally important notes using music theory concepts such as downbeats, accents and tonal tension. In the first stage, a Transformer-XL network is trained on the extracted melodic skeleton to generate new skeleton sequences. The second stage comprises a Transformer encoder-decoder model that generates the full melody conditioned on encoding the skeleton to guide the process. Experiments on the Wikifonia dataset and subsequent subjective evaluations show WuYun produces melodies with improved structure and musicality compared to prior models including Pop Music Transformer [34] and Melons [80] but still falls short with respect to the musical qualities of human compositions.

AccoMontage [78] is a two step neuro-symbolic system for generating piano accompaniments for complete songs given a lead sheet that comprises the melody and chord information of the complete song. AccoMontage combines rule-based optimization for high-level structure and deep learning for local coherence. The system first retrieves candidate accompaniment phrases from a database using dynamic programming to optimize for good transitions and phrase-level fit. It then re-harmonizes the retrieved phrases using neural style transfer to match the target chord progression of the template song. Subjective and objective evaluations show AccoMontage generates more coherent and well-structured accompaniments compared to pure learning-based and template-matching baselines. However, the dependence on manual phrase level annotations in the dataset places constraints in the applicability of the model to other datasets.

The TeleMelody model in [36] applies a two-stage neuro-symbolic approach for lyric-to-melody generation. The first stage trains a lyric-to-template model to generate the symbolic template from lyrics. The template design incorporates musical knowledge and rules revolving around musical elements such as tonality, chord progression, rhythm pattern, and cadence, providing a symbolic representation that bridges lyrics and melodies. The second stage then trains a template-to-melody model in a self-supervised manner by extracting these templates from existing melodies and training a transformer model to reconstruct the melodies from the templates. Evaluations show TeleMelody generates melodies with better structure and musical attributes than previous end-to-end models with less requirement on paired training data. This is possible due to the intermediate template which helps reduce overall difficulty and improve data efficiency.

MeloForm [44] also uses a neuro-symbolic approach to generate melodies with musical form control by combining expert systems and neural networks. MeloForm has two main components. First, an expert system generates synthetic melodies with a given musical form using handcrafted rules. It develops motifs into phrases and arranges phrases into sections with repetitions/variations based on the musical form. Second, a 4 layered transformer encoder-decoder network refines the synthetic melodies from the expert system to improve musical richness without changing the form. The Transformer is trained on the Lakh MIDI Dataset [54] using an iterative masked sequence-to-sequence approach. During refinement, it conditions on the adjacent phrase’s rhythm and harmony to predict the pitch of masked phrases while preserving overall structure. The expert system provides flexibility over musical form, while the neural network improves musical coherence. This allows MeloForm to support various musical forms such as verse-chorus, rondo, variational, sonata, etc. Subjective evaluations performed on MeloForm’s generated outputs show an improvement over Melons [80] in musical structure, thematic material, melodic richness and overall quality.

Compose & Embellish [72] introduces another two-stage linear transformer based framework to generate piano performances with lead sheets as the intermediate output. The authors extract the monophonic melodic line using the skyline algorithm from [64]. For structural information, the authors employ a structural analysis algorithm from [19] that utilizes edit similarity and the A^*

search to identify repetitive phrases within the composition. Notably, the first stage involves a pre-training strategy using a separate larger dataset (Lakh MIDI Dataset [54]) for learning new lead sheets, setting it apart from the other two stage models. The second stage fine-tunes the lead sheet model and trains the performance model using the Pop1K7 dataset [31]. Conditioning the performance model on interleaved one-bar segments from the pre-trained lead sheet model enables the generation of complete performance bars with contextual understanding from the latest lead sheet bar.

