

deepGTTM-III: Simultaneous Learning of Grouping and Metrical Structures

Masatoshi Hamanaka¹ Keiji Hirata², and Satoshi Tojo³

¹ RIKEN,
masatoshi.hamanaka@riken.jp,

² Future University Hakodate,
hirata@fun.ac.jp,

³ JAIST,
tojo@jaist.ac.jp,

<http://gttm.jp/>

Abstract. This paper describes an analyzer that simultaneously learns the grouping and metrical structures on the basis of the generative theory of tonal music (GTTM) by using a deep learning technique. The GTTM is composed of four modules that are in series. The GTTM has feedback loop in which the former module uses the result of the latter module. However, each module is independent in previous GTTM analyzers, thus they do not form a feedback loop. For example, deepGTTM-I and deepGTTM-II independently learn the grouping and metrical structures by using a deep learning technique. In light of this, we present deepGTTM-III, a new analyzer that includes the concept of feedback that enables simultaneous learning of grouping and metrical structures by integrating both networks of deepGTTM-I and deepGTTM-II. Experimental results show that deepGTTM-III outperforms deepGTTM-I and deepGTTM-II.

Keywords: A generative theory of tonal music (GTTM), grouping structure, metrical structure, deep learning.

1 Introduction

Our goal is to develop a system that enables a time-span tree of a melody to be automatically acquired on the basis of the generative theory of tonal music (GTTM) [1]. GTTM is composed of four modules, each of which assigns a separate structural description to a listener's understanding of a piece of music. These four modules respectively output a grouping structure, metrical structure, time-span tree, and prolongational tree. The grouping structure is intended to formalize the intuitive belief that tonal music is organized into groups that are in turn composed of subgroups. These groups are presented graphically as several levels of arcs below a music staff. The metrical structure describes the rhythmical hierarchy of a piece of music by identifying the position of strong beats at the levels of a quarter note, half note, measure, two measures, four measures, and so on. Strong beats are illustrated as several levels of dots below the music staff (Fig. 1).




Fig. 1. Grouping structure, Metrical structure, time-span tree, and prolongational tree

The time-span tree provides performance rendering [2], music reproduction [3], and a summarization of the music [4]. This summarization can be used as a representation of a search, resulting in music retrieval systems. It can also be used for melody morphing, which generates an intermediate melody between two melodies in a systematic order [5, 6]. These systems presently need a time-span tree analyzed by musicologists because previous analyzers [7, 8] do not perform optimally.

There are three big problems when implementing GTTM on a computer.

- Conflict between rules.

There are two types of rules in GTTM: well-formedness rules (WFRs) and preference rules (PRs). WFRs are necessary conditions to assign a structure and restrictions on these structures. When more than one structure can satisfy the WFRs, PRs indicate the superiority of one structure over another.

Because there is no strict order for applying PRs, a conflict between rules often occurs when applying them, which results in ambiguities in analysis. Figure 2 shows an example of the conflict between metrical preference rules (MPRs) 5c and 5a. The MPR5c states that a relatively long slur results in a strong beat, and MPR5a states that a relatively long pitch-event results in a strong beat. Because metrical WRF 3 (MWFR3) states that strong beats are spaced either two or three beats apart, so a strong beat cannot be perceived at both onsets of the first and second notes.




Fig. 2. Conflict between two metrical preference rules

We developed an automatic time-span tree analyzer (ATTA) [7] that had 46 adjusted parameters to control the strength of each rule. In other words, the ATTA we developed enabled us to control the priority of rules, which enabled us to obtain extremely accurate groupings and metrical structures. However, we needed musical knowledge like that of musicologists to properly tune the parameters.

The full ATTA (FATTA) [8] does not have to tune the parameters because it automatically calculates the stability of structures and optimizes the parameters so that the structures are stable. FATTA obtains excellent analysis results for metrical structures but unacceptable results for grouping structures and time-span trees.

In our deepGTTM, the deep layered network enables us to learn the priority of rules.

- Difficult to integrate bottom up and top down processes

GTTM rules include bottom up and top down rules. For example, GPR2 is a bottom up rule that prescribes the relationship of onset (attack) and offset (release) timings and a grouping boundary. In contrast, GPR5 is a top down rule that prefers that a group be divided into two subgroups of the same length.

The ATTA and FATTA output frequently wrong higher level hierarchical structures even when low-level structure is correct because they only use bottom up a process.

