

Exploring Latent Spaces of Tonal Music using Variational Autoencoders

Nádia Carvalho, Gilberto Bernardes

Faculty of Engineering of University of Porto, INESC-TEC

up201208223@up.pt

Abstract

Variational Autoencoders (VAEs) have proven to be effective models for producing latent representations of cognitive and semantic value. We assess the degree to which VAEs trained on a prototypical tonal music corpus of 371 Bach's chorales define latent spaces representative of the circle of fifths and the hierarchical relation of each key component pitch as drawn in music cognition. In detail, we compare the latent space of different VAE corpus encodings — Piano roll, MIDI, ABC, Tonnetz, DFT of pitch, and pitch class distributions — in providing a pitch space for key relations that align with cognitive distances. We evaluate the model performance of these encodings using objective metrics to capture accuracy, mean square error (MSE), KL-divergence, and computational cost. The ABC encoding performs the best in reconstructing the original data, while the Pitch DFT seems to capture more information from the latent space. Furthermore, an objective evaluation of 12 major or minor transpositions per piece is adopted to quantify the alignment of 1) intra- and inter-segment distances per key and 2) the key distances to cognitive pitch spaces. Our results show that Pitch DFT VAE latent spaces align best with cognitive spaces and provide a common-tone space where overlapping objects within a key are fuzzy clusters, which impose a well-defined order of structural significance or stability — i.e., a tonal hierarchy. Tonal hierarchies of different keys can be used to measure key distances and the relationships of their in-key components at multiple hierarchies (e.g., notes and chords). The implementation of our VAE and the encodings framework are made available online.

keywords: Symbolic Musical Encodings, Latent Spaces, Variational Autoencoders

1. Introduction

One promising avenue in music cognition is using artificial neural networks to produce latent (or embedding) representations of musical data (Kim, 2022; Qiu,

Li, & Sung, 2021). In particular, Variational Autoencoders (VAEs) have shown great potential in generating meaningful and interpretable latent spaces (Roberts et al., 2018; Guo, Kang, & Herremans, 2022). Latent spaces are a mathematical representation that allows for manipulating and analyzing data using machine learning techniques and have been successfully adopted in various applications, such as music recommendation systems, style transfer, and music generation. They have shown promising results in improving the quality of generated music (Roberts et al., 2018; Turker, Dirik, & Yanardag, 2022; Bryan-Kinns et al., 2021; Mezza, Zanoni, & Sarti, 2023).

Symbolic music is typically represented as a temporal sequence of discrete symbols. By computing its latent space, symbolic music can be represented as a **continuous** geometrical space, where multidimensional vectors represent a symbol. Geometric spaces of musical symbols have two main appeals. First, the quality of a musical symbol depends on its spatial relationship with other symbols, i.e., its configurable properties. Its hierarchical dependencies are typically shown when projecting symbols segmented on unitary pitch structures, such as notes and chords. Second, geometrical spaces of low dimensionality provide a concise summary of relations in a form that is easy to visualize and intuitive to understand (i.e., the **circle of fifths** in Figure 1. The continuity of the space allows mathematical operations to be performed on the symbolic data, such as similarity comparisons or clustering, making it easier for machine learning algorithms to learn patterns and generate new music.

The intelligibility and high explanatory power of tonal pitch spaces usually account for a variety of subjective and contextual factors. Historically, tonal spaces can be roughly divided into two categories, each anchored to a specific discipline and applied methods. We have models grounded in music theory (Cohn, 1997, 1998; Lewin, 1987; Tymoczko, 2010; Weber, 1817-1821), and models based on cognitive psychology (Krumhansl, 1990; Longuet-Higgins, 1987; Shepard 1982). Tonal pitch spaces based on music theory rely on musical knowledge, experience, and the ability to imagine complex musical structures to explain which structures work. Cognitive psychology intends to capture and assess the mental processes underlying and relating musical pitch from listening experiments. Despite their inherent methodological differences, they share the same motivation to capture intuitions about the closeness of tonal pitch, which is an important aspect of our experience of tonal music (Deutsch, 1984) and allow the quantification of pitch relations as distances (e.g., What pitch E or G is closer to the A major key?).

Recently, data-driven approaches to the construction of pitch spaces have been pursued from large datasets of symbolic music. Moss, Neuwirth, & Rohrmeier (2022) explore fundamental tonal relations in musical compositions from a corpus representative of historical periods with the aim of studying the evolution of tonal relations across history. Nardelli, Culbreth, & Fuentes (2022) propose a dynamical score network to represent harmonic progressions from an extensive musical corpus spanning 500 years of Western classical music. They found

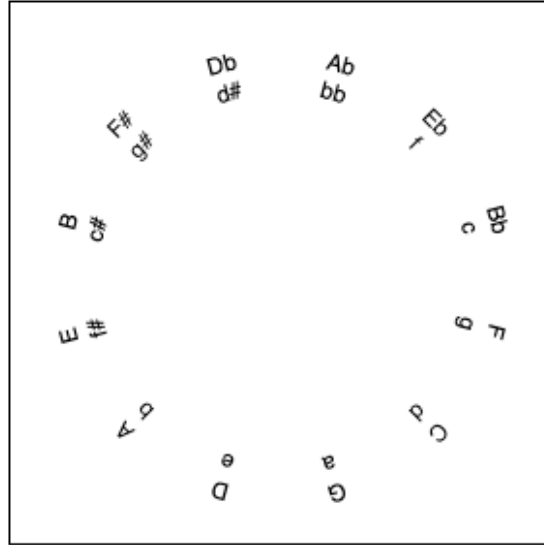


Figure 1: Dimensions 1 and 2 of the four-dimensional multidimensional scaling solution of the intercorrelations between the 24 major and minor key profiles (Circle of fifths), as presented by Krumhansl (1990) from probe-tone experiments.

increased harmonic complexity over the historical evolution of the corpora. Plitsis et al. (2020) and Prang & Esling (2021) explore a large corpus of monophonic and polyphonic musical data, respectively, to evaluate the use of symbolic music encodings in generative models. The former adopts a simple Long-Short Term Memory (LSTM) structure, for which the ABC notation presented the best results overall. The latter adopts the MusicVAE architecture, for which a signal-like representation reflects better reconstruction performance and a latent space more aligned with cognitive musical qualities. Our paper is in line with both works, particularly the last one, but our proposed encoding is simpler than theirs.

The choice of musical encoding is fundamental to the performance of machine and deep-learning techniques in symbolic music tasks, such as generation, transcription, and style recognition (Sarmiento et al., 2023). The properties of the encoding determine the amount and quality of information that can be extracted from the data and thus influence the accuracy and expressiveness of the generated output. For instance, selecting an encoding that can capture high-level semantic features of the data, such as chord progressions or melody patterns, can potentially improve the musical output (Sarmiento et al., 2023). Our paper explores the effectiveness of VAEs while conditioning the model and its ability to produce latent spaces that represent the cognitive pitch distances. To this end, we train VAEs on a prototypical tonal music corpus of 371 Johann Sebastian Bach's (JSB) chorales and compare the latent space of typical VAE corpus

encodings — piano roll, MIDI, ABC, Tonnetz, DFT of pitch, and pitch class distributions. The two latter encodings are proposed in this article and aim to leverage the potential of the DFT of pitch and pitch class distributions in exposing higher-level information on interval and pitch content (Amiot 2016).

Our evaluation adopts four-fold objective metrics to evaluate the models' performance: accuracy, mean squared error (MSE), Kullback–Leibler divergence (KL-divergence), and the computational cost (of each encoding, training the respective VAE, and extracting the information from the original source). Moreover, we quantify the degree to which the encoding's latent spaces provide a pitch space for key relations that align with cognitive distances. In detail, we adopt intra- and inter-segment distances per key and the key distances across all 12 major or minor transpositions per piece to assess the degree of key segmentation between keys and their alignment to the circle of fifths. In eliciting VAE latent spaces with a cognitive, perceptual, and musical theoretical value from tonal music corpus, we can foresee future endeavors which leverage pitch spaces for style-specific musical expressions or less studied harmonic systems, such as modal and microtonal music.

