

# Creating compositions in the style of Henryk Wieniawski using Generative AI

Reuben Andrew Ethan Menezes

Received May 27, 2025

Accepted September 6, 2025

Electronic access November 30, 2025

Music is a fundamental part of various shared human experiences. In this paper, we aim to contribute to enhancing this experience by leveraging rapidly progressing AI technology to create music. Specifically, we recreate music in the style of Henryk Wieniawski using a custom-designed LSTM Sequential Generative AI Neural Network architecture. The AI generated music successfully mimicked Wieniawski's style and created an arguably present impression. This work demonstrates how the influence of composers can persist beyond the grave, allowing their music to be enjoyed and intertwined with various modern styles and trends. It paves the steps for how AI generated and human made music can one day be indistinguishable.

## Introduction

Throughout ages, music has been closely bonded with humans. It's an intrinsic part of our shared human experience. In this new age of AI, this experience can be transferred to machines. This paper intends to explore that by looking at a prolific 19th century composer Henryk Wieniawski<sup>1</sup>. This was as his music had been an integral part of the author's journey of learning and playing music.

Since music composition is sequential in nature<sup>2-8</sup> and depends on continuing the pattern of coherence created by the previous notes, we decided that our architecture should follow this criterion. Hence we decided to use a generative sequence neural network.

Sequence models have made significant strides in the 2020s, with transformer models<sup>9</sup> such as MuseNet taking the world by storm. However transformer models are complex to work with and expensive to train. Hence we use a simpler architecture: Long Short Term Memory (LSTMs)<sup>10-12</sup>. LSTMs are a reasonably advanced and capable Recurrent Neural Network architecture that is capable of handling our sequential musical data: they can retain context over long durations and are capable of making predictions based on this past information.

Thus LSTMs are suitable for our use case with handling sequential data and creating music in the style of Henryk Wieniawski.

## Methods

For reasons of simplicity and practicality, we chose to implement our LSTM system using TensorFlow<sup>13</sup>. Particularly since much of the relevant Gen AI literature that we based our work on used TensorFlow.

## Discussion of the Dataset

Our input data was in the MIDI standard. The MIDI format is a convenient data for presenting music because of its small file size, ease of editing, and wide compatibility. MIDI files store instructions for music playback, including note pitches and timings.

We used the broad Maestro dataset<sup>14</sup>, which consists of 200 hours of paired audio and MIDI recordings from ten years of International Piano-e-Competition., to pretrain the model in order to teach it basic musical theory, alongside a smaller dataset consisting of 8 handpicked compositions for violin by Henryk Wieniawski, which are used for the purpose of finetuning the model towards his style of composition. The Maestro dataset is ~ 201.2 hours of music long, while the Wieniawski one is ~ 1.1 hours. The pieces we used are shown in Table 1.

Wieniawski Violin Concerto Number 2
Wieniawski Violin Caprice no.18
Scherzo-Tarentelle Op.16
L'école moderne, Op.10
Wieniawski Op.15 Variations On An Original Theme
Wieniawski Etudes-Caprices 1-4 from Opus 18
Wieniawski Violin Concerto 1 Movement 1

Table 1: These pieces were sourced from MuseScore: the world's largest online sheet music catalog. This was done for the sake of simplicity and because the author had experience with the platform.

The input MIDI data was transformed into a set of values (duration, pitch, and step) using `pretty_midi`<sup>15</sup>, a Python library which provides an interface for dealing with the MIDI standard.

This is a sensible embedding for both man and machine: it represents the contents of a musical sheet a human would be reading while learning how to play a piece, simultaneously capturing the core elements that a machine would require and need to understand as well. It might not necessarily be the optimal embedding possible, but it is a good starting point nevertheless.

With regards to processing the data, our Neutral Network architecture consisted of 2 LSTM layers followed by a dense layer. The purpose of the first layer is to capture the micro structure of the music i.e: short-range features and the logical flow of notes<sup>4</sup>. The second layer was meant to be a layer that captured the overarching nature of the piece, such as the mood of the melody and compositional styles<sup>4</sup>. The final dense layer takes the output of the 2nd layer and converts this into a musical output. To avoid overfitting of the model, dropout layers<sup>16</sup> were incorporated, as shown in the visualisation of the entire model in Figure 1.