In contrast, we also developed analyzers that only use a top down process called σ GTTM [9], σ GTTMII [10], and σ GTTMIII [11]. σ GTTM and σ GTTMII can detect the local grouping structure in GTTM analysis by combining the GTTM with statistical learning. However, σ GTTM and σ GTTMII are only suitable for grouping structures and cannot acquire time-span trees. σ GTTMIII enabled us to automatically analyze time-span trees by learning with a time-span tree of 300 pieces from the GTTM database [12] on the basis of probabilistic context-free grammar (PCFG). σ GTTMIII performed the best at acquiring time-span trees. However, these analyzers [7-11] do not perform sufficiently for use in application systems [2-6].

In our deepGTTM, the deep layered network learns both top down and bottom up rules from learning data.

- Feedback loops

The four modules in the GTTM are in series, that is, the latter module uses the result of the former module. The GTTM also has a feedback loop in which the former module uses the result of the latter module. For example, GPR7 (time-span and prolongational stability) prefers a grouping structure that results in a more stable time-span and/or prolongation reduction. For another example, MPR9 (time-span interaction) prefers a metrical analysis that minimizes conflict in the time-span reduction. However, each module is independent in previous GTTM analyzers, thus they do not form a feedback loop. For example, deepGTTM-I [13] and deepGTTM-II [14] independently learn the grouping and metrical structures by using a deep learning technique.

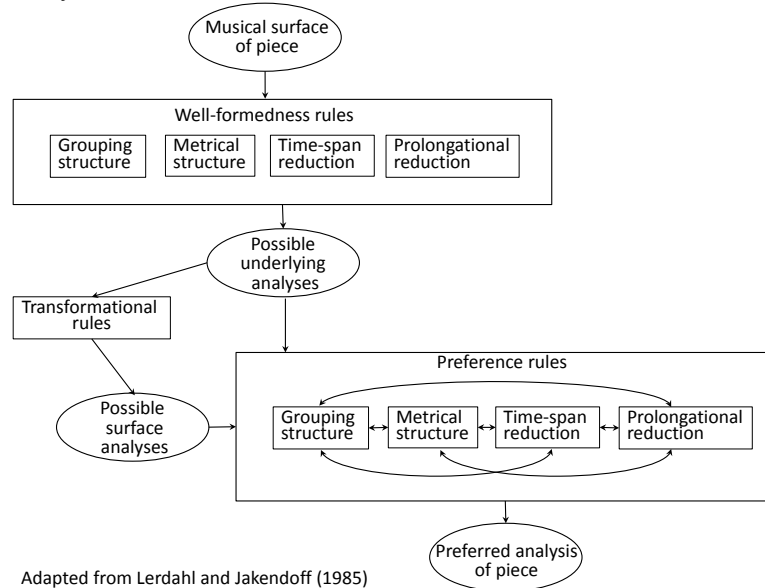
Figure 3 summarizes the theory of Lerdahl and Jakendoff [1]. The lower right shows the preference rules. If we naively implement this theory, the analysis process is endless looping when the output is divergence.

In light of this, we present deepGTTM-III, a new analyzer that solves the above problems and enables simultaneous learning of grouping and metrical structures by integrating both networks of deepGTTM-I and deepGTTM-II. The deep

layered network learns the priority of rules that solves the conflict between rules. The network also learns both bottom up and top down rules. By integrating both networks of deepGTTM-I and deepGTTM-II, the integrated network possesses information for acquiring the metrical structure and the grouping structure. Therefore, the information for acquiring the metrical structure can be used for acquiring the grouping structure, and vice versa. Therefore, a feedback loop is implicitly constructed inside the network.

The network is pre-trained by using 15,000 pieces of music formatted in musicXML acquired by web crawling. We use 300 pieces from the GTTM database: 200 for fine-tuning and 100 for evaluation [15]. Experimental results show that the integrated network outperforms the independent network.

The paper is organized as follows. Section 2 describes related work, and Section 3 explains our GTTM analyzers: deepGTTM-I, II, and III. Section 4 explains how we evaluated the performances of deepGTTM-I, II, and III, and Section 5 concludes with a summary and an overview of future work.



