





Beam Search (SBBS) to generate a sequence of symbols  $y$  that matches the desired length  $l$  and emotion with arousal  $a$  and valence  $v$ . SBBS receives as input the models  $L$  and  $E_m$ , the current sequence  $x$ , the desired emotion values  $v$  and  $a$ , SBBS’s parameter values  $b$  and  $k$ , which are explained below, and the desired length  $l$  of the piece to be generated.

In the first call to Composer, the sequence  $x$  is initialized with the symbols of the first 4 timesteps of a random human-composed piece with the emotion  $v, a$ , as returned by  $E_s$ . Every time there is a transition from one emotion to another, we reinitialize the sequence  $x$  using the same process. This is used to bias the generative process and to emphasize emotion transitions.

To be used in real-time, Composer is invoked with the most recently captured speech signal  $v$  and returns a composed piece of music. While the most recent piece is being played at the game table, Composer receives another signal  $v$  and composes the next excerpt. One also needs to define the length of the signal  $v$ . In our implementation, similar to Padovani et al. (2017), we use YouTube’s subtitle system as the speech recognition system  $S$ . Therefore, signals  $v$  are long enough to form a subtitle.

### Classifying the Story’s Emotion

In order to have a common model of emotion between stories and music, we use a mapping from Bardo’s four emotion model to the valence-arousal model. Namely, we have Suspenseful mapping to low valence and arousal ( $v = 0, a = 0$ ); Agitated to low valence and high arousal ( $v = 0, a = 1$ ); Calm to high valence and low arousal ( $v = 1, a = 0$ ); and Happy to high valence and arousal ( $v = 1, a = 1$ ).

For example, in the context of the game Dungeons and Dragons, the sentence “Roll initiative” is normally said at the beginning of battles and it can be considered ( $v = 0, a = 1$ ), once a battle is a negative (dangerous) moment with high energy. “Roll initiative” is normally classified as Agitated in Padovani et al.’s dataset. This mapping allows us to use the valence-arousal model with the labeled CotW dataset.

The valence-arousal mapping is based on the model used to annotate the VGMIDI dataset. When human subjects annotated that dataset, they used a continuous valence/arousal model with labels defining a fixed set of discrete basic emotions (e.g. happy or sad) (Ferreira and Whitehead 2019).

Given the limited amount of TRPG stories labeled according to emotion (there are only 5,892 sentences in the CotW dataset), we use a transfer learning approach to classify the sentences (Radford et al. 2018). We fine-tune a high-capacity BERT architecture with the CotW dataset (Devlin et al. 2018). We use BERT because it outperforms several other transformers across different NLP tasks (Devlin et al. 2018). Although in Algorithm 1 we depict the classifier for story emotions as a single  $E_s$  model, in our implementation we treat valence and arousal independently, thus we fine-tune a pre-trained BERT for each dimension.

### Classifying the Music’s Emotion

As was the case with the TRPG stories, given the limited amount of MIDI pieces labeled according to emotion, we also apply a transfer learning approach to classify emotion

in music ( $E_m$ ). However, different than the  $E_s$  model where we fine-tune a BERT architecture, for  $E_m$  we fine-tune a GPT-2 architecture (Radford et al. 2019). We use GPT-2 for  $E_m$  because it is better suited for sequence generation than BERT. Similarly to  $E_s$ , model  $E_m$  also treats valence and arousal independently. Thus, we fine-tune a pre-trained GPT-2 for each of these dimensions.

To the best of our knowledge, in the symbolic music domain, there are no publicly available high-capacity LM pre-trained with large (general) datasets. Typically, models in this domain are trained with relatively small and specific datasets. For example, the MAESTRO dataset (Hawthorne et al. 2019), the Bach Chorales (Hadjeres, Pachet, and Nielsen 2017) and the VGMIDI (Ferreira and Whitehead 2019) dataset. We pre-train a general high-capacity GPT-2 architecture as a language model (Radford et al. 2019) using a new dataset we created called ADL (Augmented Design Lab) Piano MIDI dataset<sup>1</sup>.