Our loss function L was chosen such that the machine would focus on capturing the pitch, step and duration of the notes in a balanced way:

$$L = a \cdot L_{ce}(\text{pitch}) + b \cdot L_{mse}(\text{duration}) + c \cdot L_{mse}(\text{step})$$

with L being the cross-entropy loss, L\_mse being the mean squared error loss and a,b,c human-chosen weights.

This architecture is capable of taking the embedded data and processing it in a systematic and sensible manner.

We pretrained our model until L converged for the base Maestro dataset for different values of a, b and c hyperparameters. Each training run 24 epochs, at a batch-size of 45 minutes of music. Our evaluation of the model performance consisted of listening to the output manually instead of seeking to minimise the loss on some train/test data: this is because finding the absolute minimum of our (or, to our knowledge, any) quantitative loss function would not necessarily represent an ideal, objective human impression of the music. Quantifying how a human experiences music through a loss function is a highly challenging task indeed<sup>17</sup> and well beyond the scope of this work. In the end, a favourite (a,b,c) combination was determined [see appendix] and fixed.

Following the creation of the pretrained NN, we incorporated Wieniawski's musical style through transfer learning. This was done by allowing for select layers of the NN to be retrained on the 7 Wieniawski compositions. Three configurations were explored: only dense1 layer trainable, dense1 and lstm2 layers, and finally dense1, lstm2 and lstm1. Each retraining was done using 50 epochs, each time starting from the pretrained model.

## Results / Discussion

The output of our generative NN were four musical MIDI files: the output of pretraining on Maestro, along with the results of the three transfer learning combinations.

Layer (type)	Output Shape	Param #	Connected to
input_layer_3 (InputLayer)	(None, 100, 3)	0	-
lstm1 (LSTM)	(None, 100, 128)	67,584	input_layer_3[0]...
lstm2 (LSTM)	(None, 128)	131,584	lstm1[0][0]
dense1 (Dense)	(None, 128)	16,512	lstm2[0][0]
dropout_3 (Dropout)	(None, 128)	0	dense1[0][0]
duration (Dense)	(None, 1)	129	dropout_3[0][0]
pitch (Dense)	(None, 128)	16,512	dropout_3[0][0]
step (Dense)	(None, 1)	129	dropout_3[0][0]

Fig. 1: Our architecture was 2 LSTM layers (lstm1 and lstm2) followed by a Dense layer (dense1). We added a dropout layer to correct overtraining.

We first discuss the pretraining output, where all layers were trained on Maestro. The audio is characteristic of a piano due to the Maestro dataset being a collection of piano recordings (in contrast to the violin recordings of our Wieniawski). We regarded the audio to be mildly incoherent. In our opinion, it does not sound as pleasing as the other three outputs. Furthermore, we make the obvious observation that this audio displays no similarity to Wieniawski because it was trained on a dataset which included no pieces by the composer.

Now we discuss the outputs from the transfer-learning on Wieniawski's 7 compositions.

The output when only the dense1 layer was trainable suffers from a persistently low note, although the audio is coherent and structured. In other words, it is a sort of composition. However, the sections of the composition, i.e. the melodies, do not blend well with each other, making the overall piece sound subpar. This is reminiscent of the "stream of thought" issue that overly simple sequence models are known to suffer from, where parts of sentences or even whole sentences make sense, while the actual paragraph does not<sup>18</sup>.

When the dense1 and lstm2 layers are trainable, the audio is still both slightly incoherent in some parts and suffers from a persistent low note, but to a lesser extent than when just the dense1 layer was trainable. There is variance in the tempo of the output but none in terms of dynamics. The sequences of notes follow each other logically in the short-term, but, again, not in the long-term.