5 Future Directions

Exploring Structural Encoding: While many music generation systems have used bar level representations such as REMI [34,72] and REMI+ [56] for modelling music with structure, incorporating phrase level representations or meta data may be an interesting direction to explore. Recent work by [49] extends the bar-level REMI encoding to incorporate phrase-level and bar countdown features. Their model demonstrates the value of encoding pop musical structure at the phrase level, aligning more closely with human music perception. However, phrase-based modelling and repetitive motif detection remains an open challenge in algorithmic music generation and composition. One barrier is the lack of standard practices or implementations for extracting musical phrases and structural annotations across genres. Moving forward, the development of robust tools to systematically extract and evaluate musical phrases across corpora would provide the necessary components to assemble hierarchical phrase structures, thereby enabling systems to construct forms at the phrase level. Additional annotations demarcating phrase boundaries within datasets would also facilitate progress in this area.

Mastering Advanced Compositional Techniques: Developing sophisticated long-term structure requires moving beyond basic repetitions and successions. Current models such as [33] and [58] struggle with nuanced variations, hindering the conveyance of complex musical narratives. Advancements demand models to master musical development through compositional techniques such as fragmentation, inversion, augmentation, diminution, stretto, sequence, and modulation across various time frames. Learning how to manifest a motif across the piece in a more structural way by incorporating such melodic and rhythmic transformations may warrant for a similarity conscious generative framework in which the generator understands how similar the last generated phrase is to the initial motif, previous phrase and overall form and has the ability to go back to refine it. Building a pause-think-refine framework along similar lines to [26] would help research advance beyond sequential note-to-note generations that do not explicitly model higher-level motivic structures and transformations.

Integrated Neuro-Symbolic Approaches in Music Generation: Music generation with long term structure is a complex endeavor that requires the coordination of multiple components at the preprocessing stage. For example, preprocessing data to extract melodic/rhythmic/phrase level features into mean-

ingful musical representations provides the groundwork for modelling music, as seen in [44,58,72,76]. However, recent advances in music generation have not yet effectively integrated insights from symbolic methodologies and cognitive musicology to achieve long-term structural coherence. While musical information retrieval has actively researched structural analysis algorithms for decades - spanning tempo estimation, key and modulation detection, cadence identification, and phrase similarity quantification, etc. - modern neural approaches surprisingly overlook these methods, opting for simplistic data preprocessing. Furthermore, cognitive studies elucidating how humans internalize and anticipate musical structure are imperative for designing evaluation protocols. Yet current assessments of generative systems’ musicality, both objective and subjective, lack foundations in perceptual research. Effectively melding the methodological tools of symbolic analysis and lessons from music cognition remains an open challenge in paving the way for a truly unified neuro-symbolic approach.

Top Down Structure: Sub-task decomposition studies in section 4 such as [72,76,80] clearly demonstrate the viability of top-down hierarchical plan-then-generate frameworks for algorithmic composition. By separating high-level musical form construction from lower-level sequence generation, these two-stage systems aim to improve global coherence. While we expect this trend to continue, we also hope to see studies that explore alternatives to the existing task breakdown involving pre-training, multi-task learning, unsupervised and contrastive learning methods. What data, features, representations, learning strategies and model architectures the sub-tasks consist of and how many stages of hierarchy make the framework most effective are questions worth looking into.

6 Conclusion

This literature review has charted elements of the evolution of modelling musical structure in AI-powered symbolic composition, from early symbolic systems to contemporary deep learning techniques and subtask decomposition frameworks. While progress is evident in capturing motifs, progressions, and global forms, modelling nuanced development and variation of themes much like human compositions across extended periods remains challenging. Key future directions involve moving beyond rigid structural units, integrating neuro-symbolic approaches at various stages in the system, and further exploring top down structures to shape musical narratives. Continued progress across complementary AI techniques provides hope for achieving the creative development of themes reflective of human composers and musical culture, which may someday be attainable.

Acknowledgements

This work was supported by the UKRI and EPSRC under grant EP/S022694/1. We gratefully acknowledge the use of generative AI in drafting and revising certain sections of this literature review. Finally, we owe much appreciation to our reviewers for their insightful critiques which greatly strengthened the work.