The remainder of this paper is structured as follows. Section 2 describes the methodology of our paper. Section 3 presents each encoding's implementation and characteristics. Section 4 presents our implementation of the VAE model. Section 5 describes our two-folded approach for evaluating both the model performance and the effectiveness of latent spaces in representing cognitive distances between musical pitches. Finally, Section 6 presents the conclusions and avenues for future work.

2. Methodology

Figure 2 shows the architecture of a system we implemented to compare several symbolic music encodings. To process encoding based on the same general properties, we developed a Python 3¹ framework, relying on the music21² library to parse music from different sources (e.g., MusicXML, MIDI, ABC).

For each encoding, we perform musical data augmentations by transposing each piece to all 12 key transpositions per mode. In Section 3, we detail the augmentation strategies implementation per encoding.

The proposed multi-encoding framework allows the training of our VAE model using the same methods (i.e., encoding extraction, decoding of the predictions, augmentation, one-hot encoding, storage, and retrieval of the encoded dataset). The VAE model's implementation, loss functions, and measures rely on the Tensorflow³ framework. We detail its implementation in Section 4.

¹<https://www.python.org/>, Last accessed on 07/03/2023.

²<http://web.mit.edu/music21/>, Last accessed on 07/03/2023.

³<https://www.tensorflow.org/>, Last accessed on 07/03/2023.

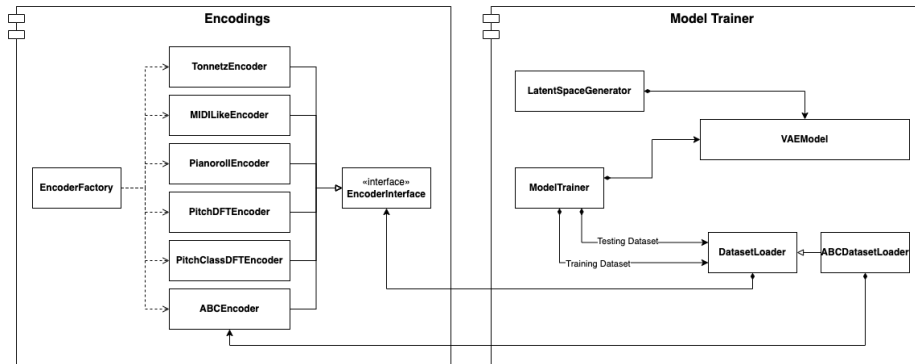


Figure 2: Architecture of the proposed multi-encoding framework. It is constructed in a modular way, allowing the training of our VAE model using the same methods (i.e., encoding extraction, decoding of the predictions, augmentation, one-hot encoding, storage, and retrieval of the encoded dataset). In an application such as the one we constructed, we would simply need to call the EncoderFactory class with the encoding’s name and the ModelTrainer class with the preferred parameters for training the encoding’s dataset.

3. Symbolic Music Encodings

Departing from previous research on symbolic music encodings (Prang & Esling, 2021; Briot, Hadjeres, & Pachet, 2017), we compare four popular corpus encodings — piano roll, MIDI, ABC, and Tonnetz — in providing a pitch space with optimal model reconstruction performance and pitch relations that align with cognitive distances. Furthermore, we present two new encodings, relying on the ability of the Fourier space to describe musical objects and their intrinsic relations: the DFT of pitch class distributions and the DFT of the piano roll. Sections 3.1 to 3.5 present each encoding and its implementation within our work. We offer an engaging platform⁴ for users to explore the encodings through various symbolic music compositions.

3.1. Piano roll

The piano roll encoding uses a binary vector of ones and zeros representing each note sequence’s timestep. Ones denote note activation, and zeros represent the non-activated notes. This method is widely used for encoding melodic and polyphonic music structures and is known for its simplicity. The most noteworthy limitation, shown in Figure 3, is its inability to determine the end of each represented note. As shown in Figure 4, we address this limitation by extending the piano roll encoding to twice its length. The first half of the encoding pertains to notes starting at the timestep, while the second half refers

⁴<https://nadiacarvalho.github.io/Latent-Tonal-Music/>

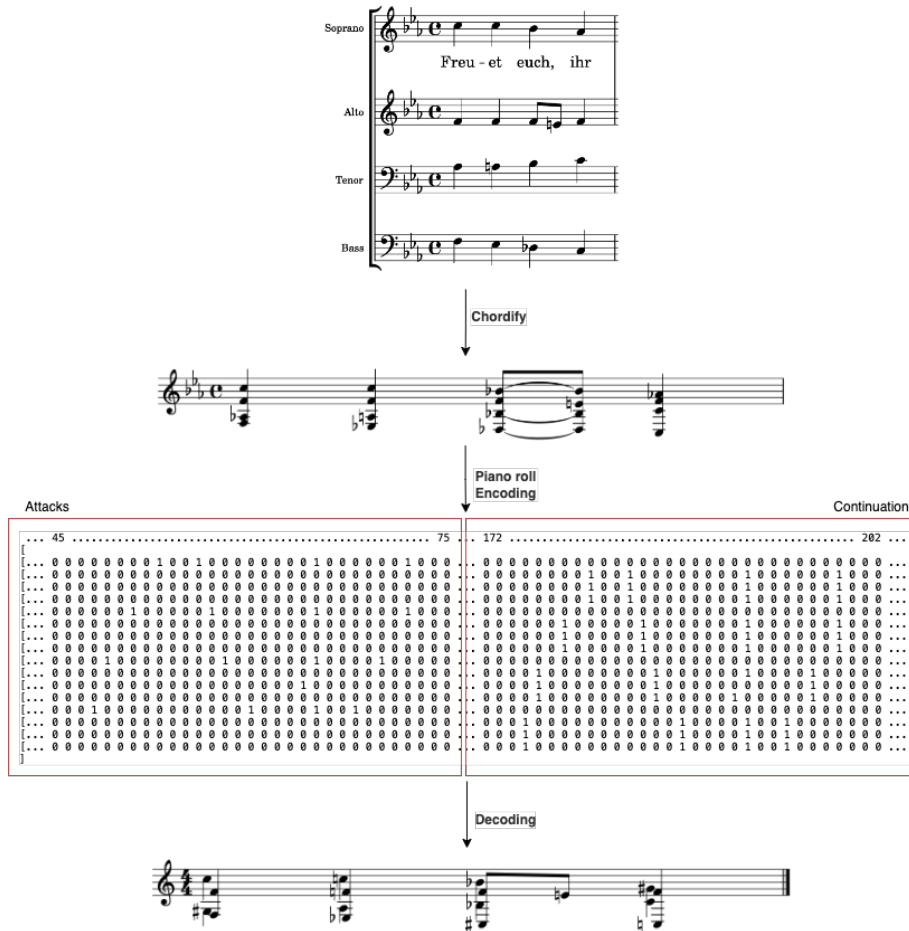


Figure 4: Process of encoding and decoding the first measure of a "Freuet euch, ihr Christen alle Bach" (BWV 40/8) as a piano roll. (Using our "continuation"-based encoding, 0-128 for attack notes and 129-256 for continuation notes).

To compute the augmentations, i.e., transposing the piano roll encoding to a different key, we rotate the representation vector by the number of half-tones corresponding to the transposing interval. This process is done separately for attacks and continuations. First, we rotate the first part of the vector containing the attacks. Second, we rotate the second part containing the continuations to maintain the original key's encoding continuity.

3.2. MIDI-like

Our work adopts the MIDI (Musical Instrument Digital Interface) protocol as an encoding, as first proposed by Oore et al. (2018). This approach relies on a vocabulary of four main MIDI events, namely the NOTE_ON event, the corresponding NOTE_OFF event, the SET_VELOCITY event, and the TIME_SHIFT event. These events represent each timestep of an input sequence as a discrete event, handling any form of music with varying degrees of polyphony and metrical variation. Figure 5 shows the extraction process of the MIDI-like encoding.

