

# MMM : Exploring Conditional Multi-Track Music Generation with the Transformer

Jeff Ens and Philippe Pasquier \*

Simon Fraser University  
jeffe@sfu.ca

**Abstract.** We propose MMM, a generative system based on the Transformer architecture that is capable of generating multi-track music. In contrast to previous work, which represents musical material as a single time-ordered sequence, where the musical events corresponding to different tracks are interleaved, we create a time-ordered sequence of musical events for each track and concatenate several tracks into a single sequence. This takes advantage of the attention-mechanism, which can adeptly handle long-term dependencies. We explore how various representations can offer the user a high degree of control at generation time, providing an interactive demo that accommodates track-level and bar-level inpainting, and offers control over track instrumentation and note density.

**Keywords:** Symbolic Music Generation, Multi-Track

## 1 Introduction

Research involving generative music systems has focused on modelling musical material as an end-goal, rather than on the affordances of such systems in practical scenarios (Sturm et al., 2019). As a result, there has been a focus on developing novel architectures and demonstrating that music generated with these architectures is of comparable quality to human-composed music, often via a listening test. Although this is a necessary first step, as systems must be capable of generating compelling material before they can be useful in a practical context, given the impressive capabilities of the Transformer-based models in the music domain (Donahue, Mao, Li, Cottrell, & McAuley, 2019; Huang et al., 2019), we shift our focus to increasing the affordances of a Transformer-based system. To achieve this goal, we develop a novel representation for multi-track musical material that accommodates a variety of methods for generation.

To avoid any confusion, we will first define what constitutes multi-track music. We consider a *track* to be a collection of notes played on a single instrument. Although alternate terminology has been employed to describe tracks, which

---

\* We acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC), and the Helmut & Hugo Eppich Family Graduate Scholarship

may be referred to as voices or instruments in various contexts, we believe that the term track is clearest, as there is a clear analog to the tracks present in a digital audio workstation. We avoid using the term voice, as it commonly connotes a monophonic musical line, while we wish to refer to musical material that may contain multiple notes sounding simultaneously. Furthermore, within a piece there may be multiple tracks featuring the same instrument, each playing a different musical part, which would make usage of the term instrument problematic. Consequently, multi-track music refers to material containing two or more tracks, where each track is played by a single instrument and may optionally contain multiple notes that sound simultaneously. It is also important to note the difference between *polyphonic tracks*, which contain simultaneously sounding notes, and *monophonic tracks*, which contain a single sequence of non-overlapping notes.

Given our interest in enhancing the usability of a system at generation time, it is worth reviewing different methods for generation, which we group into four categories: unconditioned, continuation, inpainting, and attribute-control. Unconditioned generation is analogous to generating music from scratch. Besides changing the data that the model is trained on, the user has limited control over the output of the model. Continuation involves conditioning the model with musical material that precedes (temporally) the music that is to be generated. Since both unconditioned generation and continuation come for free with any autoregressive model trained on a temporally ordered sequence of musical events, most systems are capable of generating musical material in this manner. Inpainting conditions generation on a subset of musical material, asking the model to fill in the blanks, so to speak. Note that inpainting can occur at different levels (i.e. note-level, bar-level, track-level). CoCoNet (Huang, Cooijmans, Roberts, Courville, & Eck, 2017) allows for inpainting of Bach chorales on the note-level, while InpaintNet (Pati, Lerch, & Hadjeres, 2019) allows for inpainting of 2-8 bars of monophonic musical material. Attribute-control involves conditioning generation on high-level attributes such as style, tempo or density. For example, music generated by MuseNet (Payne, 2019) can be conditioned on a set of instruments and a musical style. In some circumstances, generation methods can be chained, resulting in an iterative generation process. For example, a musical segment exhibiting a particular style could be generated via attribute control, and the user could then select various sections they are unsatisfied with for inpainting.

Our primary contribution is a novel representation for musical material, which when coupled with state-of-the-art transformer architectures, results in a powerful and expressive generative system. In contrast to previous work, which represents musical material as a single time-ordered sequence, where the musical events corresponding to different tracks are interleaved, we create a time-ordered sequence of musical events for each track and concatenate several tracks into a single sequence. Although the difference is subtle, this enables track-level inpainting, and attribute control over each track. We also explore variations on this representation which allow for within-track bar-level inpainting. Unto our

knowledge, both inpainting and attribute control have not been integrated into a single model.