References

1. Alvarado, F.H.C., Lee, W.H., Huang, Y.H., Chen, Y.S.: Melody similarity and tempo diversity as evolutionary factors for music variations by genetic algorithms. In: 11th International Conference on Computational Creativity (ICCC). pp. 251–254 (2020)
2. Amaral, G., Baffa, A., Briot, J.P., Feijó, B., Furtado, A.: An adaptive music generation architecture for games based on the deep learning transformer model. In: 2022 21st Brazilian Symposium on Computer Games and Digital Entertainment (SBGames). pp. 1–6. IEEE (2022)
3. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. arXiv preprint arXiv:1409.0473 (2014)
4. de Berardinis, J., Barrett, S., Cangelosi, A., Coutinho, E.: Modelling long-and short-term structure in symbolic music with attention and recurrence. In: Proceedings of The 2020 Joint Conference on AI Music Creativity. pp. 1–11 (2020)
5. Briot, J.P., Hadjeres, G., Pachet, F.D.: Deep learning techniques for music generation—a survey. arXiv preprint arXiv:1709.01620 (2017)
6. Burns, G.: A typology of ‘hooks’ in popular records. *Popular music* **6**(1), 1–20 (1987)
7. 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)
8. Chen, K., Zhang, W., Dubnov, S., Xia, G., Li, W.: The effect of explicit structure encoding of deep neural networks for symbolic music generation. In: 2019 International Workshop on Multilayer Music Representation and Processing (MMRP). pp. 77–84. IEEE (2019)
9. Chew, E., Chen, Y.C.: Real-time pitch spelling using the spiral array. *Computer Music Journal* **29**(2), 61–76 (2005)
10. Child, R., Gray, S., Radford, A., Sutskever, I.: Generating long sequences with sparse transformers. arXiv preprint arXiv:1904.10509 (2019)
11. Christine, P.: MuseNet, <https://openai.com/research/musenet>
12. Civit, M., Civit-Masot, J., Cuadrado, F., Escalona, M.J.: A systematic review of artificial intelligence-based music generation: Scope, applications, and future trends. *Expert Systems with Applications* p. 118190 (2022)
13. Collins, T., Laney, R.: Computer-generated stylistic compositions with long-term repetitive and phrasal structure. *Journal of Creative Music Systems* **1**(2) (2017)
14. Collins, T., Laney, R., Willis, A., Garthwaite, P.H.: Developing and evaluating computational models of musical style. *AI EDAM* **30**(1), 16–43 (2016)
15. Collins, T., Thurlow, J., Laney, R.C., Willis, A., Garthwaite, P.H.: A comparative evaluation of algorithms for discovering translational patterns in baroque keyboard works. In: Downie, J.S., Veltkamp, R.C. (eds.) Proceedings of the 11th International Society for Music Information Retrieval Conference, ISMIR. pp. 3–8 (2010), <http://ismir2010.ismir.net/proceedings/ismir2010-2.pdf>
16. Cont, A., Dubnov, S., Assayag, G.: Anticipatory model of musical style imitation using collaborative and competitive reinforcement learning. In: Workshop on Anticipatory Behavior in Adaptive Learning Systems. pp. 285–306. Springer (2006)
17. Dai, S., Jin, Z., Gomes, C., Dannenberg, R.B.: Controllable deep melody generation via hierarchical music structure representation. arXiv preprint arXiv:2109.00663 (2021)

