

Generating Music Medleys via Playing Music Puzzle Games

Yu-Siang Huang,[†] Szu-Yu Chou,^{†‡} Yi-Hsuan Yang[†]

[†]Research Center for IT innovation, Academia Sinica, Taiwan

[‡]Graduate Institute of Networking and Multimedia, National Taiwan University, Taiwan

Abstract

Generating music medleys is about finding an optimal permutation of a given set of music clips. Toward this goal, we propose a self-supervised learning task, called the music puzzle game, to train neural network models to learn the sequential patterns in music. In essence, such a game requires machines to correctly sort a few multisecond music fragments. In the training stage, we learn the model by sampling multiple non-overlapping fragment pairs from the same songs and seeking to predict whether a given pair is consecutive and is in the correct chronological order. For testing, we design a number of puzzle games with different difficulty levels, the most difficult one being music medley, which requiring sorting fragments from different songs. On the basis of state-of-the-art Siamese convolutional network, we propose an improved architecture that learns to embed frame-level similarity scores computed from the input fragment pairs to a common space, where fragment pairs in the correct order can be more easily identified. Our result shows that the resulting model, dubbed as the similarity embedding network (SEN), performs better than competing models across different games, including music jigsaw puzzle, music sequencing, and music medley. Example results can be found at our project website, <https://remyhuang.github.io/DJnet>.

Introduction

Recent years have witnessed a growing interest in unsupervised methods for sequential pattern learning, notably in the computer vision domain. This can be approached by the so-called *self-supervised learning*, which exploits the inherent property of data for setting the learning target. For example, (Misra et al. 2016), (Fernando et al. 2016) and (Lee et al. 2017) leveraged the temporal coherence of video frames as a supervisory signal and formulated representation learning as either an order verification or a sequence sorting task. (Lotter, Kreiman, and Cox 2017), on the other hand, explored prediction of future frames in a video sequence as the supervisory signal for learning the structure of the visual world. These prior arts demonstrate that learning discriminative visual features from massive unlabeled videos is possible.

From a technical standpoint, this paper studies how such a self-supervised learning methodology can be extended to




Figure 1: Similarity matrix between: (a) two identical music fragments, (b) fragments from a song and its cover version (i.e. the same song but different singer), (c) fragments from two different songs of the same singer. The goal of the proposed network is to learn patterns from such matrices.

audio, which has been less attempted. In particular, we focus on learning from music sequences. Music is known for its multilevel, hierarchical organization, with higher-level building blocks made up of smaller recurrent patterns (Widmer 2016; Hudson 2011). While listening to music, human beings can discern those patterns, make predictions on what will come next, and hope to meet their expectations. Asking machines to do the same is interesting on its own, and it poses interesting challenges that do not present, or have not been considered, in the visual domain.

First, the input instances to existing models are usually video frames (which are images) sampled from each video sequence. Each frame can be viewed as a snapshot of a temporal moment, and the task is to correctly order the frames per video. In contrast, meaningful basic unit to be ordered in music has to be an audio sequence itself. Therefore, the input instances in our case are multisecond, non-overlapping music fragments, which have a temporal dimension.

Second, existing works in the visual domain considered at most four frames per video (Lee et al. 2017), mainly due to the concern that the possible permutations increase exponentially along with the number of sampled frames. However, as a song is typically a few minutes long, we consider up to ten (multisecond) music fragments per song.

Lastly, while it makes less sense to order video frames sampled from different video sequences, for music it is interesting and practically useful if we can find an ordering of a bag of music fragments sampled from different songs.

Music Puzzle Games

Background and Significance

Indeed, music fragments of different songs, when properly ordered, can be listened to consecutively with pleasure (Lin et al. 2015), given that every pair of consecutive fragments share or follow some harmonic or melodic patterns. For example, Disc Jockeys (DJs) are professional audio engineers who can nicely perform such a *music medley generation* task. Therefore, we also include this task to evaluate the performance of our model. In doing so, from an application standpoint, this paper also contributes to the realization of automatic music medley generation, an important step toward making an AI DJ (Huang, Chou, and Yang 2017a). The hope is that one day AI would possess certain level of music connoisseurship and can serve as a DJ to create music remixes or mashups professionally.

Drawing the analogy that a fragment is like a puzzle piece, we propose to refer to the task of assembling multiple music fragments in proper order as the *music puzzle games*. Similar to previous work in the visual domain, we exploit the temporal coherence of music fragments as the supervisory signal to train our neural networks via a music puzzle game. What’s different is that we differentiate four aspects of a music puzzle game and investigate the performance of our models with different types of games. The four aspects are: 1) number of fragments to be ordered, 2) temporal length of the fragments (whether the length is fixed or arbitrary), 3) whether there is a clear cut at the boundary of fragment pairs, and 4) whether the fragments are from the same song or not. For example, other than uniformly sample a song for fragments, we also employ downbeat tracking (Böck et al. 2016) to create musically meaningful fragments (Upham and Farbood 2013).

