

PRUNING SUBSEQUENCE SEARCH WITH ATTENTION-BASED EMBEDDING

Colin Raffel and Daniel P. W. Ellis

LabROSA
Columbia University
New York, NY

ABSTRACT

Searching a large database to find a sequence that is most similar to a query can be prohibitively expensive, particularly if individual sequence comparisons involve complex operations such as warping. To achieve scalability, “pruning” heuristics are typically employed to minimize the portion of the database that must be searched with more complex matching. We present an approximate pruning technique which involves embedding sequences in a Euclidean space. Sequences are embedded using a convolutional network with a form of attention that integrates over time, trained on matching and non-matching pairs of sequences. By using fixed-length embeddings, our pruning method effectively runs in constant time, making it many orders of magnitude faster than full dynamic time warping-based matching for large datasets. We demonstrate our approach on a large-scale musical score-to-audio recording retrieval task.

Index Terms— Sequence Retrieval, Attention-Based Models, Convolutional Networks, Dynamic Time Warping, Pruning

1. INTRODUCTION

Sequences¹ are a natural way of representing data in a wide range of fields, including multimedia, environmental sciences, natural language processing, and biology. The proliferation of sensors and decreasing cost of storage in recent years has resulted in the creation of very large databases of sequences. Given such a database, a fundamental task is to find the database entry which is the most similar to a query. However, this task can be prohibitively expensive depending on both the method used to compare sequences and the size of the database.

An effective way to compare sequences is the dynamic time warping (DTW) distance, first proposed in the context of speech utterances [1]. DTW first computes an optimal alignment of the two sequences of feature vectors being compared, and then reports their distance as the total distance between aligned feature vectors. The alignment step makes DTW much more robust to timing and rate differences than more common metrics such as Euclidean distance [2]. DTW also naturally extends to the setting where subsequence matching is allowed. Using dynamic programming, the optimal DTW alignment can be found in time quadratic in the length of the sequences, but even with this optimization, naive application of DTW to nearest-neighbor retrieval can be prohibitively expensive. Other dynamic programming-based “edit distance” measures are used when appropriate, and are similar in both their effectiveness and expense.

To mitigate this issue, a variety of “pruning methods” have been proposed. The idea behind these techniques is to use heuristics to limit the most expensive comparisons to a small portion of the database. [2] gives a broad overview of different pruning methods, and shows that their application makes exact retrieval feasible in extremely large databases (e.g. trillions of sequences). However, these methods rely on various assumptions, including that the query sequence is a subsequence of its correct match and the length of the alignment path is known a priori.

An alternative approach is to replace the DTW search with a surrogate problem that is faster to solve. For example, [3] proposes a technique which maps sequences to a fixed-length embedding to avoid DTW calculation when matching the query to the database. The embedding is constructed by pre-computing the DTW distance between each sequence in the database and a small collection of “reference” sequences that are chosen to optimize retrieval accuracy. Then, to match a query sequence to the database, its DTW distance is computed against the reference sequences, and the resulting vector of DTW distances is matched against the database of embeddings using a standard Euclidean distance-based nearest-neighbor search. The resulting algorithm is approximate, but provided state-of-the-art results both in terms of accuracy and speed. This technique relies on the same assumptions as the pruning methods described in [2]. In addition, it presents a trade-off between efficiency and accuracy, since more reference sequences generally improve accuracy but require more full DTW calculations to be made, increasing the expense for very long or very high-dimensional sequences.

In [4], a more general approximate method is proposed where a neural network model is used to map and downsample sequences of feature vectors to shorter sequences of bit vectors. Computing the distance between bit vectors can be achieved by a single exclusive-or operation followed by a table lookup, so this approach vastly improves the efficiency of the local distance calculations involved in DTW. Furthermore, downsampling the sequences gives quadratic speed gains. In an experiment matching transcriptions of musical pieces to audio recordings in a large database, using this method allowed for 99% of the entries in the database to be pruned with high confidence. However, the efficiency of this approach is still limited by its reliance on DTW.