18. Dai, S., Ma, X., Wang, Y., Dannenberg, R.B.: Personalised popular music generation using imitation and structure. *Journal of New Music Research* **51**(1), 69–85 (2022)
19. Dai, S., Zhang, H., Dannenberg, R.B.: Automatic analysis and influence of hierarchical structure on melody, rhythm and harmony in popular music. *arXiv preprint arXiv:2010.07518* (2020)
20. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
21. Donahue, C., Mao, H.H., Li, Y.E., Cottrell, G.W., McAuley, J.: Lakhnes: Improving multi-instrumental music generation with cross-domain pre-training. *arXiv preprint arXiv:1907.04868* (2019)
22. Fitch, W.T., Rosenfeld, A.J.: Perception and production of syncopated rhythms. *Music Perception* **25**(1), 43–58 (2007)
23. Frid, E., Gomes, C., Jin, Z.: Music creation by example. In: *Proceedings of the 2020 CHI conference on human factors in computing systems*. pp. 1–13 (2020)
24. Gabriellsson, A.: The relationship between musical structure and perceived expression (2014)
25. Herremans, D., Chew, E.: Tension ribbons: Quantifying and visualising tonal tension. In: Hoadley, R., Nash, C., Fober, D. (eds.) *Proceedings of the International Conference on Technologies for Music Notation and Representation – TENOR’16*. pp. 8–18
26. Herremans, D., Chew, E.: Morpheus: generating structured music with constrained patterns and tension. *IEEE Transactions on Affective Computing* **10**(4), 510–523 (2017)
27. Herremans, D., Weisser, S., Sörensen, K., Conklin, D.: Generating structured music for bagana using quality metrics based on markov models. *Expert Systems with Applications* **42**(21), 7424–7435 (2015)
28. Honing, H., et al.: Structure and interpretation of rhythm and timing. *Tijdschrift voor Muziektheorie* **7**(3), 227–232 (2002)
29. Hörnel, D.: Melonet i: Neural nets for inventing baroque-style chorale variations. *Advances in Neural Information Processing Systems* **10** (1997)
30. Hosna, A., Merry, E., Gyalmo, J., Alom, Z., Aung, Z., Azim, M.A.: Transfer learning: a friendly introduction. *Journal of Big Data* **9**(1), 102 (2022)
31. Hsiao, W.Y., Liu, J.Y., Yeh, Y.C., Yang, Y.H.: Compound word transformer: Learning to compose full-song music over dynamic directed hypergraphs. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 35, pp. 178–186 (2021)
32. Hu, Z., Ma, X., Liu, Y., Chen, G., Liu, Y.: The beauty of repetition in machine composition scenarios. In: *Proceedings of the 30th ACM International Conference on Multimedia*. pp. 1223–1231 (2022)
33. Huang, C.Z.A., Vaswani, A., Uszkoreit, J., Shazeer, N., Simon, I., Hawthorne, C., Dai, A.M., Hoffman, M.D., Dinculescu, M., Eck, D.: Music transformer. *arXiv preprint arXiv:1809.04281* (2018)
34. Huang, Y.S., Yang, Y.H.: Pop music transformer: Beat-based modeling and generation of expressive pop piano compositions. In: *Proceedings of the 28th ACM international conference on multimedia*. pp. 1180–1188 (2020)
35. Ji, S., Yang, X., Luo, J.: A survey on deep learning for symbolic music generation: Representations, algorithms, evaluations, and challenges. *ACM Computing Surveys* (2023)