In view of the second challenge mentioned above, we propose to take fragment pairs as input to our neural network models and lastly use a simple heuristic to decide the final ordering. For a music puzzle game with n fragments, this pair-wise approach requires our models to evaluate in total $2 \cdot \binom{n}{2} = n(n-1)$ pairs, which is much fewer than the $n!$ number of possible permutations and accordingly opens up the possibility to consider $n > 4$ fragments.

Moreover, in view of the first challenge mentioned above, we propose an novel model, called the similarity embedding network (SEN), to solve the music puzzle games. The main idea is to compute frame-level similarity scores between each pair of short-time frames from the two input fragments, and then learn to embed the resulting similarity matrix (Serrà et al. 2012) into a common space, where coherent and incoherent fragment pairs can be more easily distinguished.

The idea of learning from the similarity matrices has roots in the so-called recurrence plot (Marwan et al. 2007), which provides a way to visualize the periodic nature of a trajectory (i.e. a time-series data) through a phase space. Given two trajectories, we can similarly compute their point-by-point (or frame-by-frame) similarity to mine patterns from the resulting matrix. For our tasks, learning from the similarity matrices is promising, for we can examine temporal correspondence between fragments in more details, as suggested by the example similarity matrices shown in Figure 1. Our experiments show that SEN performs consistently better than competing models across different music puzzle games.

In academia, some researchers have investigated the design of music-based puzzle games, mostly for education purposes. A notable example is the work presented by (Hansen et al. 2013), which experimented with a number of designs of sound-based puzzle game to train the listening abilities of visually impair people. A music clip was divided into several fragments and a player had to rearrange them in order to reconstruct the original song. For advanced players, they further applied pitch and equalization shifts randomly on fragments, requiring the players to detect those transpositions to complete the puzzle. However, in this work a music clip was divided into pieces at arbitrary timepoints. This way, there may be no clear cut at the boundary of the fragments, providing strong temporal cues that make the game easier: when the fragments are in incorrect order, the result will not only sound unmusical but also unnatural.

More recently, (Smith et al. 2017) improved upon this work by dividing songs at downbeat positions, which often coincide with chord changes (Böck, Krebs, and Widmer 2016) and provides clearer cut among the fragments. Moreover, based on a “mashability” measure (Davies et al. 2014), they proposed an algorithm to create cross-song puzzles for more difficult games. They claimed that the game can train the musical and logical reasoning of ordinary people.

Although these previous works are interesting, their focus is on the design of the puzzle games for human beings, rather than on training machines to solve such games. In contrast, we let machine learn sequential patterns (and logic) in the musical world in a self-supervised learning manner by playing and solving such games.

Another benefit of experimenting with the music puzzle games is that the input to such games are sequences (not images). Therefore, similar network architecture may be applied to time series data in other domains as well.

In what follows, we firstly discuss the design of music puzzle games for machines, and then present a mathematical formulation of the learning problem.

Game Design

As shown in Table 1, we consider four aspects in designing the music puzzle games. First, the number of fragments n to be ordered; larger n implies more computational cost and a more difficult game, as potentially more orderings of the fragments would look plausible. Second, whether the length of the fragments in a game is the same. Our machine model has to deal with input sequences of arbitrary length, if the

Table 1: Characteristics of the music fragments that are to be ordered by our model in various music puzzle games

	Jigsaw Puzzle	Sequencing	Medley
number	3, 4, 6, 8	10	7–11
length	fixed / arbitrary	arbitrary	arbitrary
boundary	unclear / clear	clear	clear
from	same song	same song	cross song




Figure 2: Illustration of the proposed similarity embedding network and its application to solving music jigsaw puzzle.

length of the fragments is not fixed. Third, whether there is a clear cut at the boundary of fragment pairs. Arbitrary segmentation of a song may lead to obvious continuity at the boundary of two fragments and make the game too easy. It is conceivable that when the boundaries are clearer, the puzzle game will be more difficult. Fourth, whether the fragments are from the same song or not.

As also shown in Table 1, we consider three different music puzzle games in this paper, with progressively increasing level of difficulty. For **music jigsaw puzzle**, we create fragments by dividing a 24-second music clip at either equally-spaced or at downbeat-informed timepoints. Because we are interested in comparing the performance of the proposed SEN model against those proposed to solve video puzzles, we vary the value of n from 3 to 8 in this game.