Adapted from Lerdahl and Jakendoff (1985)

Fig. 3. Summary of Lerdahl and Jakendoff's theory

2 Related Work

Deep learning has recently been used for tasks in the area of music information retrieval [15-19] and shows the potential to solve various kinds of tasks in the area. An automatic tagging system using a fully convolutional network was developed that predicts high-level information about a music clip, such as emotion, genre, and instrumentation [15]. An automatic chord detection system using bottleneck architecture of a deep layered network outperforms previous systems based on support vector machines (SVMs) and hidden Markov models (HMMs) [16]. An automatic chord estimation system based on a hybrid Gaussian HMM and deep learning approach enables very large chord progression to be estimated [17]. A music

recommendation system using deep convolutional neural networks to predict latent factors from music audio shows that deep convolutional neural networks significantly outperform more traditional approaches [18]. An automatic polyphonic music transcription using a supervised neural network model performed the best across the two most common unsupervised acoustic models [19]. These systems [15-19] replace other machine learning techniques with deep learning and performed better than traditional machine learning techniques.

The traditional machine learning techniques cannot work well in our task of acquiring hierarchical musical structure. In other words, only deep learning enables the relationship between inputting score and outputting structure to be learned. Direct learning between inputs to output does not work well because they have gaps that are too wide. Therefore, we made two steps for learning. First, the network learns each rule's application. The deep layered network can easily learn rules in the GTTM. After learning the rules' applications, the network can learn the relationship between inputting score and outputting structure. That is, the network gains musical knowledge by learning the GTTM rules.

3 deepGTTM-I, II, and III

deepGTTM-I, II, and III are GTTM analyzers based on deep learning. deepGTTM-I analyzes local grouping boundaries of a grouping structure [13], and deepGTTM-II analyzes the metrical structure [14]. In this paper, we present deepGTTM-III, which integrates deepGTTM-I and II.

There are three main advantages of using deep learning for GTTM analysis.

- Learning applications of both bottom up and top down rules
Previous analysis systems based on the GTTM were constructed by a human researcher or programmer. Some rules in the GTTM are very ambiguous, and their implementations might differ depending on the person. However, deepGTTM is a learning based system where the quality of the analyzer depends on the training data and trained network. To learn both bottom up and top down rules, the input of the network includes the score information of the whole analysis area.
- Learning priority of rules
 σ GTTM and σ GTTMII do not work well because they only determine the priority of rules from applied rules because the priority of rules depends on the context of a piece. The input of the network in deepGTTM, on the other hand, is the score, and the network learns the priority of the rules as the weight and bias of the network on the basis of the context of the score.
- Feedback loop in deep layered network
There are several kinds of feedback process in deepGTTM-III because a deep layered network is trained by multi-task learning with grouping and metrical structures. When the grouping structure is learned, the important information for acquiring grouping and metrical structures is propagated because the deepGTTM-III shares hidden nodes for acquiring both grouping and metrical structures. Similarly, when the metrical structure is learned, the important information for acquiring grouping structure is also propagated.

3.1 Structure of Network

We used a deep belief network (DBN) for deepGTTM-I, II, and III. Figure 4 outlines the structure for the DBN of deepGTTM-I. The input of the DBN was the onset time, offset time, pitch, and velocity of note sequences from musicXML. All inputs are normalized from 0 to 1. The output of DBN formed multi-tasking learning, which had 10 outputs: 9 kinds of grouping preference rules (GPR2a, 2b, 3a, 3b, 3c, 4, 5, 6, and 7) and local grouping.

Figure 5 outlines the structure of deepGTTM-II. The inputs of deepGTTM-II are the onset time, offset time, pitch, and velocity, and grouping structure manually analyzed by musicologists. Each hierarchical level of the grouping structure is separately inputted by a note neighboring the grouping boundary as 1; otherwise, 0. There are eight outputs of deepGTTM-II that enable multi-tasking learning in each hierarchical level of the metrical structure, i.e., seven MPRs (MPR2, 3, 4, 5a, 5b, 5c, and 5d), and one level of the metrical structure. Individual outputs have two units, e.g., rules that were not applicable (=0) and rules that were applicable (=1), or weak beats (=0) and strong beats (=1). A metrical structure consists of hierarchical levels, and we added one hidden layer to generate the next structure level. We used logistic regression to connect the final hidden layer ($n, n+1, \dots, n+h$) and outputs. All outputs shared the hidden layers from 1 to the final hidden layer.

Figure 6 outlines the structure of the DBN we call deepGTTM-III to generate a grouping and metrical structure. deepGTTM-III has the same input as deepGTTM-I, and its output is the same as the merged output of deepGTTM-I and II.

