Revised Selected Papers

Accademia Musicale Studio Musica Michele Della Ventura, *editor*

2024

Proceedings of the International Conference on New Music Concepts

Vol. 11



Accademia Musicale Studio Musica

International Conference on New Music Concepts

Proceedings Book Vol. 11

Accademia Musicale Studio Musica Michele Della Ventura Editor

Published in Italy First edition: April 2024

©2024 Accademia Musicale Studio Musica www.studiomusicatreviso.it Accademia Musicale Studio Musica – Treviso (Italy) ISBN: 978-88-944350-5-4

Unsupervised Learning of Harmonic Analysis Based on Neural HSMM with Code Quality Templates

Yui Uehara

Kanagawa University yuiuehara@kanagawa-u.ac.jp

Abstract. This paper presents a method of unsupervised learning of harmonic analysis based on a hidden semi-Markov model (HSMM). We introduce the chord quality templates, which specify the probability of pitch class emissions given a root note and a chord quality. Other probability distributions that comprise the HSMM are automatically learned via unsupervised learning, which has been a challenge in existing research. The results of the harmonic analysis of the proposed model were evaluated using existing labeled data. While our proposed method has yet to perform as well as existing models that used supervised learning and complex rule design, it has the advantage of not requiring expensive labeled data or rule elaboration. Furthermore, we also show how to recognize the tonic without prior knowledge, based on the transition probabilities of the Markov model.

Keywords. Automatic chord recognition, Harmonic analysis, Hidden Semi-Markov Models, Neural network.

1 Introduction

Harmonic analysis is the process of representing a musical piece as a sequence of chord labels, which facilitates understanding the structure of the piece. In tonal music, the chord label is called chord degree. The chord degree represents chords with a position (denoted by a Roman numeral) on the scale of a local key. This approach to labeling chords in the context of the key is based on the idea that chord progressions play a crucial role in establishing a tonality, which is fundamental in tonal music. Therefore, harmonic analysis can be applied to various tasks such as composition [6, 28] and higher-level music analysis [5, 20].

Several studies have explored automated harmonic analysis [1, 2, 12, 14, 15, 19, 22-26, 31]. Among them, unsupervised learning is advantageous as it does not require expensive labeled data. However, there have been few works on unsupervised harmonic analysis [19, 31]. The challenge with harmonic analysis is that it involves simultaneously identifying both keys and chords. Since many possible combinations

exist, obtaining an optimal result through unsupervised learning is challenging. Therefore, in order to make the unsupervised learning tractable, previous studies relied on manually designed model parameters [19, 31]. In this sense, the models were not fully unsupervised.

This study presents a new method for unsupervised harmonic analysis¹. The proposed model is based on a hidden semi-Markov model. Most model parameters can be learned automatically using non-labeled data, which is a departure from the previous approaches. As an exception, code quality templates are set manually. The chord quality templates correspond to chord labels in supervised learning, which allow comparison with existing harmonic analysis. We construct the semi-Markov model with the technique of deep latent variable models [9]. The technique is to employ neural networks to approximate probability distributions that comprise a targeted latent variable model, which helps unsupervised learning.

Although the experimental results show that our model still has room to be improved, we exemplify the potential of our model with automatic evaluations and discussions on the obtained harmonic analysis. We also discuss how transition probabilities of the model obtained by unsupervised learning can find the tonic without prior knowledge.

This paper is organized as follows. In Section 2, we review related studies. We introduce the proposed model in Section 3. We describe experimental setups and results in Section 4. Then, we summarize our contributions in Section 5.

2 Related Work

2.1 Automated harmonic analyses

2.1.1 Rule-based harmonic analysis

The Melisma Music Analyzer is one of the leading harmonic analysis models, comprising the Meter, Grouper, Streamer, Harmony, and Key programs [22-26]. Since the chord labels (Roman numerals) in harmonic analysis identify chords within the context of keys, they require information about both keys and chords. The pipeline analyses using Meter, Harmony, and Key programs can provide Roman numerals. The system is rule-based, and utilizes musical knowledge to determine chord tones. Additionally, it provides various criteria for dealing with ambiguous events, for example, preferring chord changes on strong beats and root progression on the line of fifth (circle of fifth) [24]. However, writing down all the preference rules and their

¹ https://github.com/yui-u/harmonic-analysis-chorales

priorities to deal with this ambiguity takes much work. This could pose a challenge when expanding the system to cover different types of music.

2.1.2 Harmonic analyses with supervised learning

On the other hand, supervised learning relies on high-quality labeled data to learn chord discriminators automatically. Masada and Bunescu employed a semi-Markov Conditional Random Field (semi-CRF) for supervised harmonic analysis [12]. In Conditional Random Fields (CRFs), feature functions act like preference rules. In addition, priorities (weights) for each feature function are automatically learned with the labeled data.

Furthermore, Chen and Su developed a harmonic analyzer called the Harmony Transformer that used a Transformer encoder-decoder as an input feature extractor [1, 2]. More recently, Micchi et al. proposed a model with a Convolutional Recurrent Neural Network (CRNN) encoder and Neural Autoregressive Distribution Estimator (NADE) that outperformed the Harmony Transformer [14, 15]. However, even with recent neural network-based supervised learning, the performance of harmonic analysis still has room to be improved. In addition, the dataset is biased toward piano pieces in the classical era and vocal pieces with relatively clear harmony.

2.1.3 Unsupervised harmonic analysis

Increasing the amount of training data is a way to improve performance in supervised learning. However, preparing the data can be expensive. Alternatively, unsupervised learning does not require labeled training data. However, there have been few works on unsupervised harmonic analysis [19, 31]. The challenge with harmonic analysis is that it involves simultaneously identifying both keys (determined by the combination of modes and tonics) and chords. Since many possible combinations exist, obtaining an optimal result through unsupervised learning is challenging. Consequently, past studies had to rely on manually set model parameters [19, 31]. For instance, Wang and Wechsler proposed using an Infinite Gaussian Mixture Model (IGMM) for harmonic analysis [31]. In this model, keys and chords were considered as hidden variables. Since IGMM is a type of Bayesian model, optimal keys and chords for targeted musical notes are obtained through the sampling process. One limitation of their model was that the model parameters were given manually based on musical knowledge. For example, they provided the Gaussian mean and covariance for the major and minor scales, which worked like key profiles. In addition, the IGMM could not consider chord transitions since it was a note clustering.

Raphael and Stoddard also proposed an unsupervised harmonic analysis approach based on Hidden Markov Models (HMMs) [19]. The HMM had the advantage of being

able to account for chord transitions. However, according to Raphael and Stoddard, while the chord degree transition probability was learnable, the key transition probability and the probability of the chord degree after a modulation were difficult to learn [19]. This difficulty is probable because the chord degrees appear relatively evenly, whereas keys and modulations strongly depend on individual pieces.