Motivated by the above issues, we propose a learning-based system for producing sequence embeddings for approximate matching. Our approach is similar in spirit to [3], except that it is fully general, i.e. it does not rely on any assumptions about the alignment length or whether the query is a subsequence of its match. This is thanks to the fact that we utilize an attention-based neural network model which can adapt to any problem setting according to the training data provided. Our approach is also constant-time in the query length: By mapping sequences to a fixed-length embedding, comparing a query to each database entry is a single Euclidean dis-

This work was supported by NSF project IIS-1117015.

¹In this work, we refer to “signals” with the more generic term “sequences”, but the two can be used interchangeably.

tance calculation, which does not become less efficient for longer or higher-dimensional sequences.

In the next section we discuss embedding and attention-based neural networks, which we use to motivate our proposed model in Section 3. In Section 4, we evaluate the effectiveness of our proposed model on the task of matching musical transcriptions to audio recordings of their corresponding songs in a large database. Finally, we discuss results and possibilities for improvement in Section 5.

2. EMBEDDING AND ATTENTION

The idea behind “embedding” is attractive: Given data in a representation which is either inefficient or does not readily reveal factors of interest, can we produce a mapping to fixed-length vectors that possess these desired properties? Because many simple machine learning methods are particularly effective given data where Euclidean distance corresponds to some notion of similarity, embedding can make data more amenable to learning and comparison.

As a concrete example, in [5], a neural network architecture is used to embed images of faces such that images of the same person have small pairwise distances in the embedded space and images of different people have large distances. The authors were able to achieve state-of-the-art results on the “labeled faces the wild” dataset by simply thresholding the resulting distances to denote “same person” or “different person”. In addition, the embedded representation produced qualitatively useful results when performing unsupervised clustering using k -means. Another striking application of embedding is word2vec, which is a family of models for mapping words to a Euclidean space with desirable properties [6]. For example, the difference between “China” and “Beijing” in the embedded space is roughly equivalent to the difference between “Rome” and “Italy”.

The idea of embedding has also been applied to sequences. Notably, in [7], a recurrent network was trained to map a sequence of words in a source language to a fixed-length vector, which was then used to generate a sequence of words in a target language, resulting in a quality translation of the original sentence. The resulting source-sentence embeddings encoded high-level linguistic characteristics, such as invariance to word order in sentences with the same meaning. A benefit of using recurrent networks is their ability to summarize an entire sequence of arbitrary length as a fixed-length vector, which brings the possibility of learning embeddings to sequences.

Recurrent networks, however, have difficulty in modeling long-term dependencies [8]. In the embedding setting, this may result in the end of the sequence having a greater impact on the embedding than the beginning. A recent technique for mitigating this issue has been dubbed “attention” [9]. Conceptually, an attention-based model includes a mechanism which learns to assign a weight to each sequence step based on the current and previous feature vectors. When used with recurrent networks, the addition of attention has proven effective in a wide variety of domains including speech recognition, machine translation, and image captioning [10].

Attention-based models may be compared to the common technique of summarizing a sequence of feature vectors by the mean, variance, and occasionally covariance of feature dimensions. This simple approach has seen frequent application in the domain of music information retrieval, where it has proven to be surprisingly effective at genre classification [11], instrument classification [12], artist recognition [13], and similarity rating prediction [14]. Attention-based models can be seen as a generalization of this technique where both the feature representation and the weighting used in constructing the average are optimized and data-adaptive.

3. FEED-FORWARD ATTENTION

While previous work on attention-based models has focused on recurrent networks, in the present work we will use feed-forward networks for the following reasons: First, recurrent networks have difficulty modeling very long-term dependencies [8]. In the experiment described in Section 4, sequences may have thousands of time steps, which is prohibitively long even for sophisticated models such as long short-term memory networks [15]. Second, feed-forward networks are much more efficient to train and evaluate than recurrent networks because their operations can be completely parallelized, whereas recurrent networks must evaluate each time step sequentially. Finally, attention provides its own form of temporal modeling because it integrates over the entire sequence. We therefore propose a “feed-forward” variant of attention, which computes a weighting for each sequence step independent of all other steps.