The ADL Piano MIDI dataset is based on the Lakh MIDI dataset (Raffel 2016), which, to the best of our knowledge, is the largest MIDI dataset publicly available. The Lakh MIDI dataset contains a collection of 45,129 unique MIDI files that have been matched to entries in the Million Song dataset (Bertin-Mahieux et al. 2011). Among these files, there are many versions of the same piece. We kept only one version of each piece. Given that the datasets for emotion classification in music are limited to piano only, we extracted from the Lakh MIDI dataset only the tracks with instruments from the “piano family”(MIDI program numbers 1-8 in the dataset). This process generated a total of 9,021 unique piano MIDI files. These files are mainly Rock and Classical pieces, so to increase the genre diversity (e.g. Jazz, Blues, and Latin) of the dataset, we included an additional 2,065 files scraped from public sources on the Internet<sup>2</sup>. All files in the final collection were de-duped according to their MD5 checksum. The final dataset has 11,086 pieces.

After pre-training the high-capacity GPT-2 model, we fine-tune two independent models (one for valence and one for arousal) with an extended version of the VGMIDI dataset (Ferreira and Whitehead 2019). We extended the VGMIDI dataset from 95 to 200 labeled pieces using the same annotation method of the original dataset. All the 200 pieces are piano arrangements of video game soundtracks labeled according to the valence-arousal model of emotion.

**Encoding** We encode a MIDI file by parsing all notes from the NOTE\_ON and NOTE\_OFF events in the MIDI. We define a note as a set  $z = (z_p, z_s, z_d, z_v)$ , where  $\{z_p \in \mathbb{Z} | 0 \leq z_p < 128\}$  is the pitch number,  $\{z_s \in \mathbb{Z} | z_s \geq 0\}$  is the note starting time in timesteps,  $\{z_d \in \mathbb{Z} | 0 \leq z_d \leq 56\}$  is note duration in timesteps and  $\{z_v \in \mathbb{Z} | 0 \leq z_v < 128\}$  is the note velocity. Given a MIDI NOTE\_ON event, we parse a note  $z$  by retrieving the starting time  $z_s$  (in seconds), the pitch number  $z_p$  and the velocity  $z_v$  from that event. To calculate the note duration  $z_d$ , we find the correspondent NOTE\_OFF event of the given NOTE\_ON and retrieve the NOTE\_OFF end time  $z_e$  (in seconds). We discretize  $z_s$  and  $z_e$  to compute the

<sup>1</sup><https://github.com/lucasnfe/adl-piano-midi>

<sup>2</sup><https://bushgrafts.com/midi/> and <http://midkar.com/jazz/>

---

**Algorithm 2** Stochastic Bi-Objective Beam Search

---

**Require:** Music emotion classifier  $E_m$ , LM  $L$ , previously composed symbols  $x$ , valence and arousal values  $v$  and  $a$ , number  $k$  of symbols to consider, beam size  $b$ , length  $l$  in seconds of the generated piece.

**Ensure:** Sequence of symbols of  $l$  seconds.

```
1:  $B \leftarrow [x], j \leftarrow 0$ 
2: while  $l(y[t:t+j]) < l, \forall y \in B$  do
3:    $C \leftarrow \{\}$ 
4:   for all  $m \in B$  do
5:      $C_m \leftarrow \{m \cup s | s \in V\}$ 
6:      $C_m \leftarrow k$  elements  $y$  from  $C_m$  with largest  $p_L(y)$ 
7:      $C \leftarrow C \cup C_i$ 
8:    $B \leftarrow b$  sequences  $y$  sampled from  $C$  proportionally
   to  $p_L(y)(1 - |v - E_{m,v}(y)|)(1 - |a - E_{m,a}(y)|)$ 
9:    $j \leftarrow j + 1$ 
10: return  $m \in B$  such that  $p_L(m) = \max_{y \in B} p_L(y)$  and
 $l(y[t:t+j]) \geq l$ 
```

---

note duration  $z_d = t \cdot z_e - t \cdot z_s$  in timesteps, where  $t$  is a parameter defining the sampling frequency of the timesteps.