2.2 Deep latent variable models

Deep latent variable models are methods in which deep neural networks work as the approximators of the probability distributions that make up the latent variable models [9, 13, 27]. In recent years, deep latent variable models have attracted attention as a new method for improving unsupervised learning. For example, Tran et al. proposed unsupervised neural Hidden Markov Models for the part-of-speech tag induction [27], and Miao et al. introduced neural topic models [13].

In the music field, Uehara and Tojo extended the unsupervised neural HMMs to semi-Markov models (HSMMs) and performed recognition of chord segments and their transitions [30]. However, their model was not able to recognize keys and chords simultaneously, and thus could not perform harmonic analysis. Our model is based on Uehara's Neural HSMMs but with significant changes for harmonic analysis, including output emission probability with chord quality templates, simultaneous recognition of keys and roots, and dynamic modulation detection.

3 Methodology

The proposed model for harmonic analysis requires no labeled data for parameter training. In this sense, we describe the model as "unsupervised," which is a significant step forward, given that conventional models require all or part of the model parameters to be designed manually [19, 31]. However, our model incorporates simplifications proposed in previous methods, such as the assumption of transpositional equivalency, a predefined set of chord qualities, and simplified conditional probabilities of the stochastic model.

The hidden semi-Markov model [32] forms the core of the proposed model. Unlike Markov models, semi-Markov models equip a notion of the duration of a state, thus more suitable for segment-level recognition. This property of the HSMM is advantageous since chords are recognized as a result of score segmentation [12]. Among several variants of the hidden semi-Markov model, we utilize the "Residential-time HMM" [33] that assumes independence between the duration of the current and the previous hidden states. The EM algorithm is a widely known method for learning the parameters of an HSMM [32, 33]. However, it has been reported that Neural HSMM, a type of deep latent variable model, could achieve better marginal likelihood than the EM algorithm [29]. As described

in Section 2, the deep latent variable model predicts model parameters through neural networks [9]. In the following sections, we will first overview the proposed HSMM for harmonic analysis. Then, the neural networks that approximate probability distributions that constitute HSMM are described. Table I shows the notations that will be used in the rest of this paper.

| L | Output sequence length | $\in \mathbb{N}$ |
|--------------------------|--|----------------------------------|
| t | Time step | $\in \{0,\ldots,L-1\}$ |
| $oldsymbol{x}_t$ | An observation at t | $\in \{0,1\}^{12}$ |
| pc | Pitch class ² | $\in \{0,\ldots,11\}$ |
| m | Mode index | $\in \{0,1\}$ |
| r | Root pitch class index ^{3,4} | $\in \{0, \dots, 12\}$ |
| i, j | Root pitch class indices that distinguish before/after transition. | $\in \{0, \dots, 12\}$ |
| s | Shift value. (m, s) specifies a key k . | $\in \{0,\ldots,11\}$ |
| k | Key index | $\in \{0,\ldots,23\}$ |
| q | Chord quality index | $\in \{0, \ldots, 6\}$ |
| d | Hidden state duration index | $\in \{0,\ldots,15\}$ |
| $N_{(\cdot)}$ | The number of indices (\cdot) | $\in \mathbb{N}$ |
| $v_{(\cdot)}$ | <i>Logit</i> ⁵ for the probability of (\cdot) | $\in \mathbb{R}$ |
| $oldsymbol{v}_{(\cdot)}$ | Vector of logits for the probability of (\cdot) | $\in \mathbb{R}^{N_{(\cdot)}-1}$ |
| $oldsymbol{z}_{(\cdot)}$ | Index vector with the 1 of N representation ⁶ | $\in \{0,1\}^{N_{(\cdot)}}$ |
| $z_{(\cdot)}$ | The (\cdot) -th component of the 1 of N index vector. | $\in \{0,1\}$ |
| $oldsymbol{e}_m$ | Mode embedding | $\in \mathbb{R}^{12}$ |

3.1 The hidden semi-Markov model for harmonic analysis

The proposed model is based on a hidden semi-Markov model comprising initial hidden-state, hidden-state transition, hidden-state duration, and output emission distributions [30, 32, 33]. The output for the model is a sequence of pitch classes represented in binary twelve-dimension vectors. Since the hidden states are not

 $^{^2}$ The pitch classes are the 12 numbered notes in an octave: {C, C#/Db, D, ..., B} are numbered {0, 1, 2, ..., 11} respectively.

³ In our design, the root pitch class represented by the index r indicates the root note of the chord, independent of the key.

⁴ The last 13th dimension is used as the *Rest* state.

⁵ In this paper, we use the term *logit* as a value that is derived from a neural network and parameterize a probability. The logit is fed into the sigmoid or softmax function to produce the parameter of a targeted probability distribution.

⁶ The *1 of N representation* is an *N*-dimensional vector representation of a category index. For example, for chord quality index q = 3, the corresponding vector of *1 of N representation* is $\{0, 0, 0, 1, 0, 0, 0\}$.

observed and no labeled data is used, what the hidden states represent is not specified. However, simplifications of the model, described later in this paper, make the hidden states correspond to combinations of keys and root pitch classes.

Fig. 1 shows an example of possible paths of hidden states. Note that the numbers of keys, root pitch classes, and the maximum duration are less than the actual settings for ease of reading the figure. At each time step, the hidden states are represented by a combination of a key and a root pitch class. In addition, possible remaining times of the hidden states are combined with each hidden state. Then, triplets of (key, root, duration) represent all possible states. Key or root transitions are only permitted when the remaining duration time is zero in a semi-Markov model. Otherwise, the remaining duration times decrease by one at each time step. Note that the Residential-time HMM, a variant of HSMMs, assumes a hidden state [32, 33]. Thus, the hidden state transition is represented as follows.

$$p(k_t, r_t, d_t | k_{t-1}, r_{t-1}, d_{t-1}) = \begin{cases} p(k_t, r_t, d_t | k_{t-1}, r_{t-1}) & \text{(if } d_{t-1} = 0) \\ \mathbb{1}(d_t = d_{t-1} - 1) & \text{(if } d_{t-1} > 0) \end{cases}$$
(1)

Note that transition probabilities do not include self-transitions, as illustrated in Fig. 1. In addition, we simplify the hidden-state transition probability $p(k_t, r_t, d_t | k_{t-1}, r_{t-1})$ as follows.

$$p(k_t, r_t, d_t | k_{t-1}, r_{t-1}) = \begin{cases} p(d_t) p(r_t | k_t) p(k_t | k_{t-1}) & \text{(if } k_t \neq k_{t-1}) \\ p(d_t) p(r_t | r_{t-1}, k_t) p(k_t | k_{t-1}) & \text{(if } k_t = k_{t-1}) \end{cases}$$
(2)

In the rest of the paper, we call $p(k_t|k_{t-1})$ as the key transition probability, $p(r_t|k_t)$ as the initial root probability, $p(r_t|r_{t-1}, k_t)$ as the root transition probability, and $p(d_t)$ as the duration probability.