Our simplified variant of attention can be formulated as follows: Given a matrix which comprises a sequence of M D -dimensional vectors $X \in \mathbb{R}^{M \times D}$, we compute weights $\omega \in \mathbb{R}^M$ using the following transformation:

$$\omega = \text{softmax}(\sigma(Xw + b)) \quad (1)$$

where

$$\text{softmax}(x)_j = \frac{e^{x_j}}{\sum_{i=1}^M e^{x_i}}$$

and σ is some nonlinear function (throughout this work, we will use $\sigma = \tanh$). $w \in \mathbb{R}^D$ and $b \in \mathbb{R}$ are parameters of the transformation, which can be learned as part of a larger feed-forward network model. After computing ω , the weighted average of the vectors in X is computed by

$$\hat{X} = \sum_{i=1}^M \omega_i X_{i,:}$$

Conceptually, this attention mechanism can be thought of as using w , b , and σ to compute scalar “importance” values for each vector in X , which are then normalized using the softmax function and used to compute a weighted mean \hat{X} over all vectors in X .

The main difference between this formulation and the one proposed in [9, 10] is that here we are only producing a single weighting ω rather than a different ω at each time step. This is primarily because our goal in embedding is to produce a single vector for the entire sequence whereas previous applications of attention have mainly focused on sequence-to-sequence transduction. As a result, Equation 1 does not contain any terms for a previous attention vector or a previous hidden state, because in the present contexts these quantities do not exist. The main disadvantage to using attention in this way is that it ignores temporal order by computing an average over the entire sequence. However, motivated by the surprising effectiveness of the unweighted/non-adaptive Gaussian integration results discussed in Section 2, we submit that this approach will be effective and verify its utility in the following section.

4. PRUNING EXPERIMENT

Recall that the goal of this work is to employ embeddings to speed up large-scale sequence similarity searches by avoiding comparisons for most sequences in a database. Towards this end, we will utilize embeddings as follows: First, given a training set of matching sequences, we will train a feed-forward network with attention to embed each sequence in a Euclidean space where similar sequences have a small distance and dissimilar ones have a large distance.

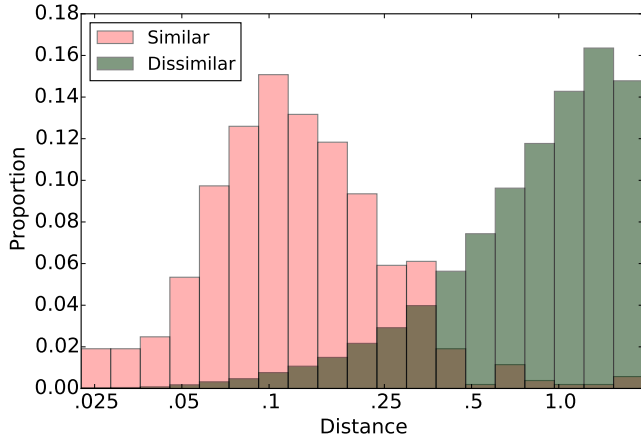


Fig. 1. Histograms of embedded distances for positive and negative example pairs in the validation set for our best-performing embedding network (c.f. [4] Figure 3). Note that the x-axis is log-scaled.

Then, we will utilize this network to pre-compute embeddings for all sequences in a database. In order to match a query sequence, we will compute the distance between its embedding and all pre-computed embeddings from the database. Given a test set of correct matches which are known a priori, the effectiveness of this approach can be measured by determining how often the correct match appears in the N entries with the N smallest embedded distances for varying values of N .

We will evaluate this approach on the problem proposed in [4]. This work seeks to efficiently match a transcription of a piece of music to a recording of the same piece in a large database of recordings. The resulting problem boils down to a large-scale subsequence matching problem, where the sequences are particularly long (up to thousands of time steps), the feature dimensionality is relatively high, and the correct match may be a subsequence of the query and vice versa. These properties make the traditional DTW pruning methods outlined in Section 1 ill-suited and provide a setting to evaluate completely general-purpose methods like the one proposed in this work. We will outline the details of the experiment below; for a complete discussion, we refer the reader to [4].