Lastly, when enabling the training of the dense1, lstm2 and lstm1 layers, the output becomes generally coherent and the sequence of notes most logical. There are a few characteristic elements of Wieniawski such as string crossings and alternations between low and high pitched notes<sup>19</sup>. More musical elements such as variation in tempo are noticeable: suggesting that the lstm1 layer was a useful addition. Something resembling a motif can be observed: there are two melodies between which the output alternates. One issue with this output is that the

ending is abrupt, unlike most classical music pieces which slow down - or become softer - towards the end.

We found the final output to sound the best out of the four. We speculate this could be due to the neural network having a hot-start from the general Maestro datasets, and then properly fine-trained on 7 musical pieces that are consistent and much more correlated with each other than the disparate music in Maestro.

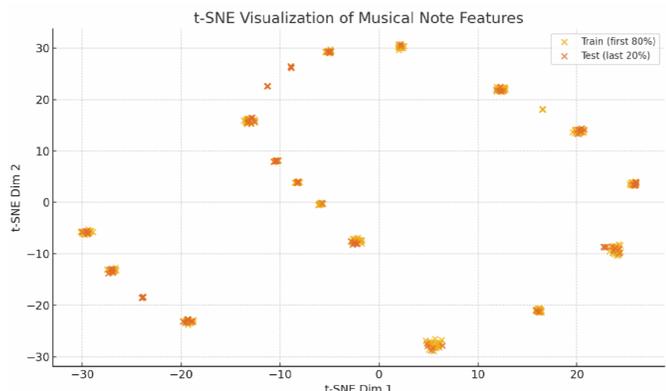


Fig. 2: tSNE curve with train/test split of 80% and 20%

The t-SNE plot shows clusters of notes with similar pitch (tone), duration (length), and step (time between notes), revealing structured musical patterns that the LSTM can learn and generalize from, with blue (train) and orange (test) dots occupying different regions, highlighting a shift in musical style. Figure 3 visualizes the model’s learning dynamics across 50

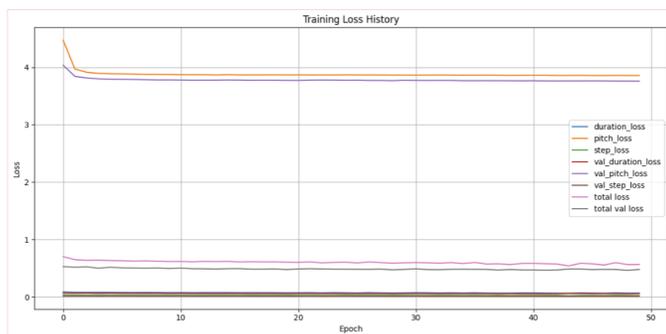


Fig. 3: Training and Validation Loss Curves for Musical LSTM Model over 50 Epochs

epochs. It shows the individual component losses (pitch\_loss, step\_loss, duration\_loss) and their counterparts on the validation set (val\_pitch\_loss, etc.). The total loss (a weighted sum of the three outputs) is also shown for both training and validation, renamed here as "total loss" and "total val loss" for clarity. The losses generally exhibit convergence over time, with relatively small generalization gaps.

## Conclusion

To summarise, we have here managed to create music, a deeply human pursuit, in the style of Wieniawski using AI. Such AIs can help a composer’s musical style continue to live on even after their death; in a sense, even resurrecting them in the form of an AI model. While our output was coherent, it lacked the true complexity and flair of Wieniawski. In the future, our model can be iterated on by using transformer layers instead of LSTMs, adjusting the loss function, improving data quality and size, and formulating an automated cross-validation scheme to minimise human involvement in evaluating our model performance. In the future, we envision further advances in generative AI for music generation<sup>20</sup>, allowing great renditions of the historical composers and their works<sup>21</sup>, and even novel compositions that will strike at the heart of humans<sup>22</sup>.