Fig. 1. An example of possible paths of hidden states where the number of keys is 2, the number of roots is 2, and the maximum duration is 3. For simplicity, the number of keys, roots, and the maximum duration are less than the actual settings. Thin solid lines represent root transitions, dotted lines represent key transitions (modulations), and bold solid lines represent the consumption of the remaining duration times.

Since there is no previous hidden state when t = 0, the probability of the first hidden state is computed as follows.

$$p(k_0, r_0, d_0) = p(d_0)p(r_0|k_0)p(k_0)$$
(3)

The distribution for $p(d_0)$ is the same as $p(d_t)$, and $p(r_0|k_0)$ is the same as $p(r_t|k_t)$. $p(k_0)$ is the initial key probability.

At each time step, the output x_t depends only on the hidden state at the same time step. We decompose the emission probability $p(x_t|k_t, r_t)$ into a chord quality probability $p(q_t|k_t, r_t)$ and the pitch class emission probability $p(\boldsymbol{x}_t|q_t, r_t)$, where q_t denotes a chord quality.

$$p(\boldsymbol{x}_t|k_t, r_t) = \sum_{q_t} p(\boldsymbol{x}_t|q_t, r_t) p(q_t|k_t, r_t)$$
(4)

Among all the probability distributions in the proposed HSMM, we only set manually $p(\boldsymbol{x}_t|q_t, r_t)$, which we call *chord quality templates*⁷. The details of the chord quality templates are described in 3.1.1.

With the initial (3) and transition (1)(2) probability of the hidden states and the emission probability of the observations (4) described above, the generative process of the HSMM is represented as follows, where L is the sequence length.

$$p(\boldsymbol{x}_{0}, \dots, \boldsymbol{x}_{L-1}, k_{0}, \dots, k_{L-1}, r_{0}, \dots, r_{L-1}, d_{0}, \dots d_{L-1})$$

= $P(k_{0}, r_{0}, d_{0}) \prod_{t=0}^{L-1} P(\boldsymbol{x}_{t} | k_{t}, r_{t}) \prod_{t=1}^{L-1} P(k_{t}, r_{t}, d_{t} | k_{t-1}, r_{t-1}, d_{t-1})$ (5)

3.1.1 Emission distribution with chord quality templates

In this section, we first describe the details of the calculation of output emission distribution. As described above, we decompose the output emission probability into the chord quality and pitch class emission probability. We use seven chord qualities: {major triad, minor triad, diminish triad, dominant seventh, major seventh, minor seventh, and diminish seventh}. Then, we formulate the probability of pitch class emission given the chord quality and the root as follows.

$$p(\boldsymbol{x}_{t}|q_{t}, r_{t}) = \prod_{pc=0}^{11} \text{Bernoulli}(x_{t,pc}|q_{t}, r_{t})$$
$$= \prod_{pc=0}^{11} (\mu_{pc|q,r})^{x_{t,pc}} (1 - \mu_{pc|q,r})^{1 - x_{t,pc}}$$
(6)
$$\mu_{pc|q,r} = \text{sigmoid}(v_{pc|q,r})$$
(7)

 $x_{t,pc}$ represents the value at the pitch class (pc) of an output binary vector x_t . For example, if $x_t = \{0, 0, 1, 0, 0, 0, 1, 0, 0, 1, 0, 0\}$, the values of $x_{t,2}$, $x_{t,6}$, and $x_{t,9}$ are 1, and the remaining are 0.

⁷ More precisely, in addition to this, the maximum inter-key transition probability limit is set manually.

| quality (q) | pc=0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 |
|--------------------------|------|---|---|---|---|---|---|---|---|---|----|----|
| major triad (M) | 1 | | | | 1 | | | 1 | | | | |
| minor triad (m) | 1 | | | 1 | | | | 1 | | | | |
| diminish triad (d) | 1 | | | 1 | | | 1 | | | | | |
| dominant (7) | 1 | | | | 1 | | | 1 | | | 1 | |
| major seventh (M_7) | 1 | | | | 1 | | | 1 | | | | 1 |
| minor seventh (m_7) | 1 | | | 1 | | | | 1 | | | 1 | |
| diminish seventh (d_7) | 1 | | | 1 | | | 1 | | | 1 | | |

TABLE II. SETTINGS OF CHORD QUALITY TEMPLATES.

We set the logit $v_{pc|q,r}$ manually, which we call *chord quality templates*, equivalent to setting chord labels in supervised learning. In particular, we set the quality template for the case root pitch class of 0, as in Table II. $v_{pc|q,0} = w$ when marked by one(1) in Table II, and $v_{pc|q,0} = -w$ otherwise. The weight w is set to 5.0. As described in Table I, we use a bold symbol to represent a vector of logits, for example,

$$\boldsymbol{v}_{pc|q,r} = (v_{pc=0|q,r}, v_{pc=1|q,r}, \dots, v_{pc=11|q,r})^{\top} \in \mathbb{R}^{12}.$$
(8)

If the root pitch class is not 0, the chord templates for a specific root pitch class are obtained simply by shifting the code templates in Table II. For example, when the root pitch class is 5, the chord template of minor seventh

On the other hand, the probability of chord quality given a key and a root pitch class is calculated via neural networks as follows⁸.

$$p(\boldsymbol{z}_{q|m,r}|m,r) = \text{Categorical}(\boldsymbol{z}_{q|m,r}|m,r) = \prod_{q=0}^{N_q-1} \pi_{q|m,r}^{z_{q|m,r}}$$
$$\pi_{q|m,r} = p(q|m,r) = \text{softmax}_q(V_{pc,q|r}^{\top}\boldsymbol{e}_{m,r})$$
(9)

$$\boldsymbol{e}_{m,r} = (\boldsymbol{e}_{m \times r})_r \in \mathbb{R}^{12} \tag{10}$$

$$\boldsymbol{e}_{m \times r} = \boldsymbol{W}_{m \to m \times r}^{\top} \boldsymbol{e}_{m} \in \mathbb{R}^{12 \times 12} \tag{11}$$

Here, $V_{pc,q|r} \in \mathbb{R}^{12 \times N_q}$ is a set of chord templates $v_{pc|q,r} \in \mathbb{R}^{12}$, represented with a tensor. Equation (11) represents the operation of expanding the dimension of mode embedding by a trainable linear transformation to the number of modes times the number of roots, then, eq. (10) represents extracting the components corresponding to the targeted root. The computation of the mode embeddings e_m will be explained in the following paragraph.