The training set consists of a collection of MIDI files which have been matched to audio files. MIDI files can be thought of as transcriptions, and can be synthesized in order to obtain an audio recording which is qualitatively similar to the song they are a transcription of. In [4], the MIDI and audio pairs are also aligned in time to one another using DTW before being used as training data; the system proposed in this work doesn’t require sequences which are aligned in time so we skipped this step. The resulting training set contains 5,536 sequence pairs.

As a representation, we use log-magnitude constant-Q spectrograms [16] of the audio files and MIDI syntheses. Each constant-Q spectrum spanned four octaves, using twelve bins per octave starting from MIDI note C2 (65.4 Hz). Each frequency bin in all feature sequences was z-scored (standardized) by its mean and standard deviation over the training set. In [4], using shorter sequences was beneficial so constant-Q spectra were aggregated over estimated beat locations. The present approach is meant to be able to cope with sequences which may be oversampled in time, so we did not beat-aggregate feature vectors and instead used spectra computed every 46 milliseconds. All feature analysis and processing was accom-

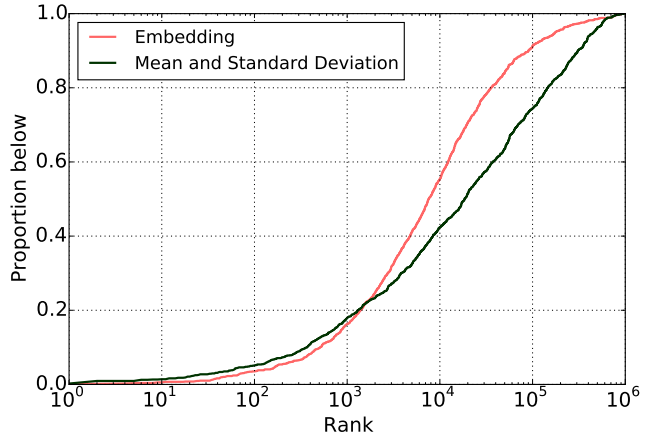


Fig. 2. Proportion of MIDI files in the test set whose correct match appeared below a given rank, based on embedding distances.

plished with `librosa` [17, 18].

For a model, we follow [4] by using a siamese architecture consisting of two parallel convolutional networks followed by two parallel fully-connected networks. The use of a convolutional “front end” allows our model to exploit both time and frequency invariances in the constant-Q spectra while reducing the number of model parameters. The main difference in models is the addition of the attention mechanism proposed in Section 3, which we place between the convolutional network and the fully-connected network. In this way, the convolutional network can be seen as transforming a “raw” input sequence to a sequence of vectors in a learned representation, which are then aggregated using the attention mechanism and mapped to the embedded space using the fully-connected network. We used an embedding space dimensionality of 128.

To train the networks, we pass pairs of “matching” and “non-matching” sequences and utilize a loss function which encourages matching pairs to have similar embeddings and non-matching pairs to have dissimilar ones. We use the same loss function as in [4]: Given a pair of positive sequences X_p, Y_p and a pair of negative sequences X_n, Y_n , we attempt to minimize

$$\mathcal{L} = \|f(X_p) - g(Y_p)\|_2^2 + \alpha \max(0, m - \|f(X_n) - g(Y_n)\|_2)^2$$

where f and g are the functions of the two networks being trained, α is a parameter controlling the importance of separating dissimilar pairs, and m is the target separation for dissimilar pairs. To construct negative examples, we simply randomly choose non-matched MIDI/audio sequence pairs from the training set.

As in [4] we utilize rectified linear units throughout the network, except on the final output layer which used tanh units. All weights were initialized using He et. al.’s method [19] and all biases were initialized to zero. The networks were optimized with respect to \mathcal{L} using Nesterov’s accelerated gradient [20] with mini-batches of 20 similar and dissimilar pairs. All sequences were either randomly cropped or padded to the median sequence length in each modality. Every 100 minibatches, we computed the mean loss \mathcal{L} on a held-out validation set consisting of 10% of the training set; when it was less than .99 times the previous lowest validation loss, we trained for 500 more minibatches. We implemented our model using `lasagne` [21], which is built on `theano` [22, 23].