We derive a sequence  $x = \{z_v^1, z_d^1, z_p^1, \dots, z_v^n, z_d^n, z_p^n\}$  of tokens for a given MIDI file by (a) parsing all notes  $z^i$  from the file, (b) sorting them by starting time  $z_s^j$  and (c) concatenating their velocity  $z_v^j$ , duration  $z_d^j$  and pitch  $z_p^j$ . We add two special tokens TS and END in the sequence  $x$ , to mark the end of a timestep and the end of a piece, respectively. This encoding yields a vocabulary  $V$  of size  $|V| = 314$ .

### Stochastic Bi-Objective Beam Search

Next, we describe how one can use a LM and a music emotion classifier to bias the process of music generation to match a particular emotion (line 3 of Algorithm 1). For that we introduce Stochastic Bi-Objective Beam Search (SBBS), a search algorithm guided by the LM  $L$  and the music emotion classifiers, denoted as  $E_{m,v}$  and  $E_{m,a}$ , for valence and arousal. The goal of SBBS is to allow for the generation of pieces that sound “good” (i.e., have high probability value according to the trained LM), but that also match the current emotion of the story being told by the players.

We call SBBS “stochastic” because it samples from a distribution instead of greedily selecting the best sequences of symbols, as a regular beam search does. The stochasticity of SBBS allows it to generate a large variety of musical pieces for the same values of  $v$  and  $a$ . We also call it “bi-objective” because it optimizes for realism and emotion.

The pseudocode of SBBS is shown in Algorithm 2. In the pseudocode we use letters  $x, y$  and  $m$  to denote sequences of musical symbols. Function  $p_L(y) = \prod_{y_t \in y} P(y_t | y_0, \dots, y_{t-1})$  is the probability of sequence  $y$  according to the LM  $L$ ; a high value of  $p_L(y)$  means that  $y$  is recognized as a piece of “good quality” by  $L$ . We denote as  $l(y)$  the duration in seconds of piece  $y$ . Finally, we write  $x[i:j]$  for  $j \geq i$  to denote the subsequence of  $x$  starting at index  $i$  and finishing at index  $j$ .

SBBS initializes the beam structure  $B$  with the sequence

$x$  passed as input (line 1). SBBS also initializes variable  $j$  for counting the number of symbols added by the search. SBBS keeps in memory at most  $b$  sequences and, while all sequences are shorter than the desired duration  $l$  (line 2), it adds a symbol to each sequence (lines 3–9). SBBS then generates all sequences by adding one symbol from vocabulary  $V$  to each sequence  $m$  from  $B$  (line 5); these extended sequences, known as the children of  $m$ , are stored in  $C_m$ .

The operations performed in lines 6 and 8 attempt to ensure the generation of good pieces that convey the desired emotion. In line 6, SBBS selects the  $k$  sequences with largest  $p_L$ -value among the children of  $m$ . This is because some of the children with low  $p_L$ -value could be attractive from the perspective of the desired emotion and, although the resulting piece could convey the desired emotion, the piece would be of low quality according to the LM. The best  $k$  children of each sequence in the beam are added to set  $C$  (line 7). Then, in line 8, SBBS samples the sequences that will form the beam of the next iteration. Sampling occurs proportionally to the values of  $p_L(y)(1 - |v - E_{m,v}(y)|)(1 - |a - E_{m,a}(y)|)$ , for sequences  $y$  in  $C$ . A sequence  $y$  has higher chance of being selected if  $L$  attributes a high probability value to  $y$  and if the music emotion model classifies the values of valence and arousal of  $y$  to be similar to the desired emotion. When at least one of the sequences is longer than the desired duration of the piece, SBBS returns the sequence with largest  $p_L$ -value that satisfies the duration constraint (line 10).

## Empirical Evaluation

Our empirical evaluation is divided into two parts. First, we evaluate the accuracy of the models used for story and music emotion classification. We are interested in comparing the fine-tuned BERT model for story emotion classification with the simpler Naïve Bayes approach of Padovani, Ferreira, and Lelis (2017). We are also interested in comparing the fine-tuned GPT-2 model for music emotion classification with the simpler LSTM of Ferreira and Whitehead (2019). In the second part of our experiments we evaluate with a user study whether human subjects can recognize different emotions in pieces generated by Composer for the CotW campaign.