## Appendix: Source Code

Here is the code used for experiments:

```

1 !pip install pretty_midi
2
3 # Library Installations
4 import os.path
5 import tensorflow as tf
6 from tensorflow.keras.layers import Input,
7     LSTM, Dense, Dropout
8 from tensorflow.keras.models import Model
9 from tensorflow.keras.losses import
10     SparseCategoricalCrossentropy
11 from tensorflow.keras.optimizers.schedules
12     import PiecewiseConstantDecay
13 from tensorflow.keras.optimizers import
14     Adam
15
16 from typing import Dict, List, Optional,
17     Sequence, Tuple
18 import collections
19 import datetime
20 import glob
21 import numpy as np
22 import pathlib
23 import pandas as pd
24 import pretty_midi
25 from matplotlib import pyplot as plt
26
27 # Hyperparameters
28 seq_length = 100
29 vocab_size = 500
30 batch_size = 256
31 epochs = 1
32 temperature = 2.0
33 num_predictions = 400
34 # Pre-train vs fine-tune flag

```

```

30 PRETRAIN = False
31 learning_rate = 0.003 if PRETRAIN else
    0.001
32 # Dataset Prep
33 seed = 42
34 tf.random.set_seed(seed)
35 np.random.seed(seed)
36
37
38 # Setting the path and loading the data
39
40
41 if PRETRAIN:
42     data_dir = pathlib.Path('maestro-v2
        .0.0-midi/maestro-v2.0.0')
43     filenames = glob.glob(str(data_dir/'
        **/*.mid*'))
44 else:
45     data_dir = pathlib.Path('finetrained')
46     filenames = glob.glob(str(data_dir/'*.
        mid*'))
47 print('Number of files:', len(filenames))
48
49
50 # analyzing and working with a sample file
51 try:
52     sample_file = filenames[1]
53 except:
54     sample_file = filenames[0]
55 pm = pretty_midi.PrettyMIDI(sample_file)
56 instrument = pm.instruments[0]
57 instrument_name = pretty_midi.
    program_to_instrument_name(instrument.
    program)
58
59
60
61
62 # Extracting the notes
63 for i, note in enumerate(instrument.notes
    [:10]):
64     note_name = pretty_midi.
        note_number_to_name(note.pitch)
65     duration = note.end - note.start
66
67
68
69
70 # Extracting the notes from the sample MIDI
    file
71
72
73 def midi_to_notes(midi_file: str) -> pd.
    DataFrame:
74     pm = pretty_midi.PrettyMIDI(midi_file)
75     instrument = pm.instruments[0]

```

```

76     notes = collections.defaultdict(list)
77
78
79     # Sort the notes by start time
80 sorted_notes = sorted(instrument.notes, key
    =lambda note: note.start)
81     prev_start = sorted_notes[0].start
82
83
84     for note in sorted_notes:
85         start = note.start
86         end = note.end
87         notes['pitch'].append(note.pitch)
88         notes['start'].append(start)
89         notes['end'].append(end)
90         notes['step'].append(start -
            prev_start)
91         notes['duration'].append(end -
            start)
92         prev_start = start
93
94
95     return pd.DataFrame({name: np.array(
        value) for name, value in notes.
        items()})
96
97
98 raw_notes = midi_to_notes(sample_file)
99 raw_notes.head()
100
101
102 # Converting to note names by considering
    the respective pitch values
103
104
105 get_note_names = np.vectorize(pretty_midi.
    note_number_to_name)
106 sample_note_names = get_note_names(
    raw_notes['pitch'])
107
108
109
110
111 # Visualizing the parameters of the musical
    notes of the piano
112 def plot_piano_roll(notes: pd.DataFrame,
    count: Optional[int] = None):
113     if count:
114         title = f'First {count} notes'
115     else:
116         title = f'Whole track'
117         count = len(notes['pitch'])
118
119
120     plt.figure(figsize=(20, 4))
121     plot_pitch = np.stack([notes['pitch