36. Ju, Z., Lu, P., Tan, X., Wang, R., Zhang, C., Wu, S., Zhang, K., Li, X., Qin, T., Liu, T.Y.: Telemelody: Lyric-to-melody generation with a template-based two-stage method. arXiv preprint arXiv:2109.09617 (2021)
37. Krumhansl, C.L., Jusczyk, P.W.: Infants' perception of phrase structure in music. *Psychological science* **1**(1), 70–73 (1990)
38. Larsen, A.B.L., Sønderby, S.K., Larochelle, H., Winther, O.: Autoencoding beyond pixels using a learned similarity metric. In: International conference on machine learning. pp. 1558–1566. PMLR (2016)
39. Lazzari, N., Poltronieri, A., Presutti, V.: Pitchclass2vec: Symbolic music structure segmentation with chord embeddings. arXiv preprint arXiv:2303.15306 (2023)
40. Li, S., Sung, Y.: Mrbert: Pre-training of melody and rhythm for automatic music generation. *Mathematics* **11**(4), 798 (2023)
41. Liu, J., Dong, Y., Cheng, Z., Zhang, X., Li, X., Yu, F., Sun, M.: Symphony generation with permutation invariant language model. arXiv preprint arXiv:2205.05448 (2022)
42. Liu, W.: Literature survey of multi-track music generation model based on generative confrontation network in intelligent composition. *The Journal of Supercomputing* **79**(6), 6560–6582 (2023)
43. Livingstone, S.R., Palmer, C., Schubert, E.: Emotional response to musical repetition **12**(3), 552–567. <https://doi.org/10.1037/a0023747>
44. Lu, P., Tan, X., Yu, B., Qin, T., Zhao, S., Liu, T.Y.: Meloform: Generating melody with musical form based on expert systems and neural networks. arXiv preprint arXiv:2208.14345 (2022)
45. Makris, D., Zixun, G., Kaliakatsos-Papakostas, M., Herremans, D.: Conditional drums generation using compound word representations. In: Martins, T., Rodríguez-Fernández, N., Rebelo, S.M. (eds.) *Artificial Intelligence in Music, Sound, Art and Design (EVOMUSART)*. pp. 179–194. Springer International Publishing, Cham (2022)
46. Medeot, G., Cherla, S., Kosta, K., McVicar, M., Abdallah, S., Selvi, M., Newton-Rex, E., Webster, K.: Structurenets: Inducing structure in generated melodies. In: International Society for Music Information Retrieval Conference (ISMIR). pp. 725–731 (2018)
47. Meredith, D.: Cosiatec and siateccompress: Pattern discovery by geometric compression. In: International Society for Music Information Retrieval Conference (ISMIR). ISMIR (2013)
48. Miller, G.A.: The magical number seven, plus or minus two: Some limits on our capacity for processing information **63**(2), 81–97. <https://doi.org/10.1037/h0043158>
49. Naruse, D., Takahata, T., Mukuta, Y., Harada, T.: Pop music generation with controllable phrase lengths. In: Proc. of the 23rd Int. Society for Music Information Retrieval Conf., Bengaluru, India (2022)
50. Ong, B.S., et al.: Structural analysis and segmentation of music signals. Citeseer (2006)
51. Oord, A.v.d., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., Kavukcuoglu, K.: Wavenet: A generative model for raw audio. arXiv preprint arXiv:1609.03499 (2016)
52. Oore, S., Simon, I., Dieleman, S., Eck, D., Simonyan, K.: This time with feeling: Learning expressive musical performance. *Neural Computing and Applications* **32**, 955–967 (2020)
53. Pachet, F., Roy, P.: Markov constraints: steerable generation of markov sequences. *Constraints* **16**, 148–172 (2011)