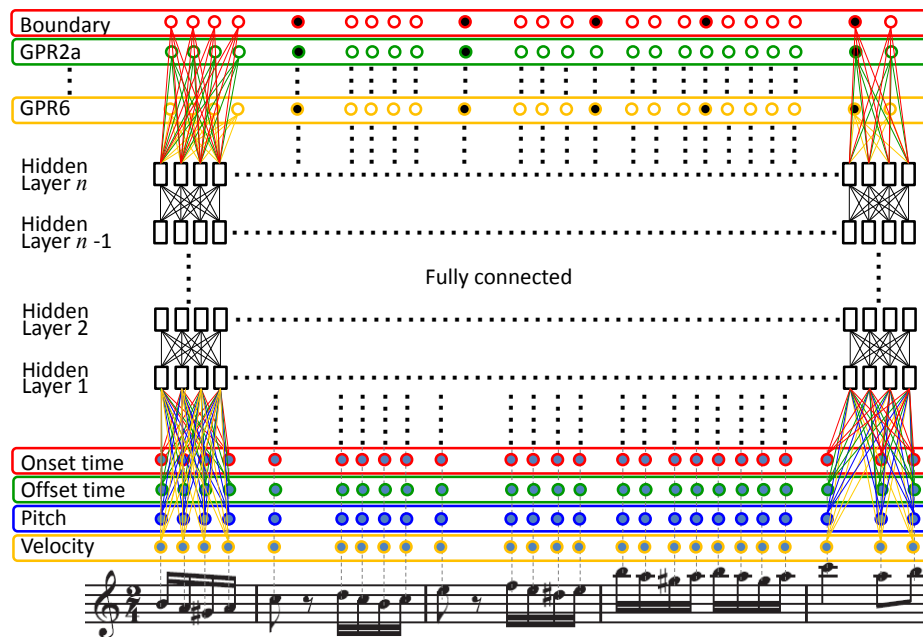


Fig. 4 DBN for deepGTTM-I.