To choose hyperparameter settings, we used the Bayesian optimization software package `spearmint` [24]. We optimized the

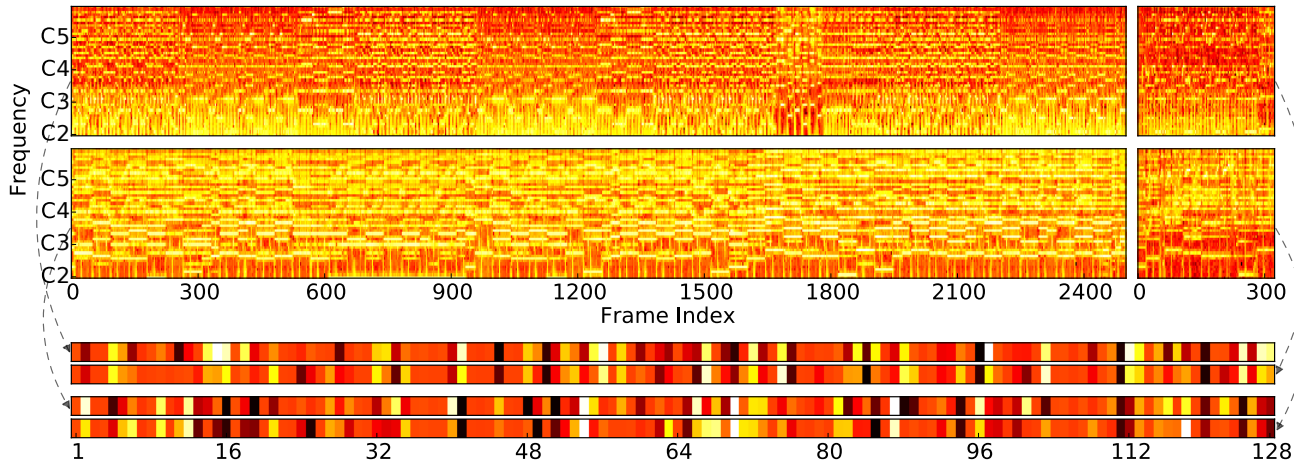


Fig. 3. Embeddings for two pairs of matching sequences from our best-performing system. The first two rows display pairs of matching constant-Q spectrograms. For both pairs, the right sequence is a subsequence of the left. The last two rows show their resulting embeddings. Dashed arrows denote which embedding corresponds to which sequence.

number of convolutional and pooling layers, the number of fully-connected layers, the learning rate and momentum, and the loss parameters α and m . We also experimented with using dropout [25] in the fully-connected layers for regularization, but did not find it to be beneficial. As an objective, we computed the Bhattacharyya coefficient [26] between the distributions of matching and non-matching pairs in the validation set. The best-performing network configuration achieved a Bhattacharyya coefficient of .471 and used a single convolutional layer with 16 filters which were 5 time steps wide by 12 semitones tall, one 2x2 max-pooling layer, two fully connected hidden layers with 2048 units each, a learning rate of 10^{-4} , a momentum of 0.997, $\alpha = 0.935$ and $m = 1$. The validation set distance distributions produced by this network after training can be seen in Figure 1.

5. RESULTS

Some qualitative intuition about our system’s behavior can be gained by inspecting the embeddings resulting from matching and non-matching pairs of sequences. Embeddings for two different matching sequence pairs from the validation set can be seen in Figure 3. Among the four embeddings displayed, the pairs of matched embeddings resulted in Euclidean distances of 0.159 and 0.347, while non-matching pairs had distances of 1.08 and 1.15. While these distances indicate the appropriate behavior, there are nevertheless some noticeable discrepancies between the embeddings of the matched sequences.