### Emotion Classifiers

**Story Emotion** The story emotion classifier we use with Composer is a pair of BERT models, one for valence and one for arousal. For both models, we use the pre-trained BERT<sub>BASE</sub> of Devlin et al. (2018), which has 12 layers, 768 units per layer, and 12 attention heads. BERT<sub>BASE</sub> was pre-trained using both the BooksCorpus (800M words) (Zhu et al. 2015) and the English Wikipedia (2,500M words).

We independently fine-tune these two BERT models as valence and arousal classifiers using the CotW dataset (Padovani, Ferreira, and Lelis 2017). Fine-tuning consists of adding a classification head on top the pre-trained model and training all the parameters (including the pre-trained ones) of the resulting model end-to-end. All these parameters were fine-tuned with an Adam optimizer (Kingma and Ba 2014) with learning rate of 3e-5 for 10 epochs. We used mini-batches of size 32 and dropout of 0.5.

Alg.	Episodes									Avg.
	1	2	3	4	5	6	7	8	9	
NB	73	88	91	85	94	81	41	74	94	80
BERT	<b>89</b>	<b>92</b>	<b>96</b>	<b>88</b>	<b>97</b>	<b>81</b>	<b>66</b>	<b>83</b>	<b>96</b>	<b>87</b>

Table 1: Valence accuracy in % of Naïve Bayes (NB) and BERT for story emotion classification.

The CotW dataset is divided into 9 episodes, thus we evaluate the accuracy of each BERT classifier using a leave-one-out strategy. For each episode  $e$ , we leave  $e$  out for testing and train in the remaining episodes. For example, when testing on episode 1, we use episodes 2-8 for training. Every sentence is encoded using a WordPiece embedding (Wu et al. 2016) with a 30,000 token vocabulary.

We compare the fine-tuned BERT classifiers with a Naïve Bayes (NB) approach (baseline), chosen because it is the method underlying the original Bardo system. NB encodes sequences using a traditional bag-of-words with tf-idf approach. Table 1 shows the accuracy of the valence classification of both these methods per episode. The best accuracy for a given episode is highlighted in bold. The BERT classifier outperforms NB in all the episodes, having an average accuracy 7% higher. For valence classification, the hardest episode for both the models is episode 7, where BERT had the best performance improvement when compared to NB. The story told in episode 7 of CotW is different from all other episodes. While the other episodes are full of battles and ability checks, episode 7 is mostly the players talking with non-player characters. Therefore, what is learned in the other episodes does not generalize well to episode 7. The improvement in accuracy of the BERT model in that episode is likely due to the model’s pre-training. Episodes 5 and 9 were equally easy for both methods because they are similar to one another. The system trained in one of these two episodes generalizes well to the other.

Table 2 shows the accuracy of arousal classification of both NB and BERT. The best accuracy for a given episode is highlighted in bold. Again BERT outperforms NB in all the episodes, having an average accuracy 5% higher. In contrast with the valence results, here there is no episode in which the BERT model substantially outperforms NB.

Alg.	Episodes									Avg.
	1	2	3	4	5	6	7	8	9	
NB	82	88	75	79	82	76	98	86	84	83
BERT	<b>86</b>	<b>90</b>	<b>77</b>	<b>86</b>	<b>89</b>	<b>88</b>	<b>99</b>	<b>90</b>	<b>88</b>	<b>88</b>

Table 2: Arousal accuracy in % of Naïve Bayes (NB) and BERT for story emotion classification.

**Music Emotion** The music emotion classifier is a pair of GPT-2 models, one for valence and one for arousal. We first pre-trained a GPT-2 LM with our ADL Piano MIDI dataset. We augmented each piece  $p$  of this dataset by (a) transpos-

Algorithm	Valence	Arousal
Baseline LSTM	69	67
Fine-tuned LSTM	74	79
Baseline GPT-2	70	76
Fine-tuned GPT-2	<b>80</b>	<b>82</b>

Table 3: Accuracy in % of both the GPT-2 and LSTM models for music emotion classification.

ing  $p$  to every key, (b) increasing and decreasing  $p$ ’s tempo by 10% and (c) increasing and decreasing the velocity of all notes in  $p$  by 10% (Oore et al. 2017). Thus, each piece generated  $12 \cdot 3 \cdot 3 = 108$  different examples.