⁸ softmax_j(\boldsymbol{x}) = $\frac{\exp(x_j)}{\sum_{j'} \exp(x'_j)}$

Since the probability p(q|m, r) is independent on the time step t, the subscript t is omitted in (9). Hereafter, we sometimes omit the subscript t when the variable is time-independent. Note that the chord quality probability is conditioned on the mode m and root r instead of key k and root r. In this paper, we set the number of modes to two. However, the two modes are not restricted to major and minor modes but are automatically learned by unsupervised learning.

$$h_m = \text{LSTM}(e_{m-1}, h_{m-1})$$
$$e_m = h_m \tag{12}$$

The $e_m \in \mathbb{R}^{12}$ in equation (11) is the mode embedding, which works as a latent feature vector of a mode. This e_m can be a learnable vector but is generated through a Recurrent Neural Network (RNN) in this study, as shown in (12). In particular, we utilize the Long-Short Term Memory [7] as an RNN. In this way, any number of mode embeddings can be generated with the same RNN. In this study, the number of modes is fixed to two, so there is no direct advantage to using an RNN, but this setting is advantageous if the number of modes is larger in the future (e.g., considering church modes) or if it is difficult to fix the number of modes in advance.

Once p(q|m, r) is calculated, p(q|k, r) is obtained by shifting p(q|m, r), where we assume transpositional equivalence in keys. For example, a key with (mode index = 1, shift = 3) is assumed to represent a key, each pitch class of which is shifted by three from (mode index = 1, shift = 0). We let the symbol *m* denote the shift(*s*) = 0 case,

and let k denote k = 0: (m = 0, s = 0), k = 1: (m = 0, s = 1), ...,

$$k = 11: (m = 0, s = 11), k = 12: (m = 1, s = 0), k = 13: (m = 1, s = 1), \dots, k = 13: (m = 1, s = 1),$$

k = 23: (m = 1, s = 11). Note that we do not restrict the tonic pitch class of a mode as r = 0. We will describe how we get the tonic pitch class later. The shift of p(q|m, r) is done without the tonic information. For example, when k = 14, p(q|k = 14, r = 2) = p(q|m = 1, r = 0).

As we have seen, the parameters in equations (11) and (12) are learnable and are utilized to generate the chord quality distribution. Thus, the neural networks generate the distributions that make up the model, which is a characteristic feature of the deep latent variable model [9]. The mode embedding e_m is also used to calculate the initial hidden state and the hidden state transition probabilities described below. The results of previous studies suggest that using such features and network elaborations can lead to better convergence than conventional methods [13, 27, 30].

3.1.2 Hidden state transition distribution

A hidden state is represented in a combination of key, root pitch class, and remaining duration. As shown in equations (1) and (2), the hidden state transition probability is

decomposed in the duration probability $p(d_t)$, root transition probability $p(r_t|r_{t-1}, k_t)$, and key transition probability $p(k_t|k_{t-1})$.

In this work, we assume p(d) is independent of keys and roots since coherent chord length would be preferred. Therefore, the logit for p(d) is a simple learnable vector, the dimension size of which is the maximum duration length. Thus, the duration distribution is obtained as follows, where $v_d \in \mathbb{R}^{N_d}$ is the learnable vector.

$$p(\boldsymbol{z}_d) = \operatorname{Cat.}(\boldsymbol{z}_d) = \prod_{d=0}^{N_d - 1} \pi_d^{\boldsymbol{z}_d}$$
(13)

$$\pi_d^{z_d} = p(d) = \operatorname{softmax}_d(\boldsymbol{v}_d) \tag{14}$$

For the root transition probability $p(r_t = j | r_{t-1} = i, k)$, we first compute $p(r_t = j | r_{t-1} = i, m)$ as follows, then shift them under the assumption of transpositional equivalency.

$$p(\mathbf{z}_{j|i,m}|r_{t-1}=i,m) = \operatorname{Cat.}(\mathbf{z}_{j|i,m}|r_{t-1}=i,m) = \prod_{j=0}^{12} \pi_{j|i,m}^{z_{j|i,m}}$$
(15)

$$\pi_{j|i,m}^{z_{j|i,m}} = p(r_t = j | r_{t-1} = i, m) = \text{softmax}_j(v_{j|i,m})$$
(16)

$$v_{r_t=j|r_{t-1}=i,m} = \mathrm{MLP}_2(\boldsymbol{e}_m \oplus \boldsymbol{v}_{pc|i,m} \oplus \boldsymbol{v}_{pc|j,m})$$
(17)

$$\boldsymbol{v}_{pc|r,m} = \sum_{q} p(q|m,r) \boldsymbol{v}_{pc|q,r}$$
(18)

The equation (18) computes a marginal emission pitch class logit of the root pitch class r. The marginal emission pitch class logit $v_{pc|r,m} \in \mathbb{R}^{12}$ is used as a feature of a chord whose root is r given the mode m. We consider the last root index (r = 12) as *Rest* and set $v_{pc|r,m}$ for it to a 12-dimensional vector with each dimension's value as -w. Then, equation (17) computes a root-transition logit $v_{r_t=j|r_{t-1}=i,m}$ where \oplus denotes a vector concatenation. The root transition probability p(j|i,m) is then computed by the softmax on the set of logits $v_{j|i,m} = (v_{j=0|i,m}, \dots, v_{j=12|i,m})$.

Once $p(r_t|r_{t-1}, m)$ is computed, $p(r_t|r_{t-1}, k)$ is obtained by shifting $p(r_t|r_{t-1}, m)$, similar to the above discussion of the chord quality probability case. For example, when k = 14, which means (m = 1, s = 2), a component of the probability $p(r_t = 9|r_{t-1} = 2, k = 14)$ equals $p(r_t = 7|r_{t-1} = 0, m = 1)$.

We use the same emission distribution, duration distribution, and root transition distribution for all observed data. On the other hand, the key distributions described hereafter can take different values for each observation sequence. Changing the probability distribution of the model depending on the observation is not possible with a conventional HSMM, but is possible with deep latent variable models [9, 13, 27, 30]. We first obtain the embedding of an observed sequence to compute the key transition distribution.

(19)

$$\boldsymbol{h}_{t}^{\text{obs.}} = \text{LSTM}^{\text{obs.}}\left(\tanh(\boldsymbol{W}^{\top}\boldsymbol{x}_{t}), \boldsymbol{h}_{t-1}^{\text{obs.}}\right)$$

$$\hat{\boldsymbol{x}} = \text{MLP}\left(\boldsymbol{h}^{\text{obs.}} \oplus \boldsymbol{t}^{ratio}\right)$$

$$(19)$$

$$\begin{aligned} a_t &= \operatorname{MLP}_2(\boldsymbol{n}_t^{-1} \oplus t^{-1}) \\ &= \exp\left(\hat{a}_t\right) \end{aligned} \tag{20}$$

$$\alpha_t = \frac{1}{\sum_{t'}^{L-1} \exp\left(\hat{a}_{t'}\right)} \tag{21}$$

$$\hat{\boldsymbol{h}}^{\text{obs.}} = \sum_{t} \alpha_t \boldsymbol{h}_t^{\text{obs.}}$$
(22)