54. Raffel, C.: Learning-based methods for comparing sequences, with applications to audio-to-MIDI alignment and matching. <https://doi.org/10.7916/D8N58MHV>
55. Roberts, A., Engel, J., Raffel, C., Hawthorne, C., Eck, D.: A hierarchical latent vector model for learning long-term structure in music. In: International conference on machine learning. pp. 4364–4373. PMLR (2018)
56. von Rütte, D., Biggio, L., Kilcher, Y., Hofmann, T.: Figaro: Generating symbolic music with fine-grained artistic control. arXiv preprint arXiv:2201.10936 (2022)
57. Sheng, Z., Song, K., Tan, X., Ren, Y., Ye, W., Zhang, S., Qin, T.: Songmass: Automatic song writing with pre-training and alignment constraint. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 13798–13805 (2021)
58. Shih, Y.J., Wu, S.L., Zalkow, F., Muller, M., Yang, Y.H.: Theme transformer: Symbolic music generation with theme-conditioned transformer. IEEE Transactions on Multimedia (2022)
59. Stevens, C.J.: Music perception and cognition: A review of recent cross-cultural research. Topics in cognitive science 4(4), 653–667 (2012)
60. Tan, N., Aiello, R., Bever, T.G.: Harmonic structure as a determinant of melodic organization. Memory & Cognition 9(5), 533–539 (1981)
61. Tang, H., Gu, Y., Yang, X.: Music generation with ai technology: Is it possible? In: 2022 IEEE 5th International Conference on Electronics Technology (ICET). pp. 1265–1272. IEEE (2022)
62. Temperley, D.: The cognition of basic musical structures. MIT press (2004)
63. Ting, C.K., Wu, C.L., Liu, C.H.: A novel automatic composition system using evolutionary algorithm and phrase imitation. IEEE Systems Journal 11(3), 1284–1295 (2015)
64. Uitdenbogerd, A.L., Zobel, J.: Manipulation of music for melody matching. In: Proceedings of the sixth ACM International Conference on Multimedia. pp. 235–240 (1998)
65. Velardo, V., Vallati, M.: A planning-based approach for music composition. Springer (2015)
66. Waite, E., Eck, D., Roberts, A., Abolafia, D.: Generating long-term structure in songs and stories, <https://magenta.tensorflow.org/2016/07/15/lookback-rnn-attention-rnn>
67. Wang, Z., Xia, G.: Musebert: Pre-training music representation for music understanding and controllable generation. In: International Society for Music Information Retrieval Conference (ISMIR). pp. 722–729 (2021)
68. Wei, I.C., Wu, C.W., Su, L.: Generating structured drum pattern using variational autoencoder and self-similarity matrix. In: International Society for Music Information Retrieval Conference (ISMIR). pp. 847–854
69. Wigram, T., Gold, C.: Music therapy in the assessment and treatment of autistic spectrum disorder: clinical application and research evidence. Child: care, health and development 32(5), 535–542 (2006)
70. Wu, J., Hu, C., Wang, Y., Hu, X., Zhu, J.: A hierarchical recurrent neural network for symbolic melody generation. IEEE transactions on cybernetics 50(6), 2749–2757 (2019)
71. Wu, J., Liu, X., Hu, X., Zhu, J.: Popmnet: Generating structured pop music melodies using neural networks. Artificial Intelligence 286 (2020)
72. Wu, S.L., Yang, Y.H.: Compose & embellish: Well-structured piano performance generation via a two-stage approach, <http://arxiv.org/abs/2209.08212>
73. Wu, X., Huang, Z., Zhang, K., Yu, J., Tan, X., Zhang, T., Wang, Z., Sun, L.: Melodyglm: Multi-task pre-training for symbolic melody generation. arXiv preprint arXiv:2309.10738 (2023)

74. Yang, L.C., Chou, S.Y., Yang, Y.H.: Midinet: A convolutional generative adversarial network for symbolic-domain music generation. arXiv preprint arXiv:1703.10847 (2017)
75. Yu, B., Lu, P., Wang, R., Hu, W., Tan, X., Ye, W., Zhang, S., Qin, T., Liu, T.Y.: Museformer: Transformer with fine-and coarse-grained attention for music generation. *Advances in Neural Information Processing Systems* **35**, 1376–1388 (2022)
76. Zhang, K., Wu, X., Zhang, T., Huang, Z., Tan, X., Liang, Q., Wu, S., Sun, L.: Wuyun: Exploring hierarchical skeleton-guided melody generation using knowledge-enhanced deep learning. arXiv preprint arXiv:2301.04488 (2023)
77. Zhang, X., Zhang, J., Qiu, Y., Wang, L., Zhou, J.: Structure-enhanced pop music generation via harmony-aware learning. In: *Proceedings of the 30th ACM International Conference on Multimedia*. pp. 1204–1213 (2022)
78. Zhao, J., Xia, G.: Accomontage: Accompaniment arrangement via phrase selection and style transfer. arXiv preprint arXiv:2108.11213 (2021)
79. Zixun, G., Makris, D., Herremans, D.: Hierarchical recurrent neural networks for conditional melody generation with long-term structure. In: *2021 International Joint Conference on Neural Networks (IJCNN)*. pp. 1–8. IEEE (2021)
80. Zou, Y., Zou, P., Zhao, Y., Zhang, K., Zhang, R., Wang, X.: Melons: generating melody with long-term structure using transformers and structure graph. In: *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 191–195. IEEE (2022)