## 2 Related Work

There are two main ways in which musical material is represented: as a matrix (i.e. a pianoroll), or as a sequence of tokens. A piano roll is a boolean matrix  $x \in \{0, 1\}^{T \times P}$ , where  $T$  is the number of time-steps and  $P$  is the number of pitches. Typically  $P = 128$ , allowing the piano roll to represent all possible MIDI pitches, however, it is not uncommon to reduce the range of pitches represented (Dong, Hsiao, Yang, & Yang, 2018). Multi-track musical material can be represented using a boolean tensor  $x \in \{0, 1\}^{M \times T \times P}$ , where  $M$  is the number of tracks. However, using this type of representation is inherently inefficient, as the number number of inputs increases by  $T \times P$  for each track that is added, and accommodating small note lengths (ex.  $32^{nd}$  triplets) substantially increases  $T$ . Despite these drawbacks, this representation has been used in practice (Boulanger-Lewandowski, Bengio, & Vincent, 2012; Dong et al., 2018; Huang et al., 2017).

The alternative approach, is to represent musical material as a sequence of tokens, where each token corresponds to a specific musical event. For example, the PerformanceRNN (Oore, Simon, Dieleman, Eck, & Simonyan, 2018) and Music Transformer (Huang et al., 2019) use a token based representation comprised of 128 distinct NOTE\_ON tokens, which are used to indicate the onset of a particular pitch; 128 NOTE\_OFF tokens, which denote the end of a particular pitch; 100 TIME\_SHIFT tokens, which correspond to different time-shifts ranging from 10ms to 1 second. Although this type of representation can accommodate polyphony, it does not distinguish between different tracks or instruments.

Our work is most similar to LahkNES (Donahue et al., 2019), MusicVAE (Roberts, Engel, Raffel, Hawthorne, & Eck, 2018) and MuseNet (Payne, 2019), which all employ token-based representations to model multi-track music. LahkNES models Nintendo Entertainment System (NES) data, which is comprised of 3 monophonic tracks and a drum track, using a transformer architecture. MusicVAE is trained on the bass, melody and drum trios extracted from the Lahk MIDI Dataset (LMD) (Raffel, 2016), which allows generation to be conditioned on a latent vector. MuseNet is trained on a superset of the LMD, and accommodates 10 different track types ranging from piano to guitar. Note that MuseNet supports both polyphonic and monophonic tracks.

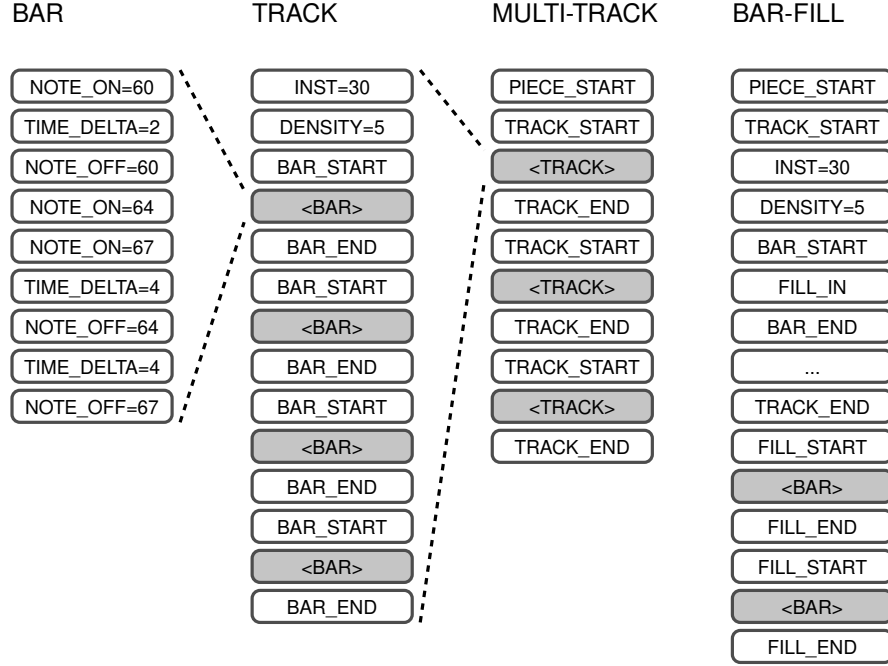
However, in contrast to these methods, where the musical events from several tracks are interleaved into a single time-ordered sequence, we concatenate single-instrument tracks, which allows for greater flexibility in several areas. First of all, we can decouple track information from NOTE\_ON and NOTE\_OFF tokens, allowing the use of the same NOTE\_ON and NOTE\_OFF tokens in each track. This differs from LahkNES, MusicVAE, and MuseNet, which use separate NOTE\_ON and NOTE\_OFF tokens for each track, placing inherent limitations on the number of tracks that can be represented. Even MuseNet, which is the largest of these networks, can