The observed data at each time step x_t s first embedded by an LSTM, as eq. (19). Next, the weight for each latent feature $h_t^{\text{obs.}} \in \mathbb{R}^{12}$ is calculated by (20) and (21), where the scalar value $t^{ratio} = \frac{t}{L}$ is an additional information of the time step. Then, the embedding of the observed sequence is obtained by the weighted summation of the latent features, as shown in (22). Using the resulting $\hat{h}^{\text{obs.}} \in \mathbb{R}^{12}$, we compute the mode probability as follows.

$$p(\boldsymbol{z}_m) = \operatorname{Cat.}(\boldsymbol{z}_m) = \prod_{m=0}^{N_m - 1} \pi_m^{z_m}$$
(23)

$$\pi_m^{z_m} = p(m) = \operatorname{softmax}_m(E^{\top} \hat{\boldsymbol{h}}^{\text{obs.}})$$
(24)

Here, $E \in \mathbb{R}^{12 \times N_m}$ is a set of mode embeddings e_m , represented with a tensor. As described in Section 3.1.1, a key is defined by specifying a combination of mode index *m* and shift value *s*. The distribution of modes is obtained by (24), and the shift value is obtained as follows.

$$p(\boldsymbol{z}_{s|m}|m) = \text{Cat.}(\boldsymbol{z}_{s|m}|m) = \prod_{s=0}^{11} \pi_{s|m}^{z_{s|m}}$$
(25)

$$\pi_{s|m}^{z_{s|m}} = p(s|m) = \operatorname{softmax}_{s}(\operatorname{MLP}_{2}(\boldsymbol{e}_{m} \oplus \hat{\boldsymbol{h}}^{\operatorname{obs.}}))$$
(26)

Then, the following equation gives the key probability.

$$p(k) = p(m, s) = p(s|m)p(m)$$
 (27)

Note that the key probability p(k) is not conditioned on the previous key; therefore, it is used as the initial key probability. There are two situations where the initial key probability is used, one at the beginning of the sequence and the other immediately after the modulation. Although the transition probabilities between keys should be considered, the number of parameters to be estimated was reduced by computing p(k) instead of $p(k_t|k_{t-1})$, as the proportions are more important than the transitions for the keys. The key transition probability is simplified as follows.

$$p(k_t|k_{t-1}) = \begin{cases} 1 - \beta & \text{(if } k_t = k_{t-1}) \\ \beta & \hat{p}(k) & \text{(if } k_t \neq k_{t-1}) \end{cases}$$
(28)

 β is a learnable value. However, we set the upper limit of β to 0.01, which means the model penalizes modulations. $\hat{p}(k)$ is the key probability modified so that $k_t \neq k_{t-1}$. At the beginning of the sequence, the key probability p(k) is combined with the initial root probability given a key in the following equation. Here, again, we first compute the initial root probability given a mode p(r|m), and the initial root probability given a key is obtained automatically by shifting p(r|m).

$$p(\boldsymbol{z}_{r|m}|m) = \text{Cat.}(\boldsymbol{z}_{r|m}|m) = \prod_{r=0}^{12} \pi_{r|m}^{z_{r|m}}$$
(29)

$$\pi_{r|m}^{z_{r|m}} = p(r|m) = \operatorname{softmax}_{r}(\boldsymbol{v}_{r|m})$$
(30)

$$v_{r|m} = \mathrm{MLP}_2(\boldsymbol{e}_m \oplus \boldsymbol{v}_{pc|r,m}) \tag{31}$$

In the above equation, e_m is the mode embedding, and $v_{pc|r,m}$ is the marginal emission pitch class logit, the same as those used in (18) and (17). The Multi-Layer Perceptron (MLP) with one hidden layer maps the concatenated vector of $e_m \oplus v_{pc|r,m}$ to a scalar logit $v_{r|m}$. The initial probability given a mode p(r|m) is then computed by applying the softmax function to the set of logits $v_{r|m} = (v_{r=0|m}, \dots, v_{r=12|m})$. The initial root probability given a key p(r|k) can be automatically computed by shifting the p(r|m). Finally, we can obtain the initial probability by p(k, r, d) = p(d)p(r|k)p(k).

3.2 Training

As a deep latent variable model, the proposed model is trained by optimizing the network parameters that produce the probability distributions of the HSMM. The loss function here is the negative log-likelihood (NLL) of the observed sequence $-\log p(x_0, \dots, x_{L-1})^9$, obtained by marginalizing all possible paths of hidden states. The technique to marginalize all possible paths of hidden states for an H(S)MM is known as the *forward algorithm* [17, 32]. The details of the forward algorithm for the HSMM, especially for the Residential-time HMM, can be found in [33] and [30]. In the proposed model, a hidden state is represented in a combination of key and root, as illustrated in Fig. 1. Then, the model can directly apply the forward algorithm for the Residential-time HMM, described in [30].

We perform unsupervised training in two phases. The proposed model is not given a feature of keys like the *Key Profile* [11]. Therefore, in the first phase, we train the score

⁹ More precisely, the loss is the average of the NLL of all the sequences in the mini-batch.

without key signatures to obtain features of two modes. After that, additional training is performed using the original key signature.

Phase 1:

Training with the *normalized* data. *Normalized* means transposing a score to a key without a key signature. If the key signature information has been lost, the score is transposed to maximize the percentage of pitch classes $\{0, 2, 4, 5, 7, 9, 11\}$. In Phase 1 training, keys are limited to the two modes only. These two modes are not necessarily limited to C major and A minor keys, but are learned under the condition that the shift value is always 0. In addition, modulation is disabled by setting the maximum inter-key transition probability limit to zero.

Phase 2:

Training with the original (not transposed) data. In the Phase 2 training, the shift values and the inter-key transition are enabled. Phase 2 training is performed as additional training using the result of Phase 1 as the initial value.

3.3 Inference (harmonic analysis)

After the training, the proposed model performs harmonic analysis by finding the most likely sequence of hidden states (combinations of keys and roots) and their residential times from an observed sequence. Thus, the harmonic analysis is the inference problem of an HSMM, and the method for it is well known as the Viterbi algorithm [4, 17, 32]. After the hidden state inference, the most likely chord quality (q_t^*) for each timestep is obtained by taking the argmax of the probability of output emission given a key and a root.

$$q_t^* = \arg\max_{q} p(\boldsymbol{x}_t|q, r_t) p(q|k_t, r_t)$$
(32)

Thus, we obtain the sequence of keys, root pitch classes, and qualities. Both chord name analysis and Roman numeral analysis are possible based on them. In particular, a chord name is derived from the information of a root pitch class¹⁰ and a quality. The Roman numeral is derived from a key, root, and quality.

To obtain the chord degree (Roman numeral), we first convert the root pitch class by $(r - s) \mod 12$. Here, *s* is the shift value that is combined with mode *m* to specify a key (k = (m, s)). Although our model does not explicitly know the diatonic scale, the transition probabilities of the learned model indicate that the chords with 0 and 7 as the roots are the most dominant. This result will be discussed later, but based on this observation, the correspondence between pitch class and degree is given in Table III. Here, the pitch classes corresponding to I and V are 0 and 7, respectively, and the rest

 $^{^{10} \}rightarrow$ footnote 3

are assigned one degree per two successive roots. After the conversion, the degree is denoted in uppercase if the predicted quality is major; otherwise, it is denoted in lowercase.