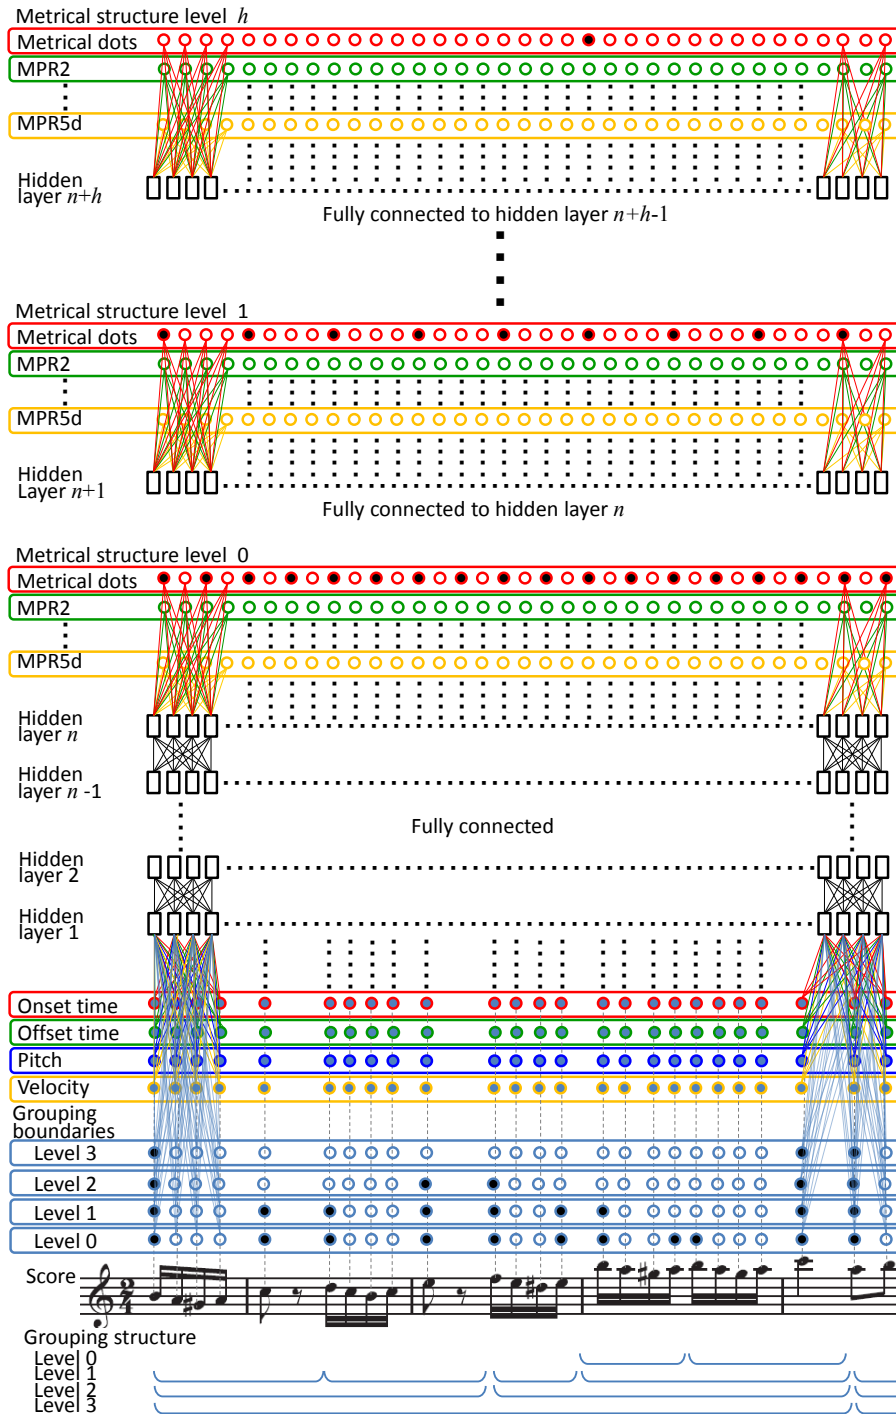


Fig. 5 DBN for deepGTTM-II.

3.2 Learning Networks

This section describes how we learned the local grouping boundaries and metrical structure by using deep layered networks.

Pre-training. For pre-training, the network learned the features of the music. A large scale dataset with no labels was needed, so we collected, 15,000 pieces of music formatted in musicXML from Web pages that were linked to on the musicXML page of MakeMusic Inc. [20]. The musicXMLs were downloaded in three steps.

- 1) A Web autopilot script made a list of URLs that were most likely files of musicXMLs from five links on the musicXML page of MakeMusic Inc.
- 2) The files in the URL list were downloaded after the URLs that were clearly not musicXMLs had been omitted.
- 3) All the downloaded files were opened using the script, and files that were not musicXML were deleted.

By using a restricted Boltzmann machine, each network of deepGTTM-I, II, and III was pre-trained.

Learning rules application and structure. The network in a fine-tuning learned with the labeled dataset. We had 300 pieces with a labeled dataset in the GTTM database, which included musicXMLs with positions of local grouping boundaries, positions of dots of each hierarchy of the metrical structure, and positions to which the grouping and metrical preference rules were applied. However, these 300 pieces were insufficient for deep learning.

Consequently, we constructed a half-labeled dataset. We automatically added the labels of six applied rules of GPR2a, 2b, 3a, 3b, 3c, and 3d, and MPR3, 5a, 5b, 5c, 5d because these rules could be uniquely applied as a score. We used our ATTA to add labels to these rules.

We also artificially increased the labeled dataset because the 300 pieces in the GTTM database were insufficient for training a deep layered network. First, we transposed the pieces for all 12 keys. Then, we changed the length of note values to two times, four times, eight times, a half time, a quarter time, and an eighth time. Thus, the total labeled dataset had 25,200 (= 300x12x7) pieces.

The priority of rules and grouping and metrical structures are learned by back propagation of the deep layered network using the half-labeled dataset and labeled dataset. deepGTTM-I and II have very complex networks. The fine-tuning of local grouping boundaries and one level of the metrical structure involves multi-task learning. The fine-tuning of each PR also involved multi-task learning. Therefore, the fine-tuning of PRs involves multi-dimensional multi-task learning. The processing flow for the learning of a GPR or local grouping boundaries has four steps. The order of music pieces was changed at every epoch in all steps.

- 1) The order of the pieces of training data is randomly shuffled, and a piece is selected from top to bottom.
- 2) The note transition of the selected piece is randomly shuffled and a note transition is selected from top to bottom.
- 3) Back propagation from output to input is carried out on the basis of whether the note transition had a boundary or the rule was applied (=1) or not (=0).