The second game is more difficult in that the fragments are taken from a *whole song*. Moreover, each fragment represents a *section* of the song, such as the intro, verse, chorus and bridge (Paulus, Miller, and Klapuri 2010; Nieto and Bello 2016). The game is challenging in that the verse and chorus section may repeat multiple times in a song, with sometimes minor variations. The boundaries are clear, and we use $n = 10$ fragments (sections) per song. In audio engineering, the task of arranging sections in a sensible way is referred to as **music sequencing**.

Lastly, we consider the **music medley** game, which aims to put together a bag of short clips from different songs to build a longer music piece (Lin et al. 2015). As the fragments (clips) are from different songs, the boundaries are also clear. This is different from the cross-song puzzle considered in (Smith et al. 2017). In music medley, we take *one* fragment per song from $m (= n)$ songs, and aim to create *an* ordering of them. In contrast, in cross-song puzzle, we take $\lfloor n/m \rfloor$ fragments per song from $m (\neq n)$ songs and aim to discern the origin of the fragments and get m orderings.

Problem Formulation

All the aforementioned games are about ordering things. While solving an image jigsaw puzzle game, human beings usually consider the structural patterns and texture information as cues by comparing the puzzle pieces one by one (Noroozi and Favaro 2016). There is no need to put all the pieces in correct order all at once. As the number of permutations grows exponentially with n , we formulate the learn-

ing problem as a binary classification problem and predict whether a given pair of fragments is consecutive and is in correct order.

In the training stage, all the fragments are segmented consecutively without overlaps per song, as shown in the left-most part of Figure 2. For each song, we get a collection of fragments $\{R_1, \dots, R_n\}$, which are in the correct order. Among the $2^{\binom{n}{2}}$ possible fragments pairs, $n - 1$ of them are in the correct order and are considered as the **positive data**, $\mathcal{P}_+ = \{(R_i, R_{i+1}) \mid i \in \{1, 2, \dots, n - 1\}\}$. While all the other possible pairs can be considered as the **negative data**, we consider only three types of them:

$$\begin{cases} \mathcal{P}_-^{(1)} = \{(R_{i+1}, R_i) \mid i \in \{1, \dots, n - 1\}\}, \\ \mathcal{P}_-^{(2)} = \{(R_i, R_{i+2}) \mid i \in \{1, \dots, n - 2\}\}, \\ \mathcal{P}_-^{(3)} = \{(R_{i+2}, R_i) \mid i \in \{1, \dots, n - 2\}\}. \end{cases}$$

Pairs of the first type is consecutive but in incorrect order. Pairs of the second and third types are not consecutive. The negative data is the union of them: $\mathcal{P}_- = \mathcal{P}_-^{(1)} \cup \mathcal{P}_-^{(2)} \cup \mathcal{P}_-^{(3)}$. Therefore, the ratio of positive and negative data $|\mathcal{P}_+|/|\mathcal{P}_-|$ is about 1/3. In our experiments, we also refer to data pairs belonging to \mathcal{P}_+ , $\mathcal{P}_-^{(1)}$, $\mathcal{P}_-^{(2)}$ and $\mathcal{P}_-^{(3)}$ as ‘R1R2,’ ‘R2R1,’ ‘R1R3’ and ‘R3R1,’ respectively.¹

Given a training set $\mathcal{D} = \{(X, y) \mid X \in \mathcal{P}, y \in \{0, 1\}\}$, where \mathcal{P} is the union of \mathcal{P}_+ and \mathcal{P}_- from all the songs and y whether a pair is positive or not, we learn the parameters θ of a neural network f_θ by solving:

$$\min_{\theta} \sum_{(X, y) \in \mathcal{D}} \mathcal{L}(f_\theta(X), y) + \mathcal{R}(\theta), \quad (1)$$

where \mathcal{L} is a loss function (e.g. cross entropy) and $\mathcal{R}(\theta)$ is a regularization term for avoiding overfitting.

Global Ordering Given a data pair $X = (R_a, R_b)$, $a, b \in \{1, \dots, n\}$, $a \neq b$, the estimate $f_\theta(X)$ is a value in $[0, 1]$, due to a softmax function. For each song in the validation set, we need to get this estimate for all the data pairs, and seek to find

¹We note that, in related work working on videos (Misra et al. 2016; Fernando et al. 2016; Lee et al. 2017), they treated R1R2 the same as R2R1, and likewise R1R2R3 the same as R3R2R1, assuming that playing a short video clip in the reverse order is fine.