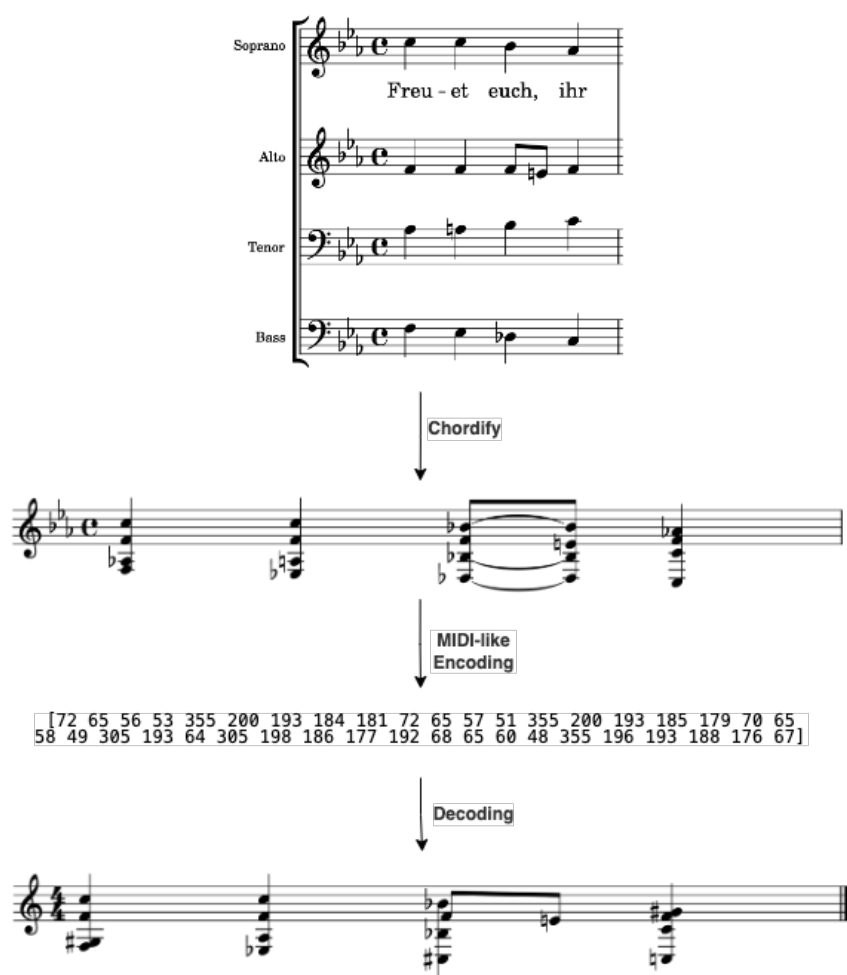


Figure 5: Process of Encoding and Decoding the first measure of a "Freuet euch, ihr Christen alle Bach" (BWV 40/8) with MIDI-like encoding.

MIDI-like encoding only has one value per timestep. Therefore, we can perform augmentation by adding the number of half-tones corresponding to the transposing interval to the MIDI-like encoded music. Similarly, when the encoding is in its one-hot form,⁵ we only need to rotate the encoded music horizontally by the same number of elements as the transposing interval.

3.3. ABC

ABC notation uses letters, numbers, and symbols to represent musical notes and rhythms. The system uses basic rules to represent each musical element, such as pitch, duration, and ornamentation. One of the advantages of ABC notation is its simplicity and ease of use, as it can be quickly learned by musicians and non-musicians alike (Sturm et al., 2018).

ABC notation has been mainly applied to monophonic music structures, i.e., melodies (Briot, Hadjeres, & Pachet, 2020). We adopt it as a form of textual representation of vertical or harmonic aggregates. To this end, we developed a parser from the music21 structures to ABC. Our parser implementation is based on the Javascript Midi2ABC parser developed by Marmoo.⁶ The process (see Figure 6) involves converting a segment of notes into ABC notation by dividing them by instrument sections and adding headers for each instrument. The chords and notes are then iterated through and converted to ABC notation, taking into account component pitches, note duration, tie marks, rests, and tuplets. The approach ensures correct timing and chord order.

We recognize some limitations in the ABC notation related to the retainment of correct durations and passing notes, namely due to the middle processing stage of chordification of the harmonic music texture, as shown in Figure 6 (e.g., the F in the second voice of the third chord would be lost). However, since the study focuses on the cognitive distances from musical pitch within a key, the former problem is not considered critical, as it ultimately may discard some non-tonal tones or aggregate them into vertical slices. To address the latter issue, we chose to incorporate passage notes by splitting the tied notes of chords in which one voice moves, treating them as distinct. This approach allowed us to retain more information from the original music source in the encoded score, regardless of possible absent notes.

3.4. Tonnetz

Chuan & Herremans (2018) introduced an extended Tonnetz musical encoding based on the Tonnetz graphical representation proposed by Euler (1739). Music theorists and musicologists have long used the Tonnetz to investigate tonality and tonal spaces.

⁵One-hot encoding is a technique used to represent categorical data as binary vectors.

⁶Implementation is available at <https://github.com/marmooo/midi2abc>, last seen 07/03/2023.

Online converter is available at <https://marmooo.github.io/midi2abc/>, last seen 07/03/2023.

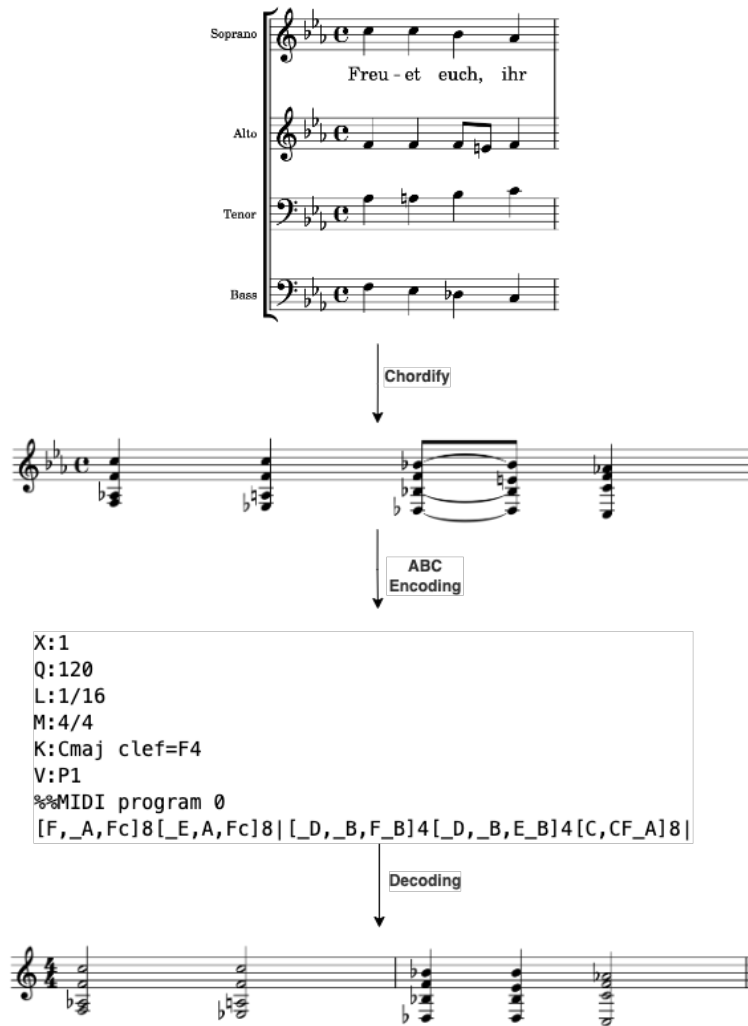


Figure 6: Process of Encoding and Decoding the first measure of a "Freuet euch, ihr Christen alle Bach" (BWV 40/8) with the ABC encoding.

The construction of the Tonnetz typically involves using 12 pitch classes, with nodes arranged in a circle-of-fifth sequence. Nodes to the right create a cycle of perfect fifths, while nodes to the left form a cycle of perfect fourths. Triangles in the network represent a triad, with the parallel major and minor triads connected vertically by sharing a baseline (see Figure 7).