```

```

122         ], notes['pitch']], axis=0)
123     plot_start_stop = np.stack([notes['
124         start'], notes['end']], axis=0)
125
126     plt.plot(plot_start_stop[:, :count
127             ], plot_pitch[:, :count], color=
128             "b", marker=".")
129     plt.xlabel('Time [s]')
130     plt.ylabel('Pitch')
131     _ = plt.title(title)
132
133 # Training Data Creation
134 def notes_to_midi(notes: pd.DataFrame,
135                 out_file: str, instrument_name: str,
136                 velocity: int = 100) ->
137     pretty_midi.PrettyMIDI
138     :
139
140     pm = pretty_midi.PrettyMIDI()
141     instrument = pretty_midi.Instrument(
142         program=pretty_midi.
143         instrument_name_to_program(
144             instrument_name))
145
146     prev_start = 0
147     for i, note in notes.iterrows():
148         start = float(prev_start + note['
149             step'])
150         end = float(start + note['duration'
151             ])
152
153         note = pretty_midi.Note(velocity=
154             velocity, pitch=int(note['pitch'
155             ]),
156             start=start
157             , end=
158             end)
159
160         instrument.notes.append(note)
161         prev_start = start
162
163     pm.instruments.append(instrument)
164     pm.write(out_file)
165     return pm
166
167 example_file = 'example.midi'
168 example_pm = notes_to_midi(
169     raw_notes, out_file=example_file,

```

```

164         instrument_name=instrument_name)
165
166 num_files = 5
167 all_notes = []
168 for f in filenames[:num_files]:
169     notes = midi_to_notes(f)
170     all_notes.append(notes)
171
172
173 all_notes = pd.concat(all_notes)
174
175 n_notes = len(all_notes)
176
177
178 key_order = ['pitch', 'step', 'duration']
179 train_notes = np.stack([all_notes[key] for
180     key in key_order], axis=1)
181
182 notes_ds = tf.data.Dataset.
183     from_tensor_slices(train_notes)
184 notes_ds.element_spec
185
186 def create_sequences(dataset: tf.data.
187     Dataset, seq_length: int,
188     vocab_size = 128) ->
189     tf.data.Dataset:
190     """Returns TF Dataset of sequence and
191     label examples."""
192     seq_length = seq_length+1
193
194     # Take 1 extra for the labels
195     windows = dataset.window(seq_length,
196         shift=1, stride=1,
197         drop_remainder
198         =True)
199
200     # `flat_map` flattens the" dataset of
201     datasets" into a dataset of tensors
202     flatten = lambda x: x.batch(seq_length,
203         drop_remainder=True)
204     sequences = windows.flat_map(flatten)
205
206     # Normalize note pitch
207     def scale_pitch(x):
208         x = x/[vocab_size,1.0,1.0]
209         return x
210
211     # Split the labels
212     def split_labels(sequences):

```

```

210     inputs = sequences[:-1]
211     labels_dense = sequences[-1]
212     labels = {key:labels_dense[i] for i
213               ,key in enumerate(key_order)}
214
215     return scale_pitch(inputs), labels
216
217
218     return sequences.map(split_labels,
219                          num_parallel_calls=tf.data.AUTOTUNE)
220
221 seq_ds = create_sequences(notes_ds,
222                          seq_length, vocab_size)
223
224
225
226
227
228 buffer_size = n_notes - seq_length # the
229 number of items in the dataset
230 train_ds = (seq_ds
231             .shuffle(buffer_size)
232             .batch(batch_size,
233                   drop_remainder=True)
234             .cache()
235             .prefetch(tf.data.experimental.
236                       AUTOTUNE))
237
238 # Making the model
239 def mse_with_positive_pressure(y_true: tf.
240                               Tensor, y_pred: tf.Tensor):
241     mse = (y_true - y_pred) ** 2
242     positive_pressure = 10 * tf.maximum(-
243     y_pred, 0.0)
244     return tf.reduce_mean(mse +
245     positive_pressure)
246
247 # Developing the model
248
249 input_shape = (seq_length, 3)
250
251 inputs = Input(input_shape)
252 # x = LSTM(128,
253 #         return_sequences=False,
254 #         dropout=0.2, #
255 #         Dropout for inputs
256 #         recurrent_dropout=0.2 #
257 #         Dropout for recurrent state
258 #         )(inputs)