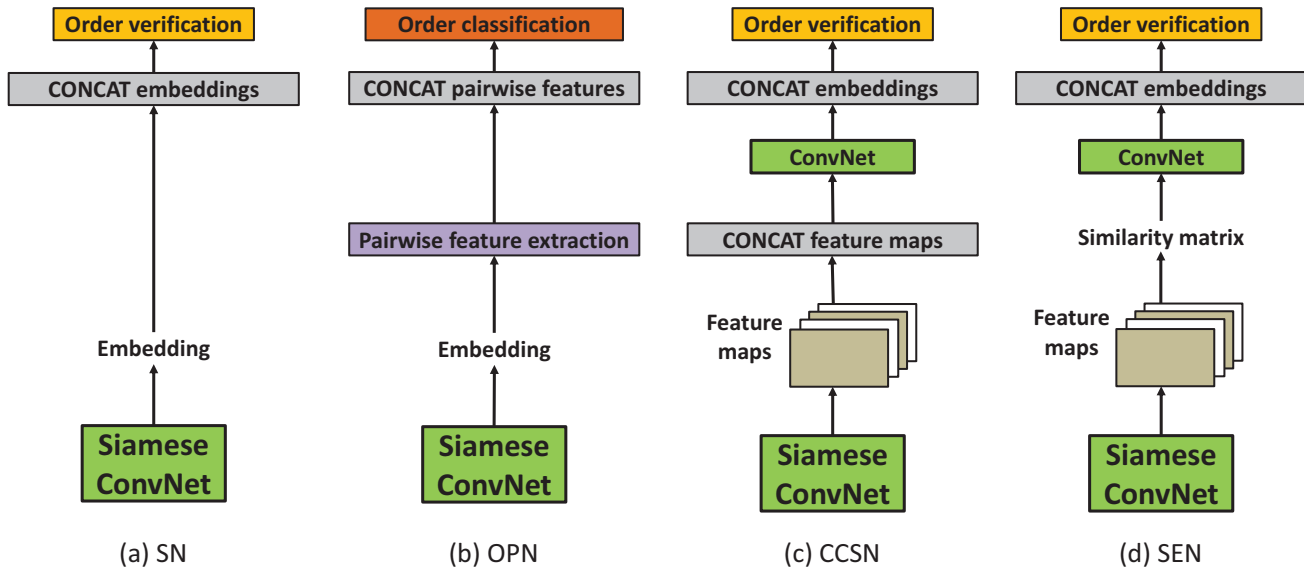


Figure 3: Four different network architectures used in our experiments: (a) A standard Siamese network (Misra et al. 2016) for order verification (i.e. binary classification); (b) the order prediction network (Lee et al. 2017), which extracts pairwise features from each embedding and concatenates all pairwise features for order classification (i.e. multiclass classification); (c) the concatenated-convolutions Siamese network, which captures joint frames by concatenating all feature maps; and (d) the proposed similarity embedding network, which learns structural patterns from the similarity matrices.

the correct global ordering of the fragments from these estimates.² While there may be other sophisticated ways doing it, we find the following simple heuristic works quite well already: we evaluate the “fitness” of any ordering of the fragments by summing the model output of the composing $n - 1$ consecutive pairs. For example, the fitness for (R_a, R_b, R_c) , for $n = 3$, will be $f_\theta(R_a, R_b) + f_\theta(R_b, R_c)$. We then simply pick the ordering with the highest fitness score as our solution for the game for that song.

Network Architecture

Similarity Embedding Network (SEN)

A Siamese network (Bromley et al. 1994) is composed of two (or more) twin subnetworks that share the same parameters. The subnetworks usually use convolutional layers (but there are exceptions (Mueller and Thyagarajan 2016)). The outputs of the last convolutional layer are concatenated and then feed to the subsequent fully-connected layers. The functions of the convolutional layers and the fully connected layers are feature learning and classifier training, respectively. Because Siamese networks can process multiple inputs at the same time, it is widely used in various metric learning problems (Chopra, Hadsell, and LeCun 2005).

As shown in the middle of Figure 2, the proposed SEN model also uses a convolutional Siamese network (Siamese ConvNet) to learn features from spectrogram-like 2D features of a pair of fragments. However, motivated by a re-

²Our model is trained by playing only the music jigsaw puzzle, but in testing time the model will be applied to different games.

cent work (Luo, Schwing, and Urtasun 2016), which used a product layer to compute the inner product between two representations of a Siamese network, we propose to compute the similarity matrix from the frame-by-frame output of the last layer of the Siamese ConvNet, and further learn features from the similarity matrix with a few more convolutional layers, as shown in the right hand side of Figure 2 (and Figure 3(d)). The output can be viewed as an embedding of the similarity matrix, therefore the name of the network.