Our approach employs the expanded Tonnetz version proposed in Chuan & Herremans (2018), utilizing a 24-by-12 matrix where each node represents a

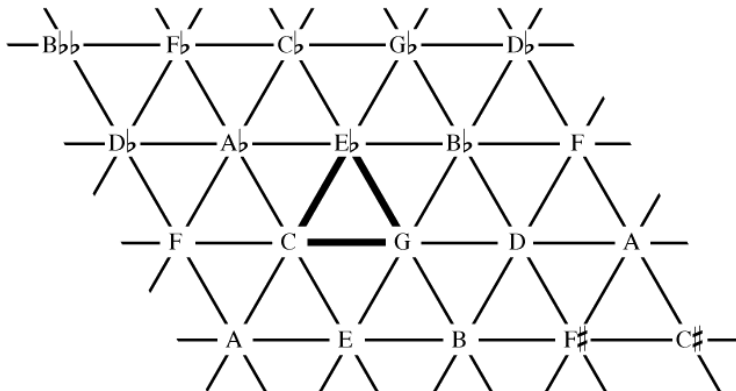


Figure 7: Prototypical Tonnetz representation. Three nodes forming a triangle in the network represent a triad, such as the ones in bold, representing the C minus chord triad.

pitch (not a pitch class, as in the traditional Tonnetz). The pitch register information is maintained (from C0 to C#8), determined by the proximity to the central column's pitch. Nodes on the same horizontal line exhibit the circle-of-fifth relationship. The expansion of the traditional one-octave Tonnetz facilitates simultaneous pitch modeling, which is critical to encoding polyphonic music.

We further extend the approach by duplicating the matrix horizontally to capture attacks and continuations, using a piano roll-like approach where active notes are encoded as ones and non-active notes as zeros (see Figure 8). However, as each pitch appears at multiple positions in the matrix, all pitch positions must be activated for each pitch. This makes computing Tonnetz augmentations challenging, so transposed versions of symbolic music are encoded as new Tonnetz encodings.

3.5. DFTs of Pitch and Pitch Class Distributions

From previous work on adopting Fourier space to describe musical objects and their intrinsic relations (Quinn, 2006; Yust, 2015; Amiot, 2016; Bernardes et al., 2016), we adopt two different encodings based on the discrete Fourier transform (DFT) of pitch distributions. The first applies the DFT on a binary pitch distribution of $m = 128$ elements, similar to a piano roll, i.e., a binary vector representation where active notes are represented as ones. The second reduces such distribution to the 12 pitch classes, where $m = 12$. We adopt the non-trivial or non-symmetrical output of the DFT, which results in $m/2 + 1$ complex numbers. The resulting vector can be converted into magnitude and phase information from which musical objects can be interpreted. Magnitudes

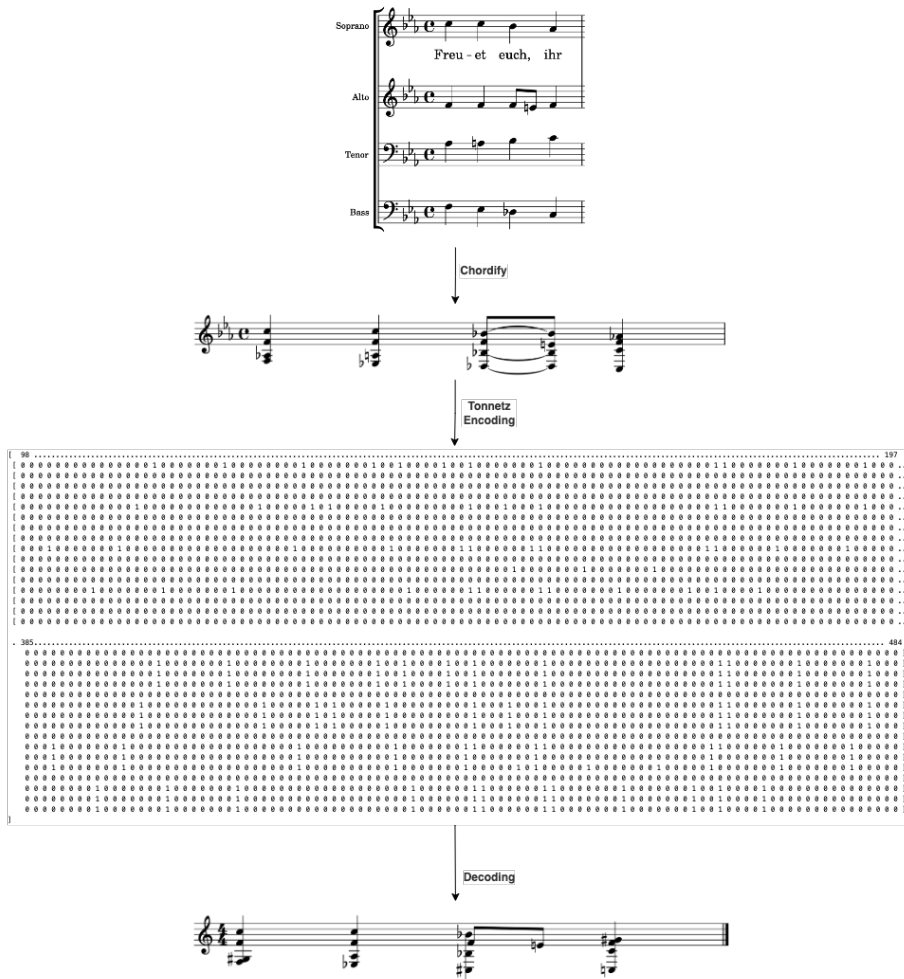


Figure 8: Process of Encoding and Decoding the first measure of a "Freuet euch, ihr Christen alle Bach" (BWV 40/8) with the Tonnetz encoding. For legibility purposes, the continuations are below the attacks.

encode the interval content of the represented pitch, and phases encode the degree to which musical objects or vectors share common tones.

To represent a pitch class distribution using the Discrete Fourier Transform (DFT), we first extract the pitch class information from the notes and chords into a two-dimensional array of 24 columns per timestep. Its position depends on whether it is attacked (first 12) or continued. The resulting pitch information is transformed using the DFT per note attacks and continuations (see Figure 9). The first component of the list represents the number of attack activations, fol-

To perform the second DFT encoding, shown in Figure 10, we use a similar process to the first half. However, instead of creating a two-dimensional array with 24 zeros per timestep, we compute a piano roll of 256 elements, with 128 elements for attack activations and 128 for continuations. We then apply the DFT separately on the attacks and continuations, extracting only the non-symmetrical components, which are the first 64 components for each DFT output. Finally, we concatenate the two DFT vectors into a unique representation.

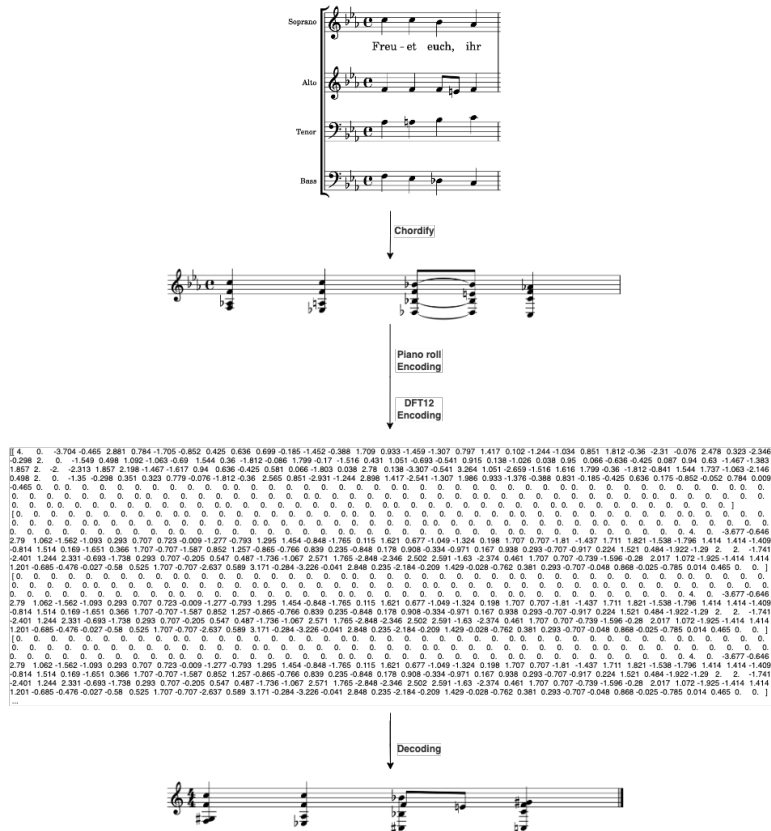


Figure 10: Process of Encoding and Decoding the first measure of a "Freuet euch, ihr Christen alle Bach" (BWV 40/8) with the encoding based on DFT from Pitch. For legibility, we only show the output of the first four timesteps.