only accommodate 10 different tracks. Secondly, using our representation, we are able to accommodate a wide variety of instruments, including all 128 general midi instruments, and a large number of tracks, without employing a prohibitively large token vocabulary. In contrast to LahkNES and MusicVAE which are designed for a fixed-schema of tracks, our system can handle an arbitrary set of tracks. MuseNet is similar in this regard, however, it supports a smaller number of instruments. Although MuseNet permits attribute control over the instrument, allowing the user to specify the set of instruments that will be featured in the generated excerpt, this information is only treated as a strong recommendation to the model, and does not guarantee which instruments will actually be used. Our system allows for specific attribute control over the instrument for each track, with the guarantee that a particular instrument will be used. Third, we offer the user control over the note-density of each track, which is not accommodated with LahkNES, MusicVAE or MuseNet. Finally, we allow for track-level and bar-level inpainting, which is not possible using LahkNES, MusicVAE and MuseNet. Collectively, these improvements afford the end-user a high-degree of control over the generated material, which has previously been proposed as a critical area of research (Briot, Hadjeres, & Pachet, 2019).

### 3 Motivation

Although systems which generate high-quality music have been proposed in recent years (Huang et al., 2019; Payne, 2019; Liang, 2016; Sturm & Ben-Tal, 2017), their usage in practical contexts is limited for two different reasons. First of all, most models place restrictions on the nature of the input. In most cases, there are limitations placed on the number and type of tracks (Roberts et al., 2018; Payne, 2019). Secondly, the user is not afforded fine-grained control over the generation process, which is critical for a system to be useful in the context of computational assisted composition. Even MusicVAE (Roberts et al., 2018), which incorporates a latent model of musical space, allowing for interpolation between examples, does not afford fine-grained control of the individual tracks. For example, it is not possible to freeze the melody and generate a new drum-part and bassline. Although one-shot generation of musical material is impressive from a technical standpoint, it is not that useful in a practical context, as the user may wish to create subtle variations on a fixed piece of music.

In contrast to time-ordered sequences, where most of the important dependencies, such as the most recently played notes, are in the recent history, non-time-ordered sequences frequently feature important dependencies in the distant history. For example, in our representation, simultaneously sounding notes in different tracks are spread far apart. The use of non-time-ordered representations is directly motivated the nature of the transformer attention mechanism (Vaswani et al., 2017). In contrast to Recurrent Neural Networks (RNN), which sharply distinguish nearby context (the most recent 50 tokens) from the distant history (Khandelwal, He, Qi, & Jurafsky, 2018), attention-based architectures allow for distant tokens to be directly attended to if they are relevant to the current pre-



**Fig. 1.** The MultiTrack and BarFill representations are shown. The **<bar>** tokens correspond to complete bars, and the **<track>** tokens correspond to complete tracks.

diction. Consequently, we do not pay a significant penalty for training models on non-time-ordered sequences, where important dependencies are predominantly in the distant history, provided the necessary tokens are within the attention window. This directly motivates the usage of non-time-ordered sequences, as they facilitate rich conditional generation.