Given the output feature maps of the Siamese ConvNet, $\mathbf{G}_a = h_\theta(R_a) \in \mathcal{R}^{N \times k}$, $\mathbf{G}_b = h_\theta(R_b) \in \mathcal{R}^{M \times k}$, where h_θ denotes the network up to the last layer of the Siamese ConvNet, N and M the (temporal) length of the output and k the dimension of the feature, the similarity matrix $\mathbf{S} \in \mathcal{R}^{N \times M}$ is computed by the Cosine score:

$$\mathbf{S}_{ij} = (\mathbf{g}_{a,i}^T \mathbf{g}_{b,j}) / (\|\mathbf{g}_{a,i}\|_2 \|\mathbf{g}_{b,j}\|_2), \quad (2)$$

where $\mathbf{g}_{a,i} \in \mathcal{R}^k$ is the i -th feature (slice) of \mathbf{G}_a .

Because we want the resulting similarity matrix to capture the temporal correspondence between the input fragments, in SEN we use 1D convolutions along the temporal dimension (Liu and Yang 2016) for the Siamese ConvNet. Moreover, we set the stride size to 1 and use no pooling layers in the Siamese ConvNet for SEN, to capture detailed temporal information of the fragments.

Baselines

Siamese CNN (SN) A valid baseline is the pairwise Siamese ConvNet, which takes the input fragment pairs and learns a binary classifier for order verification.

Concatenated-inputs CNN (CIN) An intuitive solver for the music jigsaw puzzle is to concatenate the spectrogram-like 2D features of the fragments along the time dimension, and use a CNN (instead of an SN) for order verification. We suppose this model can catch the weird boundary of an incorrectly ordered fragment pair.

Concatenated-convolutions Siamese Network (CCSN) This is a state-of-the-art network for image feature learning (Wang et al. 2016). Given the feature maps from the last convolutional layers of a Siamese ConvNet, we can simply concatenate them along the depth dimension (instead of computing the similarity matrix) and then use another stack of convolutional layers to learn features. As shown in Figure 3(c), the only difference between CCSN and SEN lies in how we extract information from the Siamese ConvNet.

Triplet Siamese Network (TSN) & Order Prediction Network (OPN) The state-of-the-art algorithms in solving video puzzle games use a *list-wise* approach instead of a *pair-wise* approach. The TSN model (Misra et al. 2016) is simply an expansion of SN by taking three inputs instead of two. In contrast, the OPN model (Lee et al. 2017), depicted in Figure 3(b), takes all the n fragments at the same time, aggregates the features from all possible feature pairs for feature learning, and seeks to pick the best global ordering out of the $n!$ possible combinations via a multiclass classification problem.

Implementation Details

As done in many previous works (Dieleman and Schrauwen 2014), we compute the spectrograms by sampling the songs at 22,050 Hz and using a Hamming window of 2,048 samples and hop size 512 samples. We then transform the spectrograms into 128-bin log mel-scaled spectrograms and use that as input to the networks, after z-score normalization.

Unless otherwise specified, in our implementation all the Siamese ConvNets use 1D convolutional filters (along the time dimension), with the number of filters being 128, 256, 512, respectively, and the filter length being 4. For SEN and CCSN, the convolutional filters for the subsequent ConvNet are 64, 128, 256, respectively, followed by 3 by 3 maximum pooling and the filter size is also 3 by 3. Here, SEN uses 2D convolutions, while CCSN uses 1D convolutions. Except for TSN and OPN, we use a global pooling layer (which is written as ‘CONCAT’ in Figure 3) after the ConvNet in SEN, CCSN, CIN, and the Siamese ConvNet in SN. The dimension of the two fully-connected layers after this pooling layer are all set to 1,024. All networks use rectified linear unit (ReLU) as the activation function everywhere. Lastly, all the models are trained using stochastic gradient descent with momentum 0.9, with batch size setting to 16.

Experiments

Data sets

Any music collection can be used in our puzzle games, since we do not need any human annotations. In this paper, we use an in-house collection of 31,377 clips of Pop music as our corpus. All these clips are audio previews downloaded from

the Internet, with unknown starting point in each song the audio preview was extracted from. All these clips are longer than 24 seconds, so we consider only the first 24 seconds per clip for simplicity of the model training process. Moreover, we randomly pick 6,000 songs as validation set, 6,000 songs for testing, and the remaining 19,377 clips for training.

Different data sets are used as the test set for different music puzzle games. For **music jigsaw puzzle**, we simply use the test set of the in-house collection. For **music sequencing**, we use the popular music subset of the RWC database (Goto et al. 2002), which contains 100 complete songs with manually labeled section boundaries (Goto 2006).³ We divide songs into fragments according to these boundaries. The number of resulting fragments per song ranges from 11 to 34. For simplicity, we consider only the first ten fragments (sections) per song. For **music medley**, we collect 16 professionally-compiled medleys of pop music (by human experts) from YouTube. Each medley contains 7 to 11 different short music clips, whose length vary from 5 to 30 seconds. For reproducibility, we have posted the YouTube links of these medleys in our project website.