TABLE III. CONVERSION TABLE OF SHIFTED ROOT PITCH CLASSES TO DEGREES.

| Shifted root pc. | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 |
|------------------|---|----|----|-----|-----|----|----|---|----|----|-----|-----|
| Degree | Ι | II | II | III | III | IV | IV | V | VI | VI | VII | VII |

Our model does not predict the chord inversion. However, we determine chord inversions in Table IV, where the blank components are assumed to be the basic chords.

TABLE IV. CONVERSION TABLE FOR DETERMINING CHORD INVERSIONS.

| $(bass pc r) \mod 12$ | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 |
|--------------------------|---|----------|--------|---|---|-----|------|---|-----|-----|----|----|
| major triad (M) | | | | | | 6/4 | | | 6 | | | |
| minor triad (m) | | | | | | 6/4 | | | | 6 | | |
| diminish triad (d) | | | | | | | 6/4 | | | 6 | | |
| dominant (7) | | | 2 | | | 4/3 | | | 6/5 | | | |
| major seventh (M_7) | | 2 | | | | 4/3 | | | 6/5 | | | |
| minor seventh (m_7) | | | 2 | | | 4/3 | | | | 6/5 | | |
| diminish seventh (d_7) | | | | 2 | | | 4/3 | | | 6/5 | | |
| | | | | _ | | | -/ - | | | -/- | | |

4 Experiments

4.1 Datasets

We used two different sets of J.S. Bach's four-part chorales: a dataset of 60 chorales formatted by Radicioni and Esposito, and a set of 371 chorales in MusicXML format from the Music21 Library. In this paper, we denote the former dataset as "JSBChorales60" and the latter as "JSBChorales371".

Radicioni and Esposito provided preprocessed scores and human-annotated chord labels in the JSBChorales60 dataset¹¹ [18]. The preprocessed scores include pitch classes, bass pitch classes, and metrical accents computed by the Meter program of the Melisma Analyzer [25]. However, the preprocessing lost some information in the original scores, for example, beat positions, key signatures, time signatures, and parts (except bass pitch classes). The dataset contains 60 chorales in total.

On the other hand, The JSBChorales371 dataset from the Music21 Library [3] is a set of scores in MusicXML format, and it retains all information as a score. In addition, the dataset contains 371 chorales, some of which are not included in the JSBChorales60 dataset. One of the reasons the JSBChorales371 dataset has not been used much is that it is just a collection of scores and does not provide labeled data. Since our model is unsupervised, we can utilize the JSBChorales371 dataset. For evaluation, we consult

¹¹ https://archive.ics.uci.edu/dataset/298/bach+choral+harmony

the human analysis of 20 pieces publicly available in the Music21 Library¹² [3]. We used fixed train, validation, and test splits for the JSBChorales371 dataset. The testing data is the set of 20 pieces where the human annotations are available. We preprocess the MusicXML scores into sequences of pitch classes. The length of each event of pitch classes is 16th-note width, which is the minimum duration of the dataset (except for very few exceptions).

4.2 Experimental setups

For the JSBChorales60 dataset, we followed the original 10-fold cross-validation splits provided by Masada and Bunescu. However, unlike the previous work, we used one fold for testing, another for development, and the remaining 8-folds for training. The random seed was fixed to 123. To train the proposed model, we only used the preprocessed pitch classes. The minibatch size was 2. Since only the annotation of chord names is provided, we performed chord name analysis for the JSBChorales60 dataset. In addition, the chord root names in the annotation were converted to pitch classes before the evaluation since our model did not distinguish enharmonic notes. However, we do not consider this a serious limitation since enharmonic distinctions are possible if the original score information has not been lost in preprocessing.

For the JSBChorales371 dataset, we used fixed train, validation, and test splits for the JSBChorales dataset. The testing data was the set of 20 pieces where the human annotations were available. However, three pieces were excluded because of the collapsed format or inconsistency of key signatures between the annotation and the original score. Thus, the resulting number of testing data was 17. In addition, we used 62 pieces randomly selected as the development data and the remaining 243 pieces for training. The total number of pieces used was 322, where 49 pieces were excluded because of the collapsed format or duplication. In both training and testing, we separated a piece into segments by the *fermata*¹³ positions and treated the segments as independent sequences. We trained with three random seeds: 123, 456, and 789. The minibatch size was 8.

Other settings were the same between JSBChorales60 and JSBChorales371. The number of training epochs was 480 for phase 1 and 240 for phase 2. The best model was chosen by the negative log-likelihood on the development data. The training was stopped prematurely if the best model was not updated in 80 consecutive epochs. The optimizer was Adam [10], and the learning rate was 1e-3.

 $^{^{12}\} https://github.com/cuthbertLab/music21/tree/master/music21/corpus/bach/choraleAnalyses$

¹³ A *fermata* represents a full-stop marker in a lyric in chorales.

4.3 Automatic evaluation results of harmonic analysis

| model | method | Full Chord | Root Chord |
|-------------------------------------|--------------|------------|------------|
| HMPerceptron [18] | supervised | 80.1 | - |
| HMPerceptron (re-experimented [12]) | supervised | 77.2 | 84.8 |
| Semi-CRF [12] | supervised | 83.2 | 88.9 |
| Melisma [25] (reported by [12]) | rule-based | - | 84.3 |
| Ours | unsupervised | 66.8 | 79.2 |

TABLE V. EVALUATION RESULTS OF ACCURACY ON JSBCHORALES60.

The *Accuracy*¹⁴ of our model on the JSBChorales60 dataset is shown in Table V. Note that the proposed model did not consider enharmonic notes. Therefore, we converted the root names to pitch classes when evaluating our model. In this sense, the evaluation is not under exactly the same conditions as the other models in the table. However, as mentioned earlier, we do not consider the enharmonic issue a limitation since it can be resolved if the information in the original score remains.

Table V shows that our model underperformed compared to supervised learning and sophisticated rule-based models. However, we have contributed to the advancement of harmonic analysis with unsupervised learning by enabling unsupervised learning of model parameters and presenting evaluation scores to demonstrate the current performance of unsupervised learning.

TABLE VI. EVALUATION RESULTS OF ACCURACY ON JSBCHORALES371.

| model | Key | Full RN | Root RN |
|-------|------|---------|---------|
| Ours | 74.2 | 61.6 | 66.9 |

As mentioned in Section 4.2, we evaluated chord names on JSBChorales60, whereas we can evaluate harmonic analysis using Roman numerals (RNs) on JSBChorale371. Analysis with Roman numerals requires simultaneous recognition of the keys and roots; therefore, it is more complex than recognizing chord symbols. The *Accuracy* scores for JSBChorales371 with our model are shown in Table VI¹⁵. To the best of our knowledge, this is the first report comparing the manual and unsupervised harmonic analyses attached to the Riemenschneider numbers 1-20 of J. S. Bach's chorales¹⁶.