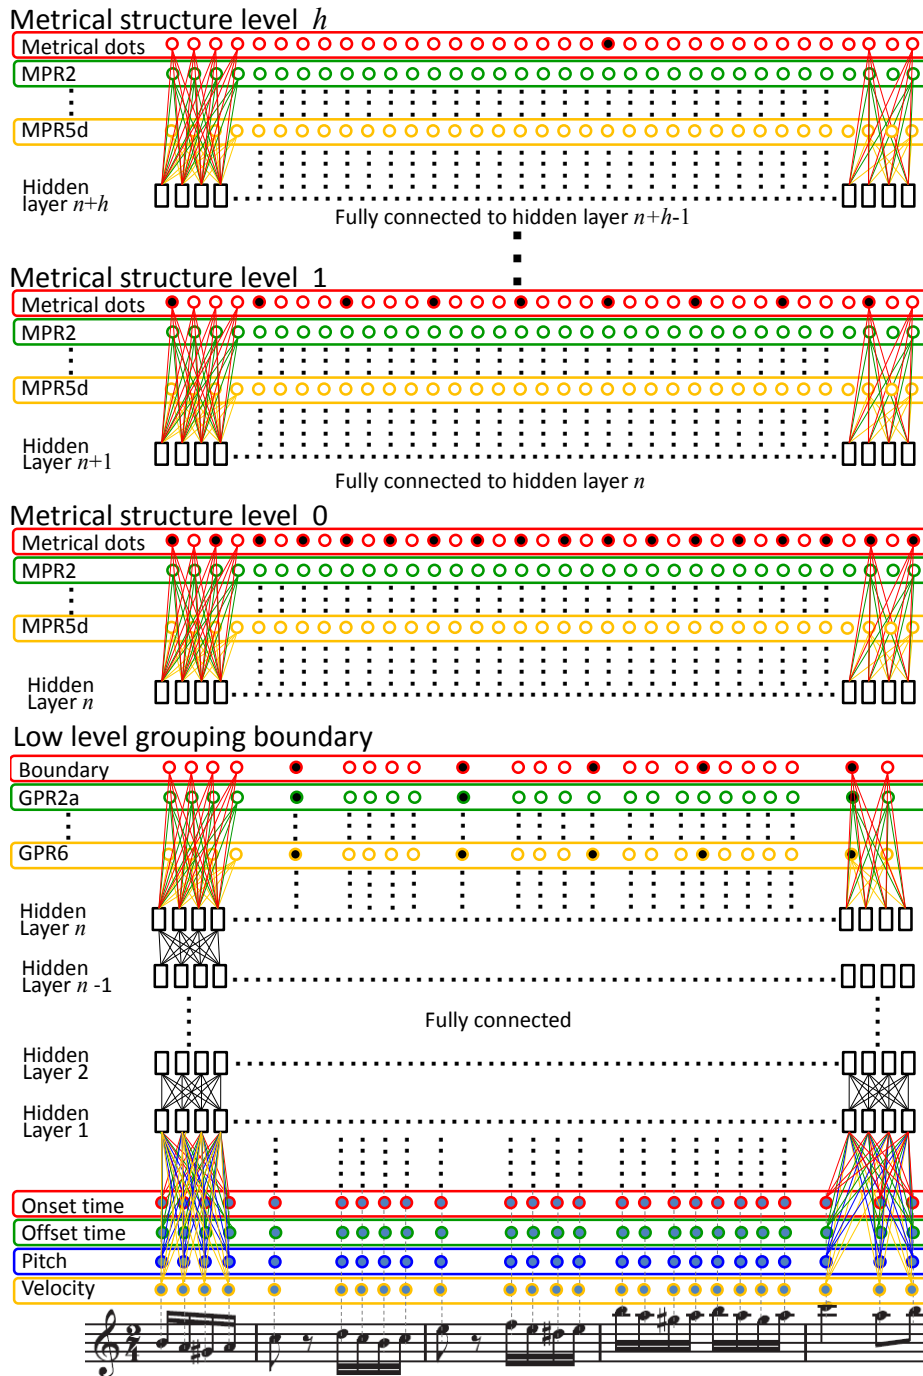


Fig. 6 DBN for deepGTTM-III.

4) The next note transition or the next piece in steps 1 and 2 is repeated.

The processing flow for the learning of an MPR or metrical dots involved four steps.

- 1) The order of the music pieces of training data is randomly shuffled, and a piece is selected from top to bottom.
- 2) The beat positions of the selected piece are randomly shuffled, and a beat position is selected from top to bottom.
- 3) Back propagation from output to input is carried out on the basis of whether the beat position had a strong beat or the rule was applied (=1) or not (=0).
- 4) The next piece in step 1 or the next beat position in step 2 is repeated.