To evaluate our system’s performance on the task of pruning large-scale subsequence search, we used it to match a held-out test set of 1,282 MIDI files to entries in the Million Song Dataset (MSD) [27]. Our test set was the same used in [4], for which the correct match in the MSD is known for each MIDI file. We first computed embeddings of constant-Q spectrograms for all of the short (30 to 60 second) preview recordings of each entry in the MSD provided by 7digital [28]. Then, we computed a constant-Q spectrogram and its embedding for each MIDI file in the test set. To determine performance, we computed the Euclidean distance between the embedding of each MIDI file and the embeddings of every entry in the MSD and determined how many entries had a distance that was smaller than

the correct match. As a baseline, we performed the same evaluation except using the concatenated mean and standard deviation of the feature dimensions of constant-Q spectrograms as the embedding. The results can be seen in Figure 2.

To summarize these results, using the proposed method the correct match appeared in the top 100,000 matches 91.3% of the time using our learned embedding and only 74.4% of the time using the mean and standard deviation of spectrograms. As noted in [4], matching a single MIDI file to the MSD using dynamic time warping distance would take about three hours; the proposed approach takes about 360 milliseconds per file using a non-optimized Python implementation. Another way of stating our result, then, is that we can discard 90% (900,000 entries) of the MSD with 91.3% confidence using a pruning method which is about 30,000 times faster than the standard DTW-based matching approach. Compared to the approach proposed in [4] which can prune 99% of the MSD with 95.9% confidence, it is substantially less accurate, but it is also over 300 times faster. In addition, our approach is fully general, and was able to achieve good performance on a task that would not be possible using the methods proposed in [2] or [3].

While our proposed approach has proven effective, there is room for improvement. The model may be able to further differentiate between similar and dissimilar sequences if the dimensionality of the embedded space was increased, at the cost of a linear increase in runtime. Alternatively, choosing negative example pairs more carefully (e.g. using the “harmonic embedding” approach of [5]) may further improve discrimination. Finally, evaluating this approach on more traditional sequence retrieval tasks which allow more stringent assumptions to be made about the correct match would facilitate comparison to existing pruning techniques. For researchers interested in utilizing and reproducing our results, code is available online.²

6. REFERENCES

- [1] Hiroaki Sakoe and Seibi Chiba, “Dynamic programming algorithm optimization for spoken word recognition,” *IEEE Trans-*