The pre-trained GPT-2 LM has 4 layers (transformer blocks), context size of 1024 tokens, 512 embedding units, 1024 hidden units, and 8 attention heads. We then fine-tuned the GPT-2 LM independently using the VGMIDI dataset, for valence and for arousal. Similarly to BERT, fine-tuning a GPT-2 architecture consists of adding an extra classification head on top of the pre-trained model and training all parameters end-to-end. Similar to the story emotion classifiers, we fine-tuned the GPT-2 classifiers for 10 epochs using an Adam optimizer with learning rate  $3e-5$ . Different from the story emotion classifiers, we used mini-batches of size 16 (due to GPU memory constrains) and dropout of 0.25. The VGMIDI dataset is defined with a train and test splits of 160 and 40 pieces, respectively. We augmented the dataset by slicing each piece into 2, 4, 8 and 16 parts of equal length and emotion. Thus, each part of each slicing generated one extra example. This augmentation is intended to help the classifier generalize for pieces with different lengths.

We compare the fine-tuned GPT-2 classifiers with LSTM models that were also pre-trained with the ADL Piano Midi dataset and fine-tuned with the VGMIDI dataset. We chose LSTMs because they are the state-of-the-art model in the VGMIDI dataset (Ferreira and Whitehead 2019). The LSTMs have same size as the GPT-2 models (4 hidden layers, 512 embedding units, 1024 hidden units) and were pre-trained and fine-tuned with the same hyper-parameters. Table 3 shows the accuracy of both models for valence and arousal. We also report the performance of these models without pre-training (i.e., trained only on the VGMIDI dataset). We call these the baseline versions of the models.

Results show that using transfer learning can substantially boost the performance of both models. The fine-tuned GPT-2 is 10% more accurate in terms of valence and 8% in terms of arousal. The fine-tuned LSTM is 5% more accurate in terms of valence and 12% in terms of arousal. Finally, the fine-tuned GPT-2 outperformed the fine-tuned LSTM by 6% and 3% in terms of valence and arousal, respectively.

## User Study

In our study we measure Composer’s performance at generating music that matches the emotions of a story. We use Composer to generate a piece for a snippet composed of 8 contiguous sentences of each of the first 5 episodes of the CotW dataset. Each snippet has one emotion transition that



Method	Episodes																Average						
	e1-p1		e1-p2		e2-p1		e2-p2		e3-p1		e3-p2		e4-p1		e4-p2		e5-p1		e5-p2		v	a	va
	v	a	v	a	v	a	v	a	v	a	v	a	v	a	v	a	v	a	v	a			
<b>Baseline</b>	56	<b>65</b>	39	56	39	62	39	<b>79</b>	<b>48</b>	<b>60</b>	<b>67</b>	<b>53</b>	<b>58</b>	70	<b>63</b>	<b>75</b>	25	36	<b>72</b>	58	<b>51</b>	<b>32</b>	<b>34</b>
<b>Composer</b>	<b>62</b>	60	<b>44</b>	<b>65</b>	<b>82</b>	<b>68</b>	<b>53</b>	68	24	55	46	43	25	<b>87</b>	37	55	<b>81</b>	<b>86</b>	51	<b>67</b>	<b>51</b>	30	<b>34</b>

Table 4: The percentage of participants that correctly identified the valence and arousal (v and a, respectively) intended by the methods for the pieces parts (p1 and p2).

happens in between sentences. The sentences are 5.18 seconds long on average. To test Composer’s ability to generate music pieces with emotion changes, we asked human subjects to listen to the 5 generated pieces and evaluate the transitions of emotion in each generated piece.<sup>3</sup>

The user study was performed via Amazon Mechanical Turk and had an expected completion time of approximately 10 minutes. A reward of USD \$1 was given to each participant who completed the study. In the first section of the study, the participants were presented an illustrated description of the valence-arousal model of emotion and listened to 4 examples of pieces from the VGMIDI dataset labeled with the valence-arousal model. Each piece had a different emotion: low valence and arousal, low valence and high arousal, high valence and low arousal, high valence and arousal.