¹⁴ Accuracy is the percentage of correctly predicted labels from the total number of events in the dataset.

¹⁵ The human annotation sometimes gives multiple interpretations of a single chord at the start or end of modulation, but in this evaluation, the last label was used.

 $^{^{16} \}rightarrow$ footnote 12

However, our model still has challenges, as seen in the scores obtained. In the next section, we will discuss these issues by presenting the results of the analysis.

4.4 Discussion on examples of the obtained analyses

Fig. 2 shows the harmonic analyses of BWV269, bars 13-20. The labels are displayed below the bass notes, if present, or at the places where the note exists, in preference of lower parts. A typical error in our model, exemplified by dotted circles in Fig. 2, is the extra annotation of passing tones. As shown in Fig. 2, the gold annotation distinguished V6-V6/5 (bar 18) and V-V7 (bar 19), while it did not distinguish vi-vi2 (bar 20), unlike our model's prediction. The difference between the V and vi cases may be suggested by whether they were on the Bass passing note or not; however, these differences were difficult to detect in our model, which is based on statistical unsupervised learning.

Another issue with our model is that it does not support borrowed chords; therefore, as shown by the dotted square (bars 15 and 16) in Fig. 2, borrowed chords are sometimes detected as modulations.



Fig. 2. Harmonic analysis of BWV269 (Riemenschneider No.1), bars 13 - 20.



Fig. 3. Harmonic analysis of BWV40.8 (Riemenschneider No.8), bars 7 - 10.

The process of identifying keys is more complex than commonly thought. Fig. 3 shows an example. As shown by dotted circles, both gold and predicted analyses detect a modulation at bar 8. However, a consistent four-degree ascending root progression (F \rightarrow Bb \rightarrow Eb \rightarrow Ab \rightarrow Db) can be observed in bars 7-9. Therefore, it is possible to interpret the progression as (V \rightarrow i \rightarrow IV \rightarrow VII \rightarrow III) without any modulation.

4.5 Tonality derived from the learned root-transition probabilities

As noted earlier, we did not assume that the two modes we set were major and minor, nor did any assumption to tonic notes. Many works have considered that the pitch class with the highest percentage of notes is the tonic [8, 11, 23, 31]. We argue, however, that the tonic should be recognized in the context of harmonic progressions. In particular, we consider that tonic is the largest component of the stationary distribution of the root-transition probability matrix. The stationary distribution of a Markov chain is also used in tasks such as measuring the importance of a Web site [16].

Here, we briefly describe the stationary distribution of a Markov chain [21]. If the initial state vector is π_0 and the transition probability matrix is M, the state probability vector after the *t* steps is as follows.

$$\boldsymbol{\pi}_t^{\top} = \boldsymbol{\pi}_0^{\top} \boldsymbol{M}^t \tag{33}$$

The following equation holds if π_t approaches a constant value π when $t \to \infty$; this π is called the stationary distribution in a Markov chain.

$$\boldsymbol{\pi}^{\top} = \boldsymbol{\pi}^{\top} \boldsymbol{M}, \quad \sum_{l} \pi_{l} = 1$$
(34)

The equation (34) means that the stationary distribution is an eigenvector for eigenvalue 1 of the matrix M.

In our study, the matrix of root transition probabilities corresponds to M in (34). However, since our model is a semi-Markov model and the root transition probabilities do not include self-transitions, we modify the root-transition matrix by combining it with the inverse of the average duration probability a^{dur} to construct M.

$$(a^{\mathrm{dur.}})^{-1} = \frac{1}{\sum_{d=0}^{N_d-1} (d+1) \pi_d^{z_d}}$$
(35)
$$M_m = (1 - (a^{\mathrm{dur.}})^{-1}) \mathbf{I} + (a^{\mathrm{dur.}})^{-1} \begin{pmatrix} \pi_{j=0|i=0,m} \cdots \pi_{j=12|i=0,m} \\ \ddots \\ \pi_{j=0|i=12,m} \cdots \pi_{j=12|i=12,m} \end{pmatrix}$$
(36)

 $\pi_d^{z_d}$ is the same duration probability in eq. (13), and $\pi_{j|i,m}$ are root-transition probabilities (16).

The obtained stationary distributions are shown in Fig. 4. In mode 0, the pitch class with the highest percentage was 2, which made D the tonic. On the other hand, in mode 1, C was the tonic. Thus, m = 0 can be interpreted as the d-minor key and m = 1 as the C-major key. The reason why the d-minor was learned instead of the a-minor can be interpreted as follows. Several pieces in the JSBChorales371 dataset were written in the Dorian mode, which can be identified by having one less key signature than modern notation. Hence, in the Phase 1 training (Section 3.2), the learned mode seemed to be a mixture of d-minor and a-minor. After that, mode 0 converged into d-minor in the Phase 2 training. The learned mode 0 had nearly equal proportions of 0 and 1 pitch classes corresponding to degree VII. Thus, the proposed model learned the ambiguity of degree VII in the minor key. It is also interesting to note that the importance of I and V was similar for both modes, but in d-minor, IV was more important than II.

Furthermore, the pitch class probability can be obtained using the stationary distribution, and the logits of the pitch class given a root computed in eq. (18).

$$p(pc|m, r) = \text{sigmoid}(v_{pc|r,m})$$
(37)

$$p(pc|m) = \sum_{r} p(pc|m, r)p(r)$$
(38)

The p(r) is the root probability calculated as the stationary distribution of the root-transition matrix described above.

Obtained pitch class probabilities are shown in Fig. 5. Unlike the *Key Profiles* [8, 11, 23], the pitch class probability of V could be larger than I. This may be partly due to the output representation as a binary vector of pitch classes. However, the stationary distribution in Fig. 4 appears to express the importance of each degree more clearly than the pitch class probabilities in Fig. 5.



Fig. 4. Result of the stationary distribution of root pitch classes (JSBChorales371, seed=789). The left side is mode m = 0, and the right is mode m = 1.



Fig. 5. Result of the pitch class probabilities (JSBChorales371, seed=789). The left side is mode m = 0, and the right is mode m = 1.

5 Conclusion

This paper proposed an unsupervised harmonic analysis based on the neural hidden semi-Markov model (HSMM). The model was constructed with neural networks that approximate probability distributions in the HSMM. This technique allowed feasible unsupervised learning of model parameters, which has been difficult in previous studies. In addition, we introduced chord quality templates, which enabled the harmonic analysis with known chord labels such as chord names and Roman numerals. Although we presupposed that the number of modes was two and keys were equivalent if transposed, we did not make any other assumption for the two modes. Nevertheless, our model could find the minor and major modes and their tonic notes properly, as was discussed in Section 4.5.