²<http://github.com/craffel/sequence-embedding>

- actions on Acoustics Speech and Signal Processing, vol. 26, no. 1, pp. 43–49, 1978.
- [2] Thanawin Rakthanmanon, Bilson Campana, Abdullah Mueen, Gustavo Batista, Brandon Westover, Qiang Zhu, Jesin Zakaria, and Eamonn Keogh, “Searching and mining trillions of time series subsequences under dynamic time warping,” in *Proceedings of the 18th ACM SIGKDD international conference on knowledge discovery and data mining*, 2012, pp. 262–270.
 - [3] Panagiotis Papapetrou, Vassilis Athitsos, Michalis Potamias, George Kollios, and Dimitrios Gunopulos, “Embedding-based subsequence matching in time-series databases,” *ACM Transactions on Database Systems*, vol. 36, no. 3, pp. 17, 2011.
 - [4] Colin Raffel and Daniel P. W. Ellis, “Large-scale content-based matching of MIDI and audio files,” in *Proceedings of the 16th International Society for Music Information Retrieval Conference*, 2015, pp. 234–240.
 - [5] Florian Schroff, Dmitry Kalenichenko, and James Philbin, “Facenet: A unified embedding for face recognition and clustering,” *arXiv preprint arXiv:1503.03832*, 2015.
 - [6] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean, “Distributed representations of words and phrases and their compositionality,” in *Advances in Neural Information Processing Systems*, 2013, pp. 3111–3119.
 - [7] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le, “Sequence to sequence learning with neural networks,” in *Advances in Neural Information Processing Systems*, 2014, pp. 3104–3112.
 - [8] Yoshua Bengio, Patrice Simard, and Paolo Frasconi, “Learning long-term dependencies with gradient descent is difficult,” *IEEE Transactions on Neural Networks*, vol. 5, no. 2, pp. 157–166, 1994.
 - [9] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
 - [10] Kyunghyun Cho, Aaron C. Courville, and Yoshua Bengio, “Describing multimedia content using attention-based encoder-decoder networks,” *arXiv preprint arXiv:1507.01053*, 2015.
 - [11] George Tzanetakis and Perry Cook, “Musical genre classification of audio signals,” *IEEE transactions on Speech and Audio Processing*, vol. 10, no. 5, pp. 293–302, 2002.
 - [12] Jeremiah D. Deng, Christian Simmermacher, and Stephen Crane-field, “A study on feature analysis for musical instrument classification,” *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 38, no. 2, pp. 429–438, 2008.
 - [13] Michael I. Mandel and Daniel P. W. Ellis, “Song-level features and support vector machines for music classification,” in *Proceedings of the 6th International Conference on Music Information Retrieval*, 2005, pp. 594–599.
 - [14] Peter Foster, Matthias Mauch, and Simon Dixon, “Sequential complexity as a descriptor for musical similarity,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 12, pp. 1965–1977, 2014.
 - [15] Sepp Hochreiter and Jürgen Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
 - [16] Judith C. Brown, “Calculation of a constant Q spectral transform,” *The Journal of the Acoustical Society of America*, vol. 89, no. 1, pp. 425–434, 1991.
 - [17] Brian McFee, Matt McVicar, Colin Raffel, Dawen Liang, Oriol Nieto, Josh Moore, and Dan P. W. Ellis et. al., “librosa: v0.4.0,” <https://github.com/bmcfee/librosa>, 2015.
 - [18] Brian McFee, Colin Raffel, Dawen Liang, Daniel P. W. Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto, “librosa: Audio and music signal analysis in python,” in *Proceedings of the 14th Python in Science Conference*, 2015.
 - [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” *arXiv preprint arXiv:1502.01852*, 2015.
 - [20] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton, “On the importance of initialization and momentum in deep learning,” in *Proceedings of the 30th international conference on machine learning*, 2013, pp. 1139–1147.
 - [21] Sander Dieleman, Jan Schlüter, Colin Raffel, Eben Olson, and Soren Kaae Sonderby et al., “Lasagne: First release,” <https://github.com/Lasagne/Lasagne>, 2015.
 - [22] James Bergstra, Olivier Breuleux, Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, Guillaume Desjardins, Joseph Turian, David Warde-Farley, and Yoshua Bengio, “Theano: a CPU and GPU math expression compiler,” in *Proceedings of the Python for scientific computing conference (SciPy)*, 2010, vol. 4, p. 3.
 - [23] Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, James Bergstra, Ian Goodfellow, Arnaud Bergeron, Nicolas Bouchard, David Warde-Farley, and Yoshua Bengio, “Theano: new features and speed improvements,” in *Deep Learning and Unsupervised Feature Learning NIPS 2012 Workshop*, 2012.
 - [24] Jasper Snoek, Hugo Larochelle, and Ryan P. Adams, “Practical bayesian optimization of machine learning algorithms,” in *Advances in Neural Information Processing Systems*, 2012, pp. 2951–2959.
 - [25] Geoffrey E. Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R. Salakhutdinov, “Improving neural networks by preventing co-adaptation of feature detectors,” *arXiv preprint arXiv:1207.0580*, 2012.
 - [26] Anil Kumar Bhattacharyya, “On a measure of divergence between two statistical populations defined by their probability distributions,” *Bulletin of the Calcutta Mathematical Society*, vol. 35, pp. 99–109, 1943.
 - [27] Thierry Bertin-Mahieux, Daniel P. W. Ellis, Brian Whitman, and Paul Lamere, “The million song dataset,” in *Proceedings of the 12th International Society for Music Information Retrieval Conference*, 2011, pp. 591–596.
 - [28] Alexander Schindler, Rudolf Mayer, and Andreas Rauber, “Facilitating comprehensive benchmarking experiments on the million song dataset,” in *Proceedings of the 13th International Society for Music Information Retrieval Conference*, 2012, pp. 469–474.