```

```

255
256
257
258
259
260
261 x = LSTM(128, return_sequences=True, name='
262 lstm1',
263         dropout=0.2, recurrent_dropout
264         =0.2)(inputs)
265 # Second LSTM layer (to be trained)
266 x = LSTM(128, name='lstm2',
267         dropout=0.2, recurrent_dropout
268         =0.2)(x)
269 x = Dense(128, name='dense1', activation = '
270 relu')(x)
271 x = Dropout(0.2)(x)
272 outputs = {'pitch': Dense(128, name='pitch'
273 ) (x),
274           'step': Dense(1, name='step')(x),
275           'duration': Dense(1, name='
276 duration')(x),
277           }
278 model = Model(inputs, outputs)
279
280 if PRETRAIN:
281     if os.path.isfile("
282 pretrained_musical_lstm.weights.h5")
283 :
284         model.load_weights("
285 pretrained_musical_lstm.weights.
286 h5")
287 else:
288     if os.path.isfile("
289 finetrained_musical_lstm.weights.h5"
290 ):
291         model.load_weights("
292 finetrained_musical_lstm.weights
293 .h5")
294 else:
295         model.load_weights("
296 pretrained_musical_lstm.weights.
297 h5")
298 loss = {'pitch':
299 SparseCategoricalCrossentropy(
300 from_logits=True),
301        'step': mse_with_positive_pressure,
302        'duration':
303        mse_with_positive_pressure,
304        }
305
306 optimizer = Adam(learning_rate=
307 learning_rate)
308 model.compile(loss=loss, optimizer=

```

```

optimizer)
291 model.summary()
292
293
294 # Compiling and fitting the model
295 model.compile(loss = loss,
296               loss_weights = {'pitch':
297                               0.05, 'step': 4.0, '
298                               duration':4.0,},
299               optimizer = optimizer)
300 history = model.fit(train_ds, epochs=epochs
301                    )
302
303 if PRETRAIN: model.save_weights("
304               pretrained_musical_lstm.weights.h5")
305 else: model.save_weights("
306               finetuned_musical_lstm.weights.h5")
307
308 # Freezing the first two LSTM layers in
309 # case of PRETRAIN being false
310 if not PRETRAIN:
311     model.get_layer('lstm1').trainable =
312         False
313     model.get_layer('lstm2').trainable =
314         True
315     # after 100 epoch, save then finetuned
316     # file as "finetuned both layers
317     # untrainable"
318     # then set lstm2 trainable to true and
319     # train for 100 more and then see how
320     # it differs. save that as finetuned
321     # last layer trainable
322
323 if PRETRAIN: model.load_weights("
324               pretrained_musical_lstm.weights.h5")
325 else: model.load_weights("
326               finetuned_musical_lstm.weights.h5")
327 # Generating the notes
328 def predict_next_note(notes: np.ndarray,
329                       keras_model: tf.keras.Model,
330                       temperature: float =
331                           1.0) -> int:
332     """Generates a note IDs using a trained
333     sequence model."""
334
335     assert temperature > 0

```

```

328
329
330 # Add batch dimension
331 inputs = tf.expand_dims(notes, 0)
332
333
334 predictions = model.predict(inputs)
335 pitch_logits = predictions['pitch']
336 step = predictions['step']
337 duration = predictions['duration']
338
339
340 pitch_logits /= temperature
341 pitch = tf.random.categorical(
342     pitch_logits, num_samples=1)
343 pitch = tf.squeeze(pitch, axis=-1)
344 duration = tf.squeeze(duration, axis
345                       =-1)
346 step = tf.squeeze(step, axis=-1)
347
348 # 'step' and 'duration' values should
349 # be non-negative
350 step = tf.maximum(0, step)
351 duration = tf.maximum(0, duration)
352
353 return int(pitch), float(step), float(
354           duration)
355
356 sample_notes = np.stack([raw_notes[key] for
357                           key in key_order], axis=1)
358
359 # The initial sequence of notes while the
360 # pitch is normalized similar to training
361 # sequences
362 input_notes = (
363     sample_notes[:seq_length] / np.array([
364         vocab_size, 1, 1]))
365
366 generated_notes = []
367 prev_start = 0
368
369 for _ in range(num_predictions):
370     pitch, step, duration =
371         predict_next_note(input_notes, model
372                           , temperature)
373     start = prev_start + step
374     end = start + duration
375     input_note = (pitch, step, duration)
376     generated_notes.append((*input_note,
377                             start, end))