In the second section of the study, participants were asked to listen to the 5 generated pieces (one per episode). After listening to each piece, participants had to answer 2 questions: (a) “What emotion do you perceive in the 1st part of the piece?” and (b) “What emotion do you perceive in the 2nd part of the piece?” To answer these two questions, participants selected one of the four emotions: low valence and arousal, low valence and high arousal, high valence and low arousal, high valence and arousal. Subjects were allowed to play the pieces as many times as they wanted before answering the questions. The final section of the study was a demographics questionnaire including ethnicity, first language, age, gender, and experience as a musician. To answer the experience as a musician, we used a 1-to-5 Likert scale where 1 means “I’ve never studied music theory or practice” and 5 means “I have an undergraduate degree in music”.

We compare Composer with a baseline method that selects a random piece from the VGMIDI dataset whenever there is a transition of emotion. The selected piece has the same emotion of the sentence (as given by the story emotion classifier). To compare these two methods, we used a between-subject strategy where Group *A* of 58 participants evaluated the 5 pieces generated by Composer and another Group *B* of 58 participants evaluated the 5 pieces from the baseline. We used this strategy to avoid possible learning effects where subjects could learn emotion transitions from one method and apply the same evaluation directly to the other method. The average age of groups *A* and *B* are 34.96 and 36.98 years, respectively. In Group *A*, 69.5% of participants are male and 30.5% are female. In Group *B*, 67.2%

are male and 32.8% are female. The average musicianship of the groups *A* and *B* are 2.77 and 2.46, respectively.

Table 4 shows the results of the user study. We consider both parts (p1 and p2 in the table) of each episode as an independent piece. The table presents the percentage of participants that correctly identified the pieces’ valence and arousal (“v” and “a” in the table, respectively), as intended by the methods. For example, 87% of the participants correctly identified the arousal value that Composer intended the generated piece for part p1 of episode 4 (e4-p1) to have. We refer to the percentage of participants that are able to identify the approach’s intended emotion as the approach’s accuracy. We also present the approaches’ average accuracy across all pieces (“Average” in the table) in terms of valence, arousal, and jointly for valence and arousal (“va” in the table). The “va”-value of 34 for Composer means that 34% of the participants correctly identified the system’s intended values for valence and arousal across all pieces generated.

Composer outperformed the Baseline in e1-p2, e2-p1, and e5-p1. Baseline outperformed Composer e3-p1, e3-p2 and e4-p2. In the other four parts, one method performed better for valence whereas the other method performance better for arousal. Overall, the average results show that both systems performed very similarly. Both of them had an average accuracy on the combined dimensions equal to 34%. The difference between these two methods and a system that selects pieces at random (expected accuracy of 25%) is significant according to a Binomial test ( $p = 0.02$ ). These results show that the participants were able to identify the emotions in the generated pieces as accurately as they were able to identify the emotions in human-composed pieces. This is an important result towards the development of a fully automated system for music composition for story-based tabletop games.

## Conclusions

This paper presented Bardo Composer, a system that automatically composes music for tabletop role-playing games. The system processes sequences from speech and generates pieces one sentence after the other. The emotion of the sentence is classified using a fine-tuned BERT. This emotion is given as input to a Stochastic Bi-Objective Beam Search algorithm that tries to generate a piece that matches the emotion. We evaluated Composer with a user study and results showed that human subjects correctly identified the emotion of the generated music pieces as accurately as they were able to identify the emotion of pieces composed by humans.

<sup>3</sup>Generated pieces can be downloaded from the following link: <https://github.com/lucasnf/bardo-composer>