## 4 Proposed Representation

To provide a comprehensive overview of the proposed representation, we first describe how a single bar of musical material is represented. Based on representations explored in previous studies (Oore et al., 2018; Huang et al., 2019), we represent musical material using 128 **NOTE\_ON** tokens, 128 **NOTE\_OFF** tokens, and 48 **TIME\_SHIFT** tokens. Since musical events are quantized using 12 subdivisions per beat, 48 **TIME\_SHIFT** tokens allow for the representation of any rhythmic unit from sixteenth note triplets to a full 4-beat bar of silence. Each bar begins with a **BAR\_START** token, and ends with a **BAR\_END** token. Tracks are simply a sequence of bars delimited by **TRACK\_START** and **TRACK\_END** tokens. At the start of each track, immediately following the **TRACK\_START** token, an **INSTRUMENT** token is used to specify the MIDI program which is to be used to play the notes

on this particular track. Since there are 128 possible MIDI programs, we have 128 distinct `INSTRUMENT` tokens. A `DENSITY_LEVEL` token follows the `INSTRUMENT` token, and indicates the note density of the current track. A piece is simply a sequence of tracks, however, all tracks sound simultaneously rather than being played one after the other. A piece begins with the `PIECE.START` token. This process of nesting bars within a track and tracks within a piece is illustrated in Figure 1. Notably, we do not use a `PIECE.END` token, as we can simply sample until we reach the  $n^{th}$  `TRACK.END` token if we wish to generate  $n$  tracks. We refer to this representation as the MultiTrack representation.

Using the MultiTrack representation, the model learns to condition the generation of each track on the tracks which precede it. At generation time, this allows for a subset of the musical material to be fixed while generating additional tracks. However, while the MultiTrack representation offers control at the track level, it does not allow for control at the bar level, except in cases where the model is asked to complete the remaining bars of a track. Without some changes, it is not possible to generate the second bar in a track conditioned on the first, third, and fourth bars. In order to accommodate this scenario, we must guarantee that the bars on which we want to condition precede the bars we wish to predict, in the sequence of tokens that is passed to the model. To do this, we remove all the bars which are to be predicted from the piece, and replace each bar with a `FILL.PLACEHOLDER` token. Then, at the end of the piece (i.e. immediately after the last `TRACK.END` token), we insert each bar, delimiting each bar with `FILL.START` and `FILL.END` tokens instead of `BAR.START` and `BAR.END` tokens. Note that these bars must appear in the same order as they appeared in the original piece representation. We refer to this representation as the BarFill representation. Note that the MultiTrack representation is simply a special case of the BarFill representation, where no bars are selected for inpainting.


## 5 Training

We use the Lahk MIDI Dataset (LMD) (Raffel, 2016), which is comprised of 176,581 MIDI files. In order to explain how we derive token sequences from MIDI files, it is necessary to provide an overview of the MIDI protocol. There are three formats for MIDI files. Type 0 MIDI files are comprised of a header-chunk and a single track-chunk. Both Type 1 and 2 MIDI files contain a header-chunk and multiple track-chunks, however, the tracks in a Type 1 MIDI file are played simultaneously, while tracks in a Type 2 MIDI file are played sequentially. Since only 0.03% of the LMD are Type 2 MIDI files, and the library we use for midi parsing does not support this encoding, we simply ignore them. Within a track-chunk, musical material is represented as a sequence of MIDI messages each which specify a channel and the time delta since the last message. In addition to note-on and note-off messages, which specify the onset and end of notes, patch-change messages specify changes in timbre by selecting one of 128 different instruments. To formally define a track, consider a Type 1 MIDI file  $F = \{t_1, \dots, t_k\}$  comprised of  $k$  tracks, where each track  $t_i = \{m_1^i, \dots, m_{n_i}^i\}$  is an ordered set of  $n_i$  MIDI