We use the circular frequency shift property of the DFT to augment each encoded timestep. Changing the pitch or pitch classes only affects the phases of the DFT and not the magnitudes. To compute the new phases of the encoding, we first compute the DFT of the pitch class (or pitch) distribution with only the second component activated. For each complex component of the output, we calculate and store its phase, using

$$a(z) = \text{atan2}(\text{Im}(z)/\text{Re}(z)) \quad (1)$$

Then, to augment a specific song by a transposition t , we rotate the angles of the original DFT components (attacks and continuations are processed individually), resulting in a new component with real and imaginary parts, such that:

$$\text{Re}(z + t) = \text{Re}(z) * \cos(t * a(z)) - \text{Im}(z) * \sin(t * a(z)) \quad (2)$$

$$\text{Im}(z + t) = \text{Re}(z) * \sin(t * a(z)) + \text{Im}(z) * \cos(t * a(z)) \quad (3)$$

4. VAE Model

We employ a VAE model (Kingma & Welling, 2013) to train and compare the latent space-generated encodings. VAEs are generative neural network models that can learn a compressed representation of data, such as images or music, by encoding it into a low-dimensional latent space. VAEs are designed to generate new data samples that resemble the input data by sampling from the learned latent space (Kingma & Welling, 2013).

A VAE consists of two main components: 1) an *encoder*, $e(x) : \mathbb{R}^x, \mathbb{R}^z$, that maps the input data to a lower-dimension latent space, and 2) a *decoder*, $d(x) : \mathbb{R}^z \mathbb{R}^x$, that maps the latent space back to the input space, so that:

$$\hat{x} = d(e(x)) \approx x \quad (4)$$

During training, VAEs minimize a loss function that ensures the generated data samples are similar to the input data while encouraging the distribution of latent variables to follow a prior distribution, such as a Gaussian distribution. This loss function is defined by two parts (Equation 5): the reconstruction loss, which evaluates how closely the output matches the input, and 2) the KL-divergence loss, which quantifies the discrepancy between the learned latent space and a prior distribution.

$$\log p_{\theta}(x) = \mathbb{E}_{q_{\phi}(z|x)}[\log[\frac{p_{\theta}(x, z)}{q_{\phi}(z|x)}]] + D_{KL}(q_{\phi}(z|x)||p_{\theta}(z|x)) \quad (5)$$

This encourages the model to generate diverse and realistic samples. In the symbolic music domain, we note the widely known MusicVAE (Roberts et al., 2018), which introduced a hierarchical decoder, the conductor, that first outputs embeddings for subsequences of the input and then uses these embeddings to generate each subsequence, independently.

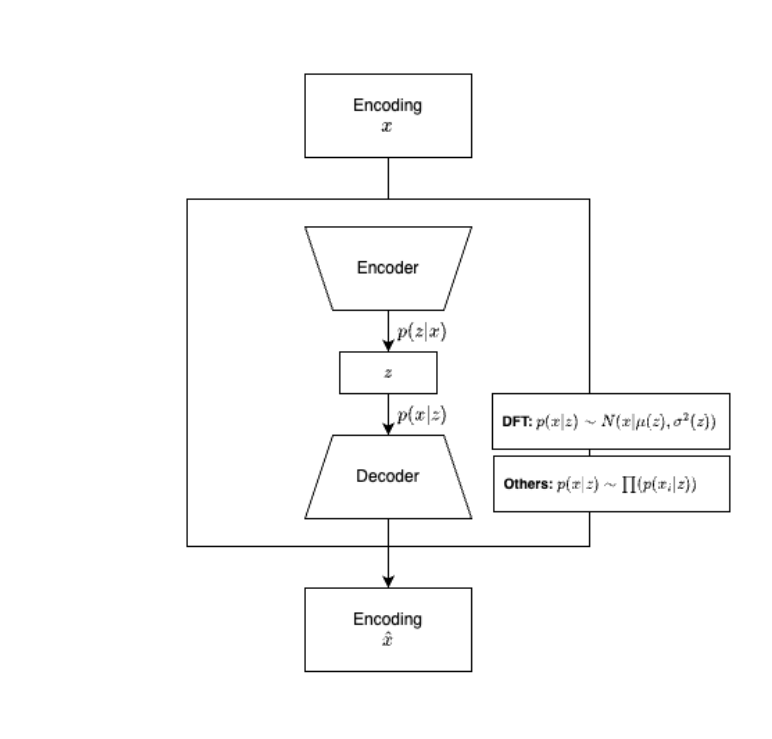


Figure 11: Variational Encoding Implementation used in our experiments. Encoder and Decoder are LSTM layers with 1024 units.

To explore the correspondence between latent space structure and human tonal perception, we opted for a flat baseline recurrent VAE model instead of the more powerful MusicVAE, enabling a comparative study of latent spaces and reconstruction possibilities across different representations. Figure 11 shows the implementation of the VAE. It consists of two recurrent LSTM layers with 1024 units, serving as encoder and decoder, respectively. We employ a categorical cross-entropy reconstruction loss function for tokenized one-hot encoded data (e.g., MIDI-like and ABC formats) and a binary cross-entropy reconstruction loss function for multi-hot encoded data (e.g., piano roll and Tonnetz formats). For DFT-encoded data, which consists of float values, we use an MSE reconstruction loss function. To minimize the loss function and optimize model parameters during training, we use the Adam optimizer with an initial learning rate of $10\text{E-}4$ and a batch size of 256. Finally, we employ a latent size of 256 for all musical encodings, which significantly reduces the input dimensionality while preserving sufficient information for accurate reconstruction (Prang & Esling, 2021). All VAE models were trained using these parameters.

5. Evaluation, Results, and Discussion

5.1. Evaluation

In this section, we present a twofold evaluation strategy of the six musical encodings described in Section 3 trained in the VAE model defined in Section 4 to assess 1) the quality of musical embeddings in reconstructing data and training the network and 2) the alignment of latent spaces to key relations and distances in cognitive-led spaces. Ultimately, we aim to assess the effectiveness of each musical encoding in training a VAE and the ability of the latent space to capture tonal relations from the input embeddings and define a data-driven pitch space. Our evaluation and generated latent spaces adopt the JSB chorales as a test dataset, extracted from the music21 library. The dataset is composed of 195 chorales in a major key (53%) and 176 in a minor key (47%), with G Major (14%), A minor (12%), and G minor (11%) being the most representative keys. On average, a chorale is constructed of 84 chords. We use 60% of the chorales from the dataset for training and the other 40% for testing. The training chorales are then augmented by transposing (up and down) the encoding to the twelve keys.

5.1.1. Model Performance

To evaluate each encoding’s ability to reconstruct the source data, we examine the 1) accuracy, 2) MSE, and 3) KL-divergence scores for every ten-timestep segment.

Accuracy measures how well the reconstructed sequence matches the original sequence and is reported in percentage. The highest accuracy score of 100% results from reconstructed sequences closely matching the original sequences (Briot & Pachet, 2018). MSE reckons the average squared difference between the reconstructed and original sequences. A lower MSE score indicates that the reconstructed sequence closely matches the original sequence (Briot & Pachet, 2018). KL-divergence estimates the difference between two probability distributions. In the context of sequence reconstruction, it measures the difference between the original and reconstructed sequence’s probability distributions. A lower KL-divergence score, close to zero, indicates reconstructed sequences closely matching the original sequence’s probability distribution (Prokhorov et al., 2019).