We need to perform downbeat tracking for the in-house collection. To this end, we use the implementation of a state-of-the-art recurrent neural network available in the Python library madmom⁴ (Böck et al. 2016). After getting the downbeat positions, we randomly choose some of them so that each fragment is about $\lfloor 24/n \rfloor$ seconds in length.

Result on 3-Piece Fixed-length Jigsaw Puzzle

As the first experiment, we consider the $n = 3$ jigsaw puzzle game, using 1,000 clips randomly selected from the test set of the in-house collection. Accordingly, we train all the neural networks (including the baselines) by playing $n = 3$ jigsaw puzzles using the training set of the same data set. All the clips are segmented at equally-spaced timepoints. Therefore, the length of fragments is fixed to 8 seconds. This corresponds to the simplest setting in Table 1.

We employ the **pairwise accuracy (PA)** and **global accuracy (GA)** as our performance metrics. For an ordering of n fragments, GA requires it to be exactly the same as the groundtruth one, whereas PA takes the average of the correctness of the $n - 1$ composing pairs. For example, ordering (R_1, R_2, R_3) as (R_2, R_3, R_1) would get 0.5 PA (for the

³<https://staff.aist.go.jp/m.goto/RWC-MDB/AIST-Annotation/>

⁴<https://github.com/CPJKU/madmom>

Table 2: The pairwise accuracy and global accuracy on $n = 3$ fixed-length jigsaw puzzle.

Method	pairwise accuracy	global accuracy
SN	0.851	0.825
CCSN (Wang et al. 2016)	0.872	0.840
CIN	0.912	0.864
TSN (Misra et al. 2016)	0.911	0.890
OPN (Lee et al. 2017)	0.929	0.916
SEN (proposed)	0.996	0.994

Table 3: The accuracy on music jigsaw puzzles with different segmentation method (fixed-length or downbeat informed) and different number of fragments. Pairwise accuracy is outside the brackets and global accuracy is inside.

	n	SN	CIN	SEN
fixed	3	0.851 (0.825)	0.912 (0.864)	0.996 (0.994)
	4	0.752 (0.641)	0.844 (0.722)	0.990 (0.982)
	6	0.609 (0.304)	0.761 (0.455)	0.989 (0.977)
	8	0.514 (0.110)	0.682 (0.229)	0.985 (0.953)
downbeat	3	0.781 (0.692)	0.912 (0.863)	0.995 (0.991)
	4	0.681 (0.472)	0.871 (0.761)	0.995 (0.987)
	6	0.536 (0.171)	0.764 (0.499)	0.991 (0.971)
	8	0.449 (0.056)	0.680 (0.297)	0.990 (0.961)

pair (R_2, R_3) is correct) and 0 GA.

The result is shown in Table 2. The performance of the baseline models seem to correlate well with their sophistication, with SN performing the worst (0.825 GA) and OPN (Lee et al. 2017) performing the best (0.916 GA). The comparison between SN and TSN (Misra et al. 2016) implies that more inputs offers more information. Moreover, the results in GA and PA seem to be highly correlated.

More importantly, by comparing the result of SEN against the baseline models, we see that SEN outperforms the others by a great margin, reaching almost 100% accuracy for both metrics. In particular, the performance gap between CCSN and SEN suggests that learning from the similarity matrix seems to be an important design.

Result on Variable-length Jigsaw Puzzle

Next, we consider music jigsaw puzzles with variable-length fragments. Thanks to the global pooling layer, in our implementation SEN, SN and CIN can take input of arbitrary length, even if the length of the training fragments is different from the length of the testing fragments. Moreover, the pair-wise strategy allows these three models to tackle $n > 4$ puzzle games, while some other models such as OPN cannot. Therefore, we only consider SEN, SN and CIN in the following experiments.

We create different games by varying n from 3, 4, 6 to 8, using the uniformly segmented fragments from the in-house data set. The length of the fragments is hence 8, 6, 4, 3 seconds, respectively. Among them, the $n = 8$ game is the most challenging one, partly due to the number of fragments and partly due to their shorter duration (implying less information per fragment). We use the same SEN, SN and CIN models trained by solving $n = 3$ jigsaw puzzles. In addition, we aim at comparing the result of using different segmentation methods to process both the training and test clips. Therefore, we re-train the SEN, SN and CIN models trained by solving $n = 3$ jigsaw puzzles segmented at downbeat positions, and apply them to also downbeat-informed jigsaw puzzles with different values of n .