However, the *Accuracy* results on labeled data and examples of obtained analyses suggest that our model still has room to improve. An important future work is to distinguish passing tones. Since our model is a generative model, it assigns probability to all notes in a score. This may cause excess labeling of passing notes, which would be more severe for instrumental music. We may consider a method of probabilistic estimation of passing notes and changing the generation mechanism depending on whether they are passing notes or not.

If the issue of passing tones is resolved, the proposed method should apply to a wider range of music. The true value of unsupervised learning would be demonstrated for music after the late Romantic period, for which there is little labeled data. Furthermore, the model may be effective for pre-Renaissance music since our model could learn modes and tonics unsupervised.

Acknowledgments

This research has been supported by JSPS KAKENHI No. 23K20011. Computational resource of AI Bridging Cloud Infrastructure (ABCI) provided by National Institute of Advanced Industrial Science and Technology (AIST) was used.

References

- [1] T.-P. Chen and L. Su. Harmony Transformer: Incorporating Chord Segmentation into Harmony Recognition. In *Proceedings of the 20th International Society for Music Information Retrieval Conference*, pages 259– 267, 2019.
- [2] T.-P. Chen and L. Su. Attend to chords: Improving harmonic analysis of symbolic music using transformer-based models. *Transactions of the International Society for Music Information Retrieval*, 4(1):1–13, 2021.
- [3] M. S. Cuthbert and C. Ariza. music21: A toolkit for computer-aided musicology and symbolic music data. In *Proceedings of the 11th International Society for Music Information Retrieval Conference*, 2010.
- [4] G. D. Forney. The viterbi algorithm. In *Proceedings of the IEEE*, 61(3):268–278, 1973.

- [5] M. Granroth-Wilding and M. Steedman. Statistical parsing for harmonic analysis of jazz chord sequences. In *International Computer Music Conference*, pages 478–485, 2012.
- [6] R. Groves. Automatic harmonization using a hidden semi-markov model. In AIIDE Workshop, pages 48–54, 2013.
- [7] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [8] D. Hu and L. K. Saul. A Probabilistic Topic Model for Unsupervised Learning of Musical Key-Profiles. In *Proceedings of the 10th International Society for Music Information Retrieval Conference*, pages 441–446, 2009.
- [9] Y. Kim, S. Wiseman, and A. M. Rush. A tutorial on deep latent variable models of natural language. *arXiv preprint arXiv:1812.06834*, 2018.
- [10] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In Proceedings of the International Conference on Learning Representations, 2015.
- [11] C. L. Krumhansl. *Cognitive foundations of musical pitch*. Oxford University Press, 2001.
- [12] K. Masada and R. Bunescu. Chord recognition in symbolic music: A segmental crf model, segment-level features, and comparative evaluations on classical and popular music. *Transactions of the International Society for Music Information Retrieval*, 2(1):1–13, 2019.
- [13] Y. Miao, E. Grefenstette, and P. Blunsom. Discovering discrete latent topics with neural variational inference. In Proceedings of the 34th International Conference on Machine Learning, volume 70 of *Proceedings of Machine Learning Research*, pages 2410–2419, 06–11 Aug 2017.
- [14] G. Micchi, M. Gotham, and M. Giraud. Not all roads lead to rome: Pitch representation and model architecture for automatic harmonic analysis. *Transactions of the International Society for Music Information Retrieval*, 3(1):42–54, 2020.
- [15] G. Micchi, K. Kosta, G. Medeot, and P. Chanquion. A deep learning method for enforcing coherence in Automatic Chord Recognition. In *Proceedings of the* 22nd International Society for Music Information Retrieval Conference, pages 443–451, 2021.
- [16] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: Bring order to the web. Technical report, Stanford University, 1998.
- [17] L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. In *Proceedings of the IEEE*, 77(2):257–286, 1989.
- [18] D. P. Radicioni and R. Esposito. BREVE: An HMPerceptron-Based Chord Recognition System. In Advances in Music Information Retrieval, pages 143– 164. Springer Berlin Heidelberg, 2010.

- [19] C. Raphael and J. Stoddard. Harmonic analysis with probabilistic graphical models. In *Proceedings of the 4th International Conference on Music Information Retrieval*, 2003.
- [20] M. Rohrmeier. Towards a generative syntax of tonal harmony. *Journal of Mathematics and Music*, 5(1):35–53, 2011.
- [21] R. Serfozo. *Basics of applied stochastic processes*. Springer, 2009.
- [22] D. Temperley. An Algorithm for Harmonic Analysis. *Music Perception*, 15(1):31–68, 10 1997.
- [23] D. Temperley. What's Key for Key? The Krumhansl-Schmuckler Key-Finding Algorithm Reconsidered. *Music Perception*, 17(1):65–100, 10 1999.
- [24] D. Temperley. *The Cognition of Basic Musical Structures*. The MIT Press, 2004.
- [25] D. Temperley and D. Sleator. The melisma music analyzer. https://www.link.cs.cmu.edu/music-analysis/.
- [26] D. Temperley and D. Sleator. Modeling meter and harmony: A preference-rule approach. *Computer Music Journal*, 23(1):10–27, 1999.
- [27] K. Tran, Y. Bisk, A. Vaswani, D. Marcu, and K. Knight. Unsupervised neural hidden markov models. In *Proceedings of the Workshop on Structured Prediction for NLP*, pages 63–71, 2016.
- [28] H. Tsushima, E. Nakamura, K. Itoyama, and K. Yoshii. Function- and rhythmaware melody harmonization based on tree-structured parsing and split-merge sampling of chord sequences. In *Proceedings of 18th International Society for Music Information Retrieval Conference*, pages 502–508, 2017.
- [29] Y. Uehara. *Unsupervised Recognition of Chords, Functions*, and Tonality. PhD thesis, Japan Advanced Institute of Science and Technology, 2022.
- [30] Y. Uehara and S. Tojo. Chord function recognition as latent state transition. *SN Computer Science*, 3:508, 2022.
- [31] Y.-S. Wang and H. Wechsler. Musical keys and chords recognition using unsupervised learning with infinite gaussian mixture. In *Proceedings of the 2nd* ACM International Conference on Multimedia Retrieval. Association for Computing Machinery, 2012.
- [32] S.-Z. Yu. Hidden semi-Markov models. *Artificial intelligence*, 174(2):215–243, 2010.
- [33] S.-Z. Yu and H. Kobayashi. An efficient forward-backward algorithm for an explicit-duration hidden Markov model. *IEEE signal processing letters*, 10(1):11–14, 2003

This book presents a collection of selected papers that present the current variety of all aspect of music research, development and education, at a high level. The respective chapters address a diverse range of theoretical, empirical and practical aspects underpinning the music science and teaching and learning, as well as their pedagogical implications. The book meets the growing demand of practitioners, researchers, scientists, educators and students for a comprehensive introduction to key topics in these fields. The volume focuses on easy-to-understand examples and a guide to additional literature.

Michele Della Ventura, editor **New Music Concepts and Inspired Education** Revised Selected Papers