```

```

373     input_notes = np.delete(input_notes, 0,
374                             axis=0)
375     input_notes = np.append(input_notes, np
376                             .expand_dims(input_note, 0), axis=0)
377     prev_start = start
378 generated_notes = pd.DataFrame(
379     generated_notes, columns=(*key_order, '
380     start', 'end'))
381
382 generated_notes.head(10)
383
384
385 out_file = 'output.midi'
386 out_pm = notes_to_midi(
387     generated_notes, out_file=out_file,
388     instrument_name=instrument_name)

```

- 14 C. Hawthorne, A. Stasyuk, A. Roberts, I. Simon, C.-Z. Huang, S. Dieleman, E. Elsen, J. Engel and D. Eck, Proceedings of the 35th International Conference on Machine Learning (ICML 2018), p. 710–719.
- 15 C. Raffel and D. Ellis, Proceedings of the 15th International Society for Music Information Retrieval Conference (Late-Breaking and Demo Papers, ISMIR).
- 16 N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever and R. Salakhutdinov, *Journal of Machine Learning Research*, **15**, 1929–1958.
- 17 H. Zhang, J. Liang, H. Phan, W. Wang and E. Benetos, From aesthetics to human preferences: Comparative perspectives of evaluating text-to-music systems. arXiv.
- 18 J. Guan, X. Mao, C. Fan, Z. Liu, W. Ding and M. Huang, Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), p. 6379–6393.
- 19 P. Kozak, *Pedagogical examination of Henryk Wieniawski's L'école moderne Opus 10*.
- 20 Y. Zhang, X. Li, Z. Wang and Y. Chen, *Applications and Advances of Artificial Intelligence in Music Generation: A Systematic Review*, arXiv preprint arXiv:2409.03715.
- 21 M. Kong and L. Huang, *Bach Style Music Authoring System based on Deep Learning*, arXiv preprint arXiv:2110.02640.
- 22 K. Agres, A. Dash and P. Chua, *AffectMachine-Classical: A novel system for generating affective classical music*, arXiv preprint arXiv:2304.04915.

## References

- 1 *Encyclopædia Britannica*.
- 2 C.-Z. Huang, A. Vaswani, J. Uszkoreit, N. Shazeer, I. Simon, C. Hawthorne, A. Dai, M. Hoffman, M. Dinculescu and D. Eck, Proceedings of the 7th International Conference on Learning Representations (ICLR).
- 3 A. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, K. Kavukcuoglu and A. Oord, *WaveNet: A Generative Model for Raw Audio*.
- 4 H.-W. Dong, W.-Y. Hsiao, L.-C. Yang and Y.-H. Yang, Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI 2018), p. 34–41.
- 5 P. Dhariwal, H. Jun, C. Payne, J. Kim, A. Radford and I. Sutskever, Proceedings of the 37th International Conference on Machine Learning (ICML 2020), p. 2701–2710.
- 6 D. Cope, *Experiments in Musical Intelligence*, A-R Editions.
- 7 J.-P. Briot, G. Hadjeres and F.-D. Pachet, *IEEE Computational Intelligence Magazine*, **13**, 14–24.
- 8 M. Cuthbert and C. Ariza, Proceedings of the 11th International Society for Music Information Retrieval Conference (ISMIR 2010), p. 637–642.
- 9 A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. Gomez, Kaiser and I. Polosukhin, *Advances in Neural Information Processing Systems 30 (NIPS 2017)*, p. 5998–6008.
- 10 S. Hochreiter and J. Schmidhuber, *Neural Computation*, **9**, 1735–1780.
- 11 D. Eck and J. Schmidhuber, Proceedings of the IEEE Workshop on Neural Networks for Signal Processing (NNSP 2002), p. 747–752.
- 12 I. Sutskever, O. Vinyals and Q. Le, *Advances in Neural Information Processing Systems 27 (NIPS 2014)*, p. 3104–3112.
- 13 Abadi, Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 2016), p. 265–283.