messages. Note that a Type 0 MIDI file is simply a special case where  $k = 1$ . Let  $\text{chan}(x)$  (resp.  $\text{inst}(x)$ ) be a function that returns the channel (resp. instrument) on which the message  $x$  is played. Then, we can define a track as the set of midi messages  $t_{i,c,j} = \{m_k^j : \text{inst}(m_k^j) = i, \text{chan}(m_k^j) = c\}$  that are found on the  $j^{\text{th}}$  track, and played on the  $c^{\text{th}}$  channel using the  $i^{\text{th}}$  MIDI instrument. For example, given a MIDI file  $F = \{t_1, t_2\}$ , where  $t_1 = \{m_1^1\}$ ,  $t_2 = \{m_2^1, m_2^2\}$ ,  $\text{chan}(m_1^1) = 0$ ,  $\text{inst}(m_1^1) = 0$ ,  $\text{chan}(m_2^1) = 3$ ,  $\text{inst}(m_2^1) = 0$ ,  $\text{chan}(m_2^2) = 3$ , and  $\text{inst}(m_2^2) = 34$ , there would be three tracks  $(t_{0,0,1}, t_{0,3,0}, t_{0,3,34})$ .

For each of the 128 general MIDI instruments, we calculate the number of note onsets for each bar in the dataset, and use the quantiles of the resulting distributions to define distinct note-density bins for each MIDI instrument. Note that using the same note-density bins for all instrument types would be problematic as note-density varies significantly between instruments. We use 10 different note-density bins, where the  $i^{\text{th}}$  bin is bounded by the  $10i$  (lower) and  $10(i + 1)$  (upper) quantiles. We train a GPT2 (Radford et al., 2019) model using the HuggingFace Transformers library (Wolf et al., 2019) with 8 attention heads, 6 layers, an embedding size of 512, and an attention window of 2048. Each time we select a  $n$ -bar segment, we randomly order the tracks so that the model learns each possible conditional between different types of tracks. We train two separate models: MMMBar, which is trained using the BarFill representation; MMMTrack, which is trained using the MultiTrack representation.

## 6 Using MMM



**Fig. 2.** The user interface for the interactive demo. The bar shown in blue has been selected for bar-level inpainting.

In order to illustrate the flexibility of MMM, we make available <sup>1</sup> examples generated by the system, and the interactive demo shown in Figure 2. The demo was developed in Google Colab, making it accessible to all users with a compatible internet browser. The interface automatically selects the appropriate model, either MMMBar or MMMTrack, based on the bars or tracks that are selected for generation. We briefly outline the various ways that one can interact with MMM when generating musical material.

<sup>1</sup> <https://jeffreyjohnens.github.io/MMM/>

1. Track Inpainting : Given a possibly empty set of tracks  $\mathbf{t} = \{t_1, \dots, t_k\}$ , we can generate  $n$  additional tracks. When the set of tracks is empty, this is equivalent to unconditioned generation. To do this, we condition the model with the tokens representing  $k$  tracks and then sample until the  $n^{th}$  TRACK\_END token is reached.
2. Bar Inpainting : Given a set of tracks  $\mathbf{t} = \{t_1, \dots, t_k\}$  and a set of bars  $\mathbf{b} = \{b_1, \dots, b_n\}$ , we can resample each bar in  $\mathbf{b}$ . For this method, we condition the model with the tokens representing all the tracks, replacing each  $b_i$  in  $\mathbf{b}$  with the FILL\_PLACEHOLDER. Then we sample until the  $n^{th}$  FILL\_END token is reached.
3. Attribute Control for Instruments: We can specify a set of MIDI instruments for each generated track, from which the model will choose. Practically, this is accomplished by masking the MIDI instruments we wish to avoid before sampling the INSTRUMENT token at the start of a new track.
4. Attribute Control for Note Density : We can specify the note density level for each generated track.
5. Iterative Generation : The user can chain together various generation methods to iteratively compose a piece of music. Alternatively, generation methods can be chained automatically using a meta-algorithm. For example, given a set of tracks  $\mathbf{t} = \{t_1, \dots, t_k\}$ , we can progressively resample each track  $t_i$  in  $\mathbf{t}$  by asking the model to generate  $(t_i | \{t_j : t_j \in \mathbf{t}, j \neq i\})$  for each  $1 \leq i \leq k$ . This bears some similarity to Gibbs sampling. The resulting output should be more similar to the input than simply generating a set of tracks from scratch.