Additionally, we analyze the computational cost associated with training each musical encoding on the VAE and provide insights into the intelligibility and invariant procedure of the encodings.

5.1.2. Latent Space Analysis

To inspect the alignment of the VAE latent space trained from different encoding to cognitive-driven tonal space, we adopt a twofold strategy. First, we project twelve annotated transpositions of a JSB chorale onto the latent space, dividing them into ten-timestep segments. Second, assuming each key is a cluster,

we compute cluster metrics and non-parametric circular statistics methods to evaluate its performance. In other words, we aim to find how well each key transposition is segmented and whether their spatial arrangement follows the circle of fifths. Figure 12 shows the latent space of the choral “Ich dank dir, lieber Herre” (BWV 347) in A major. The choral is projected in all 12 major keys using a DFT of pitch distribution encoding. When the originally projected choral is in the minor mode, we transpose it to the remaining 11 minor keys. While plotting the latent space, we adopt the Camelot Wheel’s⁷ colors and key colors and key enumerations (numbers 1-12 are keys B to E, by fifth intervals, while the following letter, A or B, is the minor and major mode, respectively).

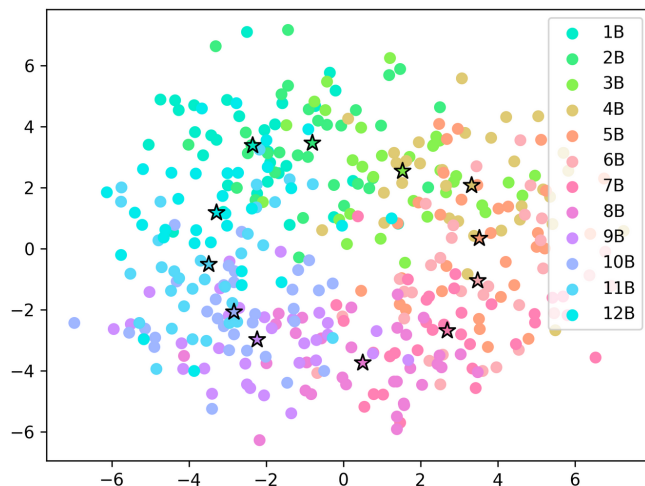


Figure 12: Latent space of the choral “Ich dank dir, lieber Herre” (BWV 347) in A major, transposed into the remaining 11 major keys from a DFT of pitch distribution encoding. We adopt the Camelot Wheel for coloring and numbering the clusters in the latent space (i.e., the keys B to E, arranged by fifths, are represented by numbers 1 to 12, followed by a B for the piece’s major mode). Each key’s cluster centroid is denoted by a colored star that matches the respective key.

In detail, the key-annotated music segments from the encoded musical data in all 12 keys are returned as a list of tuples. Then, we process each segment’s tuples through the pre-trained model’s encoder to create its latent space. To visualize and reduce the computational complexity of the evaluation metrics, we employ principal component analysis (PCA) to reduce the multidimensional

⁷Davis, Mark, “Harmonic Key Selection”, Camelot Sound, <http://www.camelotsound.com/Easymix.aspx>. Last accessed on 07/03/2023.

latent space data to two dimensions while retaining as much of the original data variation as possible (Pearson, 1901). Finally, we calculate their intra-segment distances per key and inter-key distances from the resulting space to identify if they exist as separate entities and somehow align with the expected fuzzy cluster behavior from cognitive space. Within the tonal music context, the fuzzy nature of the resulting clusters is expected due to the shared pitch between keys. Two neighbor keys typically have one different pitch element and share the remaining collection of pitches. Therefore, a perfect key cluster separation is not expected. Furthermore, we compute the order in which the keys are located in the space and thus their relationships, which, as shown in Figure 1, align with the circle of fifths. Distances in the latent space are understood as the proximity or relation between two segments. Therefore, the smaller the distances, the more related the two segments are expected to be.

We adopt two cluster-evaluation metrics, namely the Davis-Bouldin score (Davies and Bouldin, 1979) and Dunn index (Dunn†, 1974), to assess the intra-segment distances per key and inter-key distances from the latent spaces. Each of the 12 key transpositions per dataset chorale is understood here as a cluster. These measures roughly capture the silhouette per key (i.e., intra-segment distances per key) and the segregation between keys. In detail, the Davis-Bouldin score and Dunn index capture the degree to which the resulting latent space provides a compact and well-separated key cluster. The Davis-Bouldin score is lower for better-separated clusters and higher for poorly separated clusters. For the Dunn index, we adopt as intra-cluster metric (or cluster diameter) the average Euclidean distance across all key cluster segments and, as inter-cluster metric, the distance between each cluster's nearest neighbors. The Dunn Index score results from the ratio between the inter-cluster and intra-cluster distances. It aims to be maximized, indicating better separated and compact clusters.

To assess the degree to which the key positions from latent space align with the circle of fifth in cognitive spaces, we adopt a non-parametric sample circular correlation coefficient measure. In detail, we apply the circular non-parametric Kendall's Tau, as proposed by Fisher and Lee (1982), to measure the degree of association between a circular sequence of keys in fifths and the resulting order of keys in each latent space. Tau correlation coefficient ranges between -1 and 1, where -1 indicates a perfect negative association, 0 indicates no association, and 1 indicates a perfect positive association between two variables. Optimal correlations for our problem result from maximizing the absolute value of the Tau coefficient, as both -1 and 1 capture the order of the keys in circles of fifths (see Figure 13).

5.2. Results and Discussion

5.2.1. Model Performance

Table 1 presents objective results on the model's reconstruction performance. Interestingly, KL-divergence and MSE demonstrate consistent outcomes, while

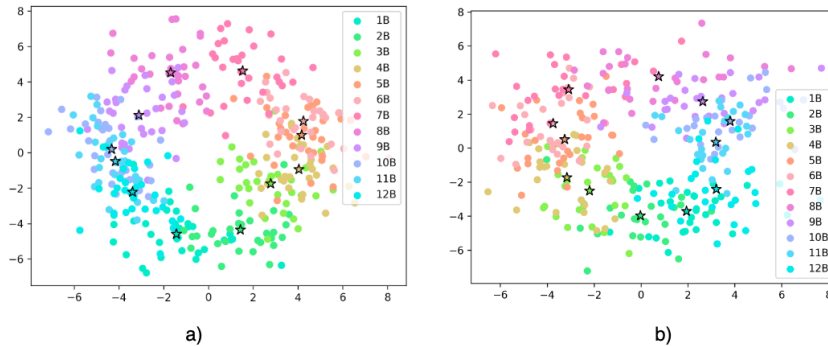


Figure 13: Latent Spaces extracted from the Pitch DFT of two chorales in all twelve key augmentations, whose Kendall’s Tau coefficient is 1 (Fig. a) and -1 (Fig. b). The order of the circle of the fifths is exactly equal, but the clusters are sequenced in either **anti-** or **clockwise** direction. We adopt the Camelot Wheel for coloring and numbering the clusters in the latent space (i.e., the keys B to E, arranged by fifths, are represented by numbers 1 to 12, followed by a B for the piece’s major mode). Each key’s cluster centroid is denoted by a colored star that matches the respective key.

accuracy does not reflect the same encoding order.

The encoding with the best accuracy in music reconstruction is ABC notation (83%), closely followed by MIDI-like (77%), the two DFT methods (77% and 76%). Both ABC and MIDI-like encodings, which use categorical one-hot encoding, performed well in all three measures by presenting the lowest values for both KL-divergence and MSE scores. Therefore, they are expected to capture the most information from the original chorales, aligning with the results of Plitsis et al. (2020).

Interestingly, the Pitch Class Distribution DFT achieved a higher accuracy score than the Pitch DFT, which suggests that the fewer features of the latter may lead to better accuracy results. However, DFT methods resulted in higher KL-divergence and MSE scores, indicating that the reconstructed sequence and its probability distribution have poorer matching when compared to the original sequence. At the same time, this may be interpreted as a sign that these models are learning a more informative latent space (Ucar, 2019).