## References

- Bergström, K., and Björk, S. 2014. The case for computer-augmented games. *Transactions of the Digital Games Research Association* 1(3).
- Bertin-Mahieux, T.; Ellis, D. P.; Whitman, B.; and Lamere, P. 2011. The million song dataset.
- Briot, J.-P.; Hadjeres, G.; and Pachet, F. 2017. Deep learning techniques for music generation—a survey. *arXiv preprint arXiv:1709.01620*.
- Dathathri, S.; Madotto, A.; Lan, J.; Hung, J.; Frank, E.; Molino, P.; Yosinski, J.; and Liu, R. 2019. Plug and play language models: a simple approach to controlled text generation. *arXiv preprint arXiv:1912.02164*.
- Davis, H., and Mohammad, S. M. 2014. Generating music from literature. *Proceedings of the 3rd Workshop on Computational Linguistics for Literature (CLfL)* 1–10.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ferreira, L. N., and Whitehead, J. 2019. Learning to generate music with sentiment. In *Proceedings of the International Society for Music Information Retrieval Conference, ISMIR'19*.
- Hadjeres, G.; Pachet, F.; and Nielsen, F. 2017. Deepbach: a steerable model for bach chorales generation. In *International Conference on Machine Learning*, 1362–1371.
- Hawthorne, C.; Stasyuk, A.; Roberts, A.; Simon, I.; Huang, C.-Z. A.; Dieleman, S.; Elsen, E.; Engel, J.; and Eck, D. 2019. Enabling factorized piano music modeling and generation with the MAESTRO dataset. In *International Conference on Learning Representations*.
- Keskar, N. S.; McCann, B.; Varshney, L. R.; Xiong, C.; and Socher, R. 2019. Ctrl: A conditional transformer language model for controllable generation. *arXiv preprint arXiv:1909.05858*.
- Kingma, D., and Ba, J. 2014. Adam: A method for stochastic optimization. *International Conference on Learning Representations*.
- Kool, W.; Hoof, H. V.; and Welling, M. 2019. Stochastic beams and where to find them: The Gumbel-top-k trick for sampling sequences without replacement. In Chaudhuri, K., and Salakhutdinov, R., eds., *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, 3499–3508. Long Beach, California, USA: PMLR.
- Monteith, K.; Martinez, T. R.; and Ventura, D. 2010. Automatic generation of music for inducing emotive response. In *International Conference on Computational Creativity*, 140–149.
- Oore, S.; Simon, I.; Dieleman, S.; and Eck, D. 2017. Learning to create piano performances. In *NIPS 2017 Workshop on Machine Learning for Creativity and Design*.
- Padovani, R.; Ferreira, L. N.; and Lelis, L. H. S. 2017. Bardo: Emotion-based music recommendation for tabletop role-playing games. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*.
- Padovani, R.; Ferreira, L. N.; and Lelis, L. H. S. 2019. Be inaccurate but don't be indecisive: How error distribution can affect user experience. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2604–2611.
- Poole, D. L., and Mackworth, A. K. 2010. *Artificial Intelligence: foundations of computational agents*. Cambridge University Press.
- Radford, A.; Narasimhan, K.; Salimans, T.; and Sutskever, I. 2018. Improving language understanding by generative pre-training. In *Arxiv*.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1(8):9.
- Raffel, C. 2016. *Learning-based methods for comparing sequences, with applications to audio-to-midi alignment and matching*. Ph.D. Dissertation, Columbia University.
- Russell, J. A. 1980. A circumplex model of affect. *Journal of personality and social psychology* 39(6):1161.
- Vijayakumar, A. K.; Cogswell, M.; Selvaraju, R. R.; Sun, Q.; Lee, S.; Crandall, D.; and Batra, D. 2018. Diverse beam search for improved description of complex scenes. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Williams, D.; Kirke, A.; Eaton, J.; Miranda, E.; Daly, I.; Hallowell, J.; Roesch, E.; Hwang, F.; and Nasuto, S. J. 2015a. Dynamic game soundtrack generation in response to a continuously varying emotional trajectory. In *Audio Engineering Society Conference: 56th International Conference: Audio for Games*. Audio Engineering Society.
- Williams, D.; Kirke, A.; Miranda, E. R.; Roesch, E.; Daly, I.; and Nasuto, S. 2015b. Investigating affect in algorithmic composition systems. *Psychology of Music* 43(6):831–854.
- Wu, Y.; Schuster, M.; Chen, Z.; Le, Q. V.; Norouzi, M.; Macherey, W.; Krikun, M.; Cao, Y.; Gao, Q.; Macherey, K.; et al. 2016. Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Zhu, Y.; Kiros, R.; Zemel, R.; Salakhutdinov, R.; Urtasun, R.; Torralba, A.; and Fidler, S. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*, 19–27.