Table 3 shows the result. We can see that the result of SN and CIN both decrease quite remarkably as the value of n increases, and that the downbeat-informed games are indeed slightly more challenging than the fixed-length games, possibly due to the clarity at the boundary. When $n = 8$, the PA

Table 4: The accuracy of SEN on three kinds of puzzle game for two segmentation methods.

game	fixed-length	downbeat-informed
puzzle ($n = 8$)	0.985 (0.953)	0.990 (0.961)
sequencing	0.789 (0.440)	0.937 (0.790)
medley	0.945 (0.688)	0.961 (0.750)

and GA of SN drop to 0.514 and 0.110, whereas the PA and GA of CIN drop to 0.682 and 0.229. However, the accuracy of the SEN model remains high even when $n = 8$ (0.985 PA and 0.953 GA), suggesting that SEN can work quite robustly against various music jigsaw puzzles.

Result on Music Sequencing and Music Medley

Lastly, we evaluate the performance of SEN on music sequencing and music medley, which are supposed to be more challenging than jigsaw puzzles. We do not consider SN and CIN here, for their demonstrated poor performance in $n = 8$ jigsaw puzzles. Instead, we compare two SEN models, one trained with uniform segmentation ($n = 3$) and the other with downbeat-informed segmentation ($n = 3$).

From Table 4, we can see that these two games are indeed more challenging than jigsaw puzzles. When using a SEN model trained with uniform segmentation, the GA can drop to as low as 0.440 for music sequencing and 0.688 for music medley. However, more robust result can be obtained by training SEN using downbeat-informed segmentation: the GA would be improved to 0.790 and 0.750 for the two games, respectively. This is possibly because the downbeat-informed segmentation can avoid SEN from learning only low-level features at the boundary of fragments.

We perform an error analysis looking into the incorrect prediction in the music sequencing game, which has some musically meaningful insights. A song is composed of several sections, such as intro (I), verse (V), chorus (C) and bridge (B), with some variations such as V_a and V_b . A correct global ordering of one of the songs in RWC is: I- V_a - B_a - V_b -C-pre-Ca- B_b - V_a -Vc-C-pre. For this song, the estimated ordering of SEN is: I- B_b - V_a -Vc-C-pre-Ca- V_a - B_a - V_b -C-pre. We can use a numerical notation and represent our result as 1-7-8-9-10-6-2-3-4-5. We can see that the local prediction of 7-8-9-10 and 2-3-4-5 is in correct order. Moreover, these two passages are fairly similar (both have the structure V-C-pre). Therefore, the predicted ordering may sound right as well. Indeed, we found that most of the incorrect predictions are correct in local ordering.

We use user-created medleys as the groundtruth in the medley game, but as the creation of a medley is an art, different orderings may sound right. Therefore, we encourage readers to visit our project website to listen to the result.

Ablation Analysis

We assess the effect of various design of downbeat-informed SEN by evaluating ablated versions. Table 5 shows the result when we (from left to right): i) replace cosine similarity in Eq. (2) by inner product, ii) increase the stride of the convolutions in Siamese ConvNet from 1 to 2, iii) use only global

Table 5: The result of a few ablated version of SEN for different music puzzle games.

game	Inner product	Conv stride 2	Global P (mean)	Global P (max)	R2R1 only	R1R3 only	R3R1 only	All
puzzle	0.90 (0.69)	0.65 (0.17)	0.96 (0.87)	0.98 (0.93)	0.84 (0.57)	0.97 (0.87)	0.96 (0.86)	0.99 (0.96)
sequencing	0.74 (0.38)	0.54 (0.06)	0.81 (0.49)	0.92 (0.76)	0.62 (0.22)	0.81 (0.46)	0.91 (0.69)	0.94 (0.79)
medley	0.88 (0.50)	0.73 (0.13)	0.81 (0.56)	0.93 (0.69)	0.86 (0.44)	0.93 (0.69)	0.90 (0.63)	0.96 (0.75)




Figure 4: Embeddings of different data pairs learned by (from left to right) SEN, CCSN and SN, respectively. The embeddings are projected to a 2D space for visualization via t-SNE (Maaten and Hinton 2008). The figure is best viewed in color.

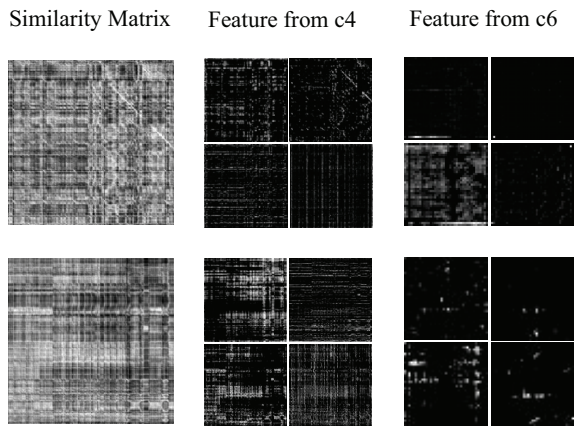


Figure 5: Visualizations of the features learned from the c4 and c6 layers of SEN for two pairs of fragments.