## 7 Conclusion

One current limitation, is that the model only allows for a fixed number of bars to be generated. However, with some slight modifications, it should be possible to represent much longer sequences. The transformer-XL architecture (Dai et al., 2019) allows for extremely distant tokens to influence the current prediction via a hidden state, combining the strengths of the attention and recurrent mechanisms. Using this type of model, the current  $n$ -bar window could be conditioned on previous and future (if they are known)  $n$ -bar windows via the hidden state, affording rich conditional control over an arbitrary local window within a larger piece. In addition, future work will include a more rigorous evaluation of the system itself. We have introduced a novel approach to representing musical material that offers increased control over the generated output. This offers a new and exciting avenue for future work, harnessing the strengths of the transformer architecture to provide fine-grained control for the user at generation time.



## References

- Boulanger-Lewandowski, N., Bengio, Y., & Vincent, P. (2012). Modeling temporal dependencies in high-dimensional sequences: Application to polyphonic music generation and transcription. *International Conference on Machine Learning*. Retrieved from <https://arxiv.org/abs/1206.6392>
- Briot, J., Hadjeres, G., & Pachet, F. (2019). *Deep learning techniques for music generation*. Springer International Publishing. Retrieved from <https://books.google.ca/books?id=flrswEACAAJ>
- Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q. V., & Salakhutdinov, R. (2019). Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*.
- Donahue, C., Mao, H. H., Li, Y. E., Cottrell, G. W., & McAuley, J. (2019). Lakhnes: Improving multi-instrumental music generation with cross-domain pre-training. *arXiv preprint arXiv:1907.04868*.
- Dong, H.-W., Hsiao, W.-Y., Yang, L.-C., & Yang, Y.-H. (2018). Musegan: Multi-track sequential generative adversarial networks for symbolic music generation and accompaniment. In *Thirty-second aaai conference on artificial intelligence* (pp. 34–41). Retrieved from <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/viewPaper/17286>
- Huang, C. A., Cooijmans, T., Roberts, A., Courville, A. C., & Eck, D. (2017). Counterpoint by convolution. In *Proceedings of the 18th international society for music information* (pp. 211–218). Retrieved from [https://ismir2017.smcnus.org/wp-content/uploads/2017/10/187\\_Paper.pdf](https://ismir2017.smcnus.org/wp-content/uploads/2017/10/187_Paper.pdf)
- Huang, C. A., Vaswani, A., Uszkoreit, J., Simon, I., Hawthorne, C., Shazeer, N., ... Eck, D. (2019). Music transformer: Generating music with long-term structure. In *7th international conference on learning representations*.
- Khandelwal, U., He, H., Qi, P., & Jurafsky, D. (2018). Sharp nearby, fuzzy far away: How neural language models use context. *arXiv preprint arXiv:1805.04623*.
- Liang, F. (2016). *Bachbot: Automatic composition in the style of bach chorales* (Unpublished doctoral dissertation). University of Cambridge.
- Oore, S., Simon, I., Dieleman, S., Eck, D., & Simonyan, K. (2018). This time with feeling: learning expressive musical performance. *Neural Computing and Applications*, 1–13.
- Pati, A., Lerch, A., & Hadjeres, G. (2019). Learning to traverse latent spaces for musical score inpainting. *arXiv preprint arXiv:1907.01164*.
- Payne, C. (2019, April). Musenet. *OpenAI*. ([openai.com/blog/musenet](https://openai.com/blog/musenet))
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & 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* (Unpublished doctoral dissertation). Columbia University.

- Roberts, A., Engel, J. H., Raffel, C., Hawthorne, C., & Eck, D. (2018). A hierarchical latent vector model for learning long-term structure in music. In *Proceedings of the 35th international conference on machine learning* (pp. 4361–4370). Retrieved from <http://proceedings.mlr.press/v80/roberts18a.html>
- Sturm, B. L., & Ben-Tal, O. (2017). Taking the models back to music practice: Evaluating generative transcription models built using deep learning. *Journal of Creative Music Systems*, 2(1).
- Sturm, B. L., Ben-Tal, O., Monaghan, Ú., Collins, N., Herremans, D., Chew, E., ... Pachet, F. (2019). Machine learning research that matters for music creation: A case study. *Journal of New Music Research*, 48(1), 36–55.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems* (pp. 5998–6008).
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., ... Brew, J. (2019). Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, *abs/1910.03771*.