The multi-hot encoded musical encodings (piano roll and tonnetz) have the poorest model performance, as reflected by the remarkably low accuracy scores. Moreover, these representations are still prone to overfitting despite our effort to minimize it through the regularisation techniques, such as dropout and data augmentation.

Musical Encoding	Reconst. Accuracy (%)	KL	MSE	Time per Epoch (min) ⁸	Embedding Time (min) ⁹
Piano roll	<1	12.48	1.4E-02	~50	~2
MIDI-like	77.4	.80	8.0E-04	~2h30	~2
ABC	82.6	.54	6.7E-03	~5h30	~3
Tonnetz	<1	35.29	1.9E-02	~55	~3
PC DFT	77.3	136.02	6.5E-01	~1h10	~5
P DFT	75.6	472.79	9.1E-01	~1h10	~20

Table 1: Results on the model’s reconstruction performance. PC DFT and P DFT stand for Pitch Class DFT and Pitch DFT.

In terms of computational cost, the piano roll had the shortest training time, closely followed by the tonnetz encoding. The DFT encodings had similar training times, although not as low as the previous two. In contrast, the ABC and MIDI-like encodings require much more sequences to train than the others, even with the same sequence length. Consequently, they require significantly longer training times. However, both models converge in higher values of accuracy (and lower KL-divergence and MSE) when compared to the remaining encodings. Notably, training the musical embedding from ABC notation requires a longer time per epoch compared to the others.

During the process of extracting information from the original symbolic music, the Pitch DFT exhibits the slowest performance by a substantial margin, taking almost four times longer than the second-worst performing method, the DFT of pitch class. In contrast, the remaining musical encodings require nearly equal amounts of time for feature extraction, with the piano roll achieving the best performance by less than five seconds.

Additionally, we observed that augmenting the piano roll and MIDI-like embeddings is a simple and fast task. However, the process of augmenting the DFTs, Tonnetz, and ABC encodings is slow, particularly for the latter two. To address this, we pre-computed and saved the augmentations prior to training, allowing us to load them during the training preparation phase.

5.2.2. Latent Space Analysis

Table 2 presents the average and standard deviation values for the cluster and key-distance metrics, allowing us to evaluate the alignment between the latent space of each musical encoding and cognitive spaces.

The latent space trained on Pitch DFT encoding demonstrates the best alignment with cognitive spaces, surpassing other encodings in all three metrics, as anticipated due to its high accuracy and KL-divergence values, and in line with the preference for signal-based encoding concluded in Prang & Esling, (2021). It presents the most condensed and seamless representation in the latent space,

Musical Encoding	Davis-Bouldin Score	Dunn Index	Kendall's Tau
Piano roll	32.8 ± 17.8	.0005 ± .0005	.11 ± .08
MIDI-like	55.9 ± 35.3	.0001 ± .0001	.15 ± .14
ABC	66.9 ± 55.0	.00003 ± .00003	.11 ± .09
Tonnetz	33.0 ± 21.6	.0006 ± .0007	.11 ± .08
Pitch Class DFT	37.3 ± 59.1	.0006 ± .0006	.11 ± .08
Pitch DFT	8.1 ± 5.5	.0008 ± .0009	.44 ± .32

Table 2: Average and standard deviation values for the cluster and key-distance metrics

outperforming other encodings by a large margin. The piano roll, Tonnetz, and Pitch Class DFT encodings also show relatively good results, while the MIDI-like and ABC encodings have the worst scores in all three metrics, indicating poorer clustering performance.

Upon analyzing Pitch DFT, there are several noteworthy insights. First, around 16% of chorales display Kendall Tau's absolute values exceeding .9, indicating a strong alignment with the cognitive pitch space. Moreover, major key chorales' latent spaces appear to be more efficient in capturing these distances, with 87% of chorales that have Kendall Tau's absolute values greater than .9 being in a major key. Surprisingly, the findings also reveal that longer chorales are better at capturing these distances. About 30% of chorales containing over 114 slices show Tau absolute values greater than .9, in contrast to only 10% of chorales with less than 64 slices exhibiting such values. Based on conventional assumptions, we would anticipate that longer chorales, which are more susceptible to modulations, would have latent spaces that are more ambiguous and, therefore, less aligned with the cognitive pitch space.

6. Conclusions and Future Work

Our paper explores the performance of VAEs in reconstructing tonal symbolic music and eliciting latent representations of cognitive and musical theoretical value. We trained VAEs on a prototypical tonal music corpus of 371 Bach's chorales, represented as six different symbolic music encodings (i.e., Piano roll, MIDI, ABC, Tonnetz, DFT of pitch and pitch class distributions) and evaluated the degree to which the latent spaces defined by the different VAE corpus encodings align with cognitive distances from musical pitch, based on objective reconstruction performance metrics (accuracy, MSE, and KL-divergence), computational performance, and clustering metrics (Davis-Bouldin Score, Dunn Index, and Kendall's Tau). Our VAE implementation and the encodings framework are available online at <https://github.com/NadiaCarvalho/Latent-Tonal-Music>. Furthermore, on our website <https://nadiacarvalho.github.io/Latent-Tonal-Music/>,

we offer an engaging platform for users to explore the encodings through various symbolic music compositions.

The results showed that the ABC VAE performed best in the data reconstruction performance metrics, while the proposed Pitch DFT VAE latent space is better aligned with a common-tone space where overlapping objects within a key are fuzzy clusters, which impose a well-defined order of structural significance or stability, i.e., a tonal hierarchy. In sum, ABC encodings would be preferable when there is an interest in preserving the original symbolic musical structures, while Pitch DFT VAEs can produce more diverse and varied generative models. Moving forward, we plan to conduct a more in-depth analysis of the data reconstructed using this encoding.

The findings suggest potential for exploring pitch spaces in less structured harmonic or pitch systems, such as modal and microtonal music. While many existing pitch spaces accurately represent the distances across various hierarchies (e.g., pitches, chords, and keys), there are currently no such spaces available for non-tonal music expressions, as far as we know.

Acknowledgments

This research has been funded by the Portuguese National Funding Agency for Science, Research and Technology [2021.05132.BD].

Ethics Statement

This research study did not involve human subjects, animals, or sensitive data. Therefore, no ethics approval was required. The authors have no conflicts of interest to declare.

References

- Amiot, Emmanuel. 2016. *Music Through Fourier Space. Computational Music Science*. Springer International Publishing. <https://doi.org/10.1007/978-3-319-45581-5>.
- Bernardes, Gilberto, Diogo Cocharro, Marcelo Caetano, Carlos Guedes, and Matthew E.P. Davies. 2016. “A Multi-Level Tonal Interval Space for Modelling Pitch Relatedness and Musical Consonance.” *Journal of New Music Research* 45 (4): 281–94. <https://doi.org/10.1080/09298215.2016.1182192>.
- Briot, Jean-Pierre, Gaëtan Hadjeres, and François-David Pachet. 2017. “Deep Learning Techniques for Music Generation -- A Survey.” arXiv. <https://doi.org/10.48550/ARXIV.1709.01620>.
- Briot, Jean-Pierre, Gaëtan Hadjeres, and François-David Pachet. 2020. *Deep Learning Techniques for Music Generation. Computational Synthesis and Creative Systems*. Springer International Publishing. <https://doi.org/10.100>

7/978-3-319-70163-9.