mean pooling or global max pooling (we use the concatenation of mean, max and standard deviation in our full model), and iv) use one type of negative data only. Most of these changes decrease the accuracy of SEN. Some observations:

- Calculating the similarity matrix using the inner product cannot guarantee that the similarity scores are in the range of $[0, 1]$ and this hurts the accuracy of SEN.
- Setting the stride size larger can speed up the training process, but doing so losses much temporal information.
- Max pooling alone works quite well for the global pooling layer, but it is even better to also consider mean and standard deviation.
- Using R2R1 as the negative data alone is far from sufficient. Actually, both R1R3 only and R3R1 only seem to

work better than R2R1 only. The best result (especially in GA) is obtained by using all three types of negative pairs.

What the SEN Model Learns?

Figure 4 shows the embeddings (output of the last fully-connected layer) of different data pairs learned by SEN, CCSN and SN, respectively. We can clearly see that the positive and negative pairs can be fairly easily distinguished by the embeddings learned by SEN. Moreover, SEN can even distinguish R2R1 (consecutive) from R1R3 and R3R1 (non-consecutive). This is an evidence of the effectiveness of SEN in learning sequential structural patterns.

Finally, Figure 5 shows the features from the first (c4) and last convolution (c6) layers in the ConvNet of SEN, given two randomly chosen pairs (the first row is R1R2 and the second row is R1R3). We see that the filters detect different patterns and textures from the similarity matrices.

Conclusion

In this paper, we have presented a novel Siamese network called the similarity embedding network (SEN) for learning sequential patterns in a self-supervised way from similarity matrices. We have also demonstrated the superiority of SEN over existing Siamese networks using different types of music puzzle games, including music medley generation.

In our evaluation, however, music medley generation is viewed as just one of the evaluation tasks. In future work, we will focus more on the music medley generation task itself. For example, a listening test should be conducted for subjective evaluation. We also plan to deeply investigate the features or patterns our model learns for generating medleys and correlate our findings with those reported in related work on automatic music mashup (Davies et al. 2014) and playlist sequencing (Bittner et al. 2017). We also want to

investigate thumbnailing methods (Huang, Chou, and Yang 2017b; Kim et al. 2017) to pick fragments from different songs, and methods such as beat-match and cross-fade (Bittner et al. 2017) to improve the transition between clips.

References

- Bittner, R. M.; Gu, M.; Hernandez, G.; Humphrey, E. J.; Jehan, T.; McCurry, P. H.; and Montecchio, N. 2017. Automatic playlist sequencing and transitions. In *Proc. Int. Soc. Music Information Retrieval Conf.*
- Böck, S.; Korzeniowski, F.; Schlüter, J.; Krebs, F.; and Widmer, G. 2016. madmom: a new Python audio and music signal processing library. In *Proc. ACM MM*, 1174–1178.
- Böck, S.; Krebs, F.; and Widmer, G. 2016. Joint beat and downbeat tracking with recurrent neural networks. In *Proc. Int. Soc. Music Information Retrieval Conf.*, 255–261.
- Bromley, J.; Guyon, I.; Lecun, Y.; Sckinger, E.; and Shah, R. 1994. Signature verification using a “siamese” time delay neural network. In *Proc. Annual Conf. Neural Information Processing Systems*.
- Chopra, S.; Hadsell, R.; and LeCun, Y. 2005. Learning a similarity metric discriminatively, with application to face verification. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 539–546.
- Davies, M. E.; Hamel, P.; Yoshii, K.; and Goto, M. 2014. AutoMashUpper: Automatic creation of multi-song music mashups. *IEEE/ACM Trans. Audio, Speech and Language Processing* 22(12):1726–1737.
- Dieleman, S., and Schrauwen, B. 2014. End-to-end learning for music audio. In *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, 6964–6968.
- Fernando, B.; Bilen, H.; Gavves, E.; and Gould, S. 2016. Self-supervised video representation learning with odd-one-out networks. *arXiv preprint arXiv:1611.06646*.
- Goto, M.; Hashiguchi, H.; Nishimura, T.; and Oka, R. 2002. RWC music database: Popular, classical and jazz music databases. In *Proc. Int. Soc. Music Information Retrieval Conf.*, 287–288.
- Goto, M. 2006. AIST annotation for the RWC music database. In *Proc. Int. Soc. Music Information Retrieval Conf.*, 359–360.
- Hansen, K. F.; Hiraga, R.; Li, Z.; and Wang, H. 2013. Music puzzle: An audio-based computer game that inspires to train listening abilities. In *Proc. Advances in Computer Entertainment*. Springer