The processing flow for the multidimensional multi-task learning of PRs involves three steps.

- 1) The order of PRs is randomly shuffled, and a rule is selected from top to bottom.
- 2) Multi-task learning of the selected PR is carried out.
- 3) The next rules in step 1 are repeated.

Simultaneous learning of grouping and metrical structures in deepGTTM-III.

The deep layered network of deepGTTM-III is trained by a multi-task learning technique of the grouping and metrical structures. The main difference in learning and acquiring the metrical structure in deepGTTM-II and deepGTTM-III is that the grouping structure is needed in the input data of deepGTTM-II but not for deepGTTM-III. In other words, deepGTTM-III predicts low level grouping boundaries by itself and uses the information for predicting the low level grouping boundaries for predicting the metrical structure.

The network of deepGTTM-I and II first learns each rule application by fixed numbers epochs and then learns the structure on the same epochs. Then it repeats learning of rule applications and structure learning. In contrast, deepGTTM-III repeats learning of GPR applications, grouping structure, MPR applications, and metrical structure by fixed numbers epochs. In the network of deepGTTM-III, all learning processes of grouping and metrical structures interact, so the feedback loop described in section 3 is formed implicitly.

A GPR and a MPR are sometimes learned complementarily when the rules are learned because some GPRs and MPRs are very similar. For example, GPR6 (parallelism) prefers form parallel parts of a group, where two or more segments of the music can be construed as parallel, and MPR1 (parallelism) prefers a parallel metrical structure, where two of more groups or parts of groups can be construed as parallel.

In another example, consider a sequence of four notes: n1, n2, n3, and n4. GPR2b (Attack-Point) states the transition n2-n3 may be heard as a group boundary if the interval between the attack points of n2 and n3 is greater than that between the attack points of n1 and n2 and that between the attack points of n3 and n4. MPR5a prefers a metrical structure in which a relatively strong beat occurs at the inception of a relatively long pitch-event.

4 Experimental Results

We evaluated the deepGTTM by using 100 music pieces from the GTTM database; the remaining 200 pieces were used to train the network. The F-measure is given by

the weighted harmonic mean of precision P (proportion of selected dots that are correct) and recall R (proportion of correct dots that were identified).

Table 1 compares the results for deepGTTM-III with those for deepGTTM-I and deepGTTM-II for a network that had 11 layers with 3000 units. The results indicate that deepGTTM-III obtained a higher F-measure for acquiring local grouping boundaries than deepGTTM-I. On the other hand, deepGTTM-III obtained an F-measure for acquiring the metrical structure similar to that of deepGTTM-II and slightly higher than that of deepGTTM-III. We use the correct grouping structure in the GTTM database because the deepGTTM-II needs the grouping structure for input of the network. In contrast, the deepGTTM-III does not need the grouping structure, so it is effective even when there is no correct grouping structure.

Table 1. Performances of deepGTTM-I, II, and III.

Melodies	Low level grouping boundary		Metrical structure	
	deepGTTM-III	deepGTTM-I	deepGTTM-III	deepGTTM-II
1. Grande Valse Brillante	0.80	0.79	0.93	0.94
2. Moments Musicaux	0.80	0.81	0.99	1.00
3. Turkish March	0.77	0.76	0.96	0.98
4. Anitras Tanz	0.78	0.76	0.90	0.90
5. Valse du Petit Chien	0.80	0.78	0.99	0.99
	:	:	:	:
Total (100 melodies)	0.81	0.78	0.94	0.96

5 Conclusion

We presented deepGTTM-III, which integrates a grouping structure analyzer called deepGTTM-I and a metrical structure analyzer called deepGTTM-II. Whereas deepGTTM-I and deepGTTM-II have to learn grouping and metrical structures independently, the deepGTTM-III learns them simultaneously. Experimental results showed that deepGTTM-III obtained a higher F-measure for acquiring local grouping boundaries than deepGTTM-I and a similar F-measure for acquiring the metrical structure to deepGTTM-II. This work was one step in implementing a generative theory of tonal music (GTTM) based on deep learning. We plan to implement time-span reduction analysis on the basis of deep learning.

Acknowledgments. This work was supported by JSPS KAKENHI Grant Number 17H01847, 25700036, 16H01744, 23500145.

References

1. Lerdahl, F. and Jackendoff, R.: *A Generative Theory of Tonal Music*. MIT Press (1985)
2. Hirata, K. and Hiraga R.: Ha-Hi-Hun plays Chopin's Etude, In *Working Notes of IJCAI-03 Workshop on methods for automatic music performance and their applications in a public rendering contest* (2003)

3. Hirata, K., Matsuda, S., Kaji K., and Nagao K.: Annotated Music for Retrieval, Reproduction, and Sharing, In: *Proceedings of the 2004 International Computer Music Conference (ICMC2004)*, pp. 584–587 (2004)
4. Hirata K. and Matsuda S.: Interactive Music Summarization based on GTTM, In: *Proceeding of the 2002 International Society for Music Information Retrieval Conference (ISMIR2002)*, pp. 86–93 (2002)
5. Hamanaka, M., Hirata, K., and Tojo, S.: Melody Morphing Method based on GTTM. In: *Proceedings of the 2008 International Computer Music Conference (ICMC2008)*, pp. 155–158 (2008)
6. Hamanaka, M., Hirata, K., and Tojo, S.: Melody Extrapolation in GTTM Approach. In: *Proceedings of the 2009 International Computer Music Conference (ICMC2009)*, pp. 89–92 (2009)
7. Hamanaka, M., Hirata, K., and Tojo, S.: Implementing ‘a generative theory of tonal music’, *Journal of New Music Research*, 35(4), 249–277 (2006)
8. Hamanaka, M., Hirata, K., and Tojo, S.: FATTA: Full automatic time-span tree analyzer, In: *Proceedings of the 2007 International Computer Music Conference (ICMC2007)*, pp. 153–156 (2007)
9. Miura, Y., Hamanaka, M., Hirata, K., and Tojo, S.: Decision tree to detect GTTM group boundaries, In: *Proceedings of the 2009 International Computer Music Conference (ICMC2009)*, pp. 125–128 (2009)
10. Kanamori, K. and Hamanaka, M.: Method to Detect GTTM Local Grouping Boundaries based on Clustering and Statistical Learning, In: *Proceedings of the 2014 International Computer Music Conference (ICMC2014)*, pp. 125–128 (2014)
11. Hamanaka, M., Hirata, K., and Tojo, S.: σ GTTM III: Learning-based Time-span Tree Generator Based on PCFG, In: *Proceedings of the 11th International Symposium on Computer Music Multidisciplinary Research (CMMR 2015)*, pp. 303–317 (2015)
12. Hamanaka, M., Hirata, K., and Tojo, S.: Musical Structural Analysis Database Based on GTTM, In: *Proceeding of the 2014 International Society for Music Information Retrieval Conference (ISMIR2014)*, pp. 325–330 (2014)
13. Hamanaka, M., Hirata, K., and Tojo, S.: deepGTTM-I: Local Boundary Analyzer based on a Deep Learning Technique, In: *Proceedings of the 12th International Symposium on Computer Music Multidisciplinary Research (CMMR 2016)*, pp. 8–20 (2016)
14. Hamanaka, M., Hirata, K., and Tojo, S.: deepGTTM-II: Automatic Generation of Metrical Structure based on Deep Learning Technique, In: *Proceedings of 13th Sound and Music Computing Conference (SMC2016)*, pp. 203–210 (2014)
15. Choi, K., Fazekas, G. and Sandler, M.: Automatic Tagging Using Deep Convolutional Neural Networks, In: *Proceeding of the 2016 International Society for Music Information Retrieval Conference (ISMIR2016)*, pp. 805–811 (2016)
16. Zhou, X., and Lerch, A.: Chord Detection Using Deep Learning, In: *Proceeding of the 2015 International Society for Music Information Retrieval Conference (ISMIR2015)*, pp. 52–58 (2015)
17. Deng, J., and Kwok, Y.: Hybrid Gaussian-HMM-Deep Learning Approach for Automatic Chord Estimation with Very Large Vocabulary, In: *Proceeding of the 2016 International Society for Music Information Retrieval Conference (ISMIR2016)*, pp. 812–818 (2016)
18. Oord, A., Sander, D., and Benjamin, S.: Deep content-based music recommendation, In: *Proceeding of the Advances in Neural Information Processing Systems 26 (NIPS 2013)*, pp. 2643-2651 (2013)
19. Sigtia, S., Benetos, E., and Dixon, S.: An End-to-End Neural Network for Polyphonic Piano Music Transcription, *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, 24 (5), 927-939 (2016)
20. MakeMusic Inc., “Finale,” 2017, <http://www.finalemusic.com/>.