- Briot, Jean-Pierre, and François Pachet. 2018. “Deep Learning for Music Generation: Challenges and Directions.” *Neural Computing and Applications* 32 (4): 981–93. <https://doi.org/10.1007/s00521-018-3813-6>.
- Bryan-Kinns, Nick, Berker Banar, Corey Ford, Courtney N. Reed, Yixiao Zhang, Simon Colton, and Jack Armitage. 2021. “Exploring XAI for the Arts: Explaining Latent Space in Generative Music.” In *eXplainable AI Approaches for Debugging and Diagnosis*. https://openreview.net/forum?id=GLhY_0xMLZr.
- Chuan, Ching-Hua, and Dorien Herremans. 2018. “Modeling Temporal Tonal Relations in Polyphonic Music Through Deep Networks with a Novel Image-Based Representation.” In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence*. AAAI’18/IAAI’18/EAAI’18. New Orleans, Louisiana, USA: AAAI Press.
- Cohn, Richard. 1997. “Neo-Riemannian Operations, Parsimonious Trichords, and Their ‘Tonnetz’ Representations.” *Journal of Music Theory* 41 (1): 1. <https://doi.org/10.2307/843761>.
- Cohn, Richard. 1998. “Introduction to Neo-Riemannian Theory: A Survey and a Historical Perspective.” *Journal of Music Theory* 42 (2): 167. <https://doi.org/10.2307/843871>.
- Davies, David L., and Donald W. Bouldin. 1979. “A Cluster Separation Measure.” *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-1 (2): 224–27. <https://doi.org/10.1109/tpami.1979.4766909>.
- Deutsch, Diana. 1984. “Two Issues Concerning Tonal Hierarchies: Comment on Castellano, Bharucha, and Krumhansl.” *Journal of Experimental Psychology: General* 113 (3): 413–16. <https://doi.org/10.1037/0096-3445.113.3.413>.
- Dunn†, J. C. 1974. “Well-Separated Clusters and Optimal Fuzzy Partitions.” *Journal of Cybernetics* 4 (1): 95–104. <https://doi.org/10.1080/01969727408546059>.
- Euler, Leonhard. 1739. *Tentamen Novae Theoriae Musicae: Ex Certissimis Harmoniae Principiis Dilucide Expositae*. ex typographia Academiae scientiarum.
- Fisher, N. I., and A. J. Lee. 1982. “Nonparametric Measures of Angular-Angular Association.” *Biometrika* 69 (2): 315. <https://doi.org/10.2307/2335405>.
- Guo, Z., J. Kang, and D. Herremans. 2022. “A Domain-Knowledge-Inspired Music Embedding Space and a Novel Attention Mechanism for Symbolic Music Modeling.” arXiv. <https://doi.org/10.48550/ARXIV.2212.00973>.
- Kim, Seung-Goo. 2022. “On the Encoding of Natural Music in Computational Models and Human Brains.” *Frontiers in Neuroscience* 16 (September). <https://doi.org/10.3389/fnins.2022.928841>.
- Kingma, Diederik P, and Max Welling. 2013. “Auto-Encoding Variational Bayes.” arXiv. <https://doi.org/10.48550/ARXIV.1312.6114>.
- Krumhansl, Carol L. 1990. *Cognitive Foundations of Musical Pitch*. Oxford Psychology Series, no. 17. New York: Oxford University Press.

- Lewin, David. 1987. *Generalized Musical Intervals and Transformations*. New Haven: Yale University Press.
- Longuet-Higgins, H. C. 1987. *Mental Processes: Studies in Cognitive Science*. Explorations in Cognitive Science 1. Cambridge, Mass: MIT Press.
- Mezza, Alessandro Ilic, Massimiliano Zanoni, and Augusto Sarti. 2023. “A Latent Rhythm Complexity Model for Attribute-Controlled Drum Pattern Generation.” *EURASIP Journal on Audio, Speech, and Music Processing* 2023 (1). <https://doi.org/10.1186/s13636-022-00267-2>.
- Moss, Fabian C., Markus Neuwirth, and Martin Rohrmeier. 2022. “The Line of Fifths and the Co-Evolution of Tonal Pitch-Classes.” *Journal of Mathematics and Music* 17 (2): 173–97. <https://doi.org/10.1080/17459737.2022.2044927>.
- Nardelli, Marco Buongiorno, Garland Culbreth, and Miguel Fuentes. 2022. “Towards A Measure Of Harmonic Complexity In Western Classical Music.” *Advances in Complex Systems (ACS)* 25 (05n06): 1–11. <https://doi.org/10.1142/S0219525922400082>.
- Ore, Sageev, Ian Simon, Sander Dieleman, Douglas Eck, and Karen Simonyan. 2018. “This Time with Feeling: Learning Expressive Musical Performance.” *Neural Computing and Applications* 32 (4): 955–67. <https://doi.org/10.1007/s00521-018-3758-9>.
- Pearson, Karl. 1901. “LIII. On Lines and Planes of Closest Fit to Systems of Points in Space.” *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2 (11): 559–72. <https://doi.org/10.1080/14786440109462720>.
- Plitsis, Manos, Kosmas Kritsis, Maximus Kaliakatsos-Papakostas, Aggelos Pikrakis, and Vassilis Katsouros. 2020. “Towards a Classification and Evaluation of Symbolic Music Encodings for RNN Music Generation.” *Zenodo*, October. <https://doi.org/10.5281/ZENODO.4285410>.
- Prang, Mathieu, and Philippe Esling. 2021. “Signal-Domain Representation of Symbolic Music for Learning Embedding Spaces.” *arXiv*. <https://doi.org/10.48550/ARXIV.2109.03454>.
- Prokhorov, Victor, Ehsan Shareghi, Yingzhen Li, Mohammad Taher Pilehvar, and Nigel Collier. 2019. “On the Importance of the Kullback-Leibler Divergence Term in Variational Autoencoders for Text Generation.” *arXiv*. <https://doi.org/10.48550/ARXIV.1909.13668>.
- Qiu, Lvyang, Shuyu Li, and Yunsick Sung. 2021. “3D-DCDAE: Unsupervised Music Latent Representations Learning Method Based on a Deep 3D Convolutional Denoising Autoencoder for Music Genre Classification.” *Mathematics* 9 (18): 2274. <https://doi.org/10.3390/math9182274>.
- Quinn, Ian. 2006. “General Equal-Tempered Harmony (Introduction and Part I).” *Perspectives of New Music* 44 (2): 114–58. <https://doi.org/10.1353/pnm.2006.0010>.
- Roberts, Adam, Jesse Engel, Colin Raffel, Curtis Hawthorne, and Douglas Eck. 2018. “A Hierarchical Latent Vector Model for Learning Long-Term Structure in Music.” *arXiv*. <https://doi.org/10.48550/ARXIV.1803.05428>.
- Sarmiento, Pedro, Adarsh Kumar, Yu-Hua Chen, CJ Carr, Zack Zukowski, and Mathieu Barthet. 2023. “GTR-CTRL: Instrument and Genre Conditioning

- for Guitar-Focused Music Generation with Transformers.” *arXiv*. <https://doi.org/10.48550/ARXIV.2302.05393>.
- Shepard, Roger N. 1982. “Geometrical Approximations to the Structure of Musical Pitch.” *Psychological Review* 89 (4): 305–33. <https://doi.org/10.1037/0033-295x.89.4.305>.
- Sturm, Bob L., Oded Ben-Tal, Úna Monaghan, Nick Collins, Dorien Herremans, Elaine Chew, Gaëtan Hadjeres, Emmanuel Deruty, and François Pachet. 2018. “Machine Learning Research That Matters for Music Creation: A Case Study.” *Journal of New Music Research* 48 (1): 36–55. <https://doi.org/10.1080/09298215.2018.1515233>.
- Turker, Meliksah, Alara Dirik, and Pinar Yanardag. 2022. “MIDISpace: Finding Linear Directions in Latent Space for Music Generation.” In *Creativity and Cognition*. ACM. <https://doi.org/10.1145/3527927.3532790>.
- Tymoczko, Dmitri. 2010. *A Geometry of Music: Harmony and Counterpoint in the Extended Common Practice*. Oxford Studies in Music Theory. New York: Oxford University Press.
- Ucar, Talip. 2019. “Bridging the ELBO and MMD.” *arXiv*. <https://doi.org/10.48550/ARXIV.1910.13181>.
- Weber, Gottfried. 1817-1821. *Versuch Einer Geordneten Theorie Der Tonsetzkunst*. Vol. 1. B. Schott’s Söhne.
- Yust, Jason. 2015. “Applications of DFT to the Theory of Twentieth-Century Harmony.” In *Mathematics and Computation in Music*, 207–18. Springer International Publishing. https://doi.org/10.1007/978-3-319-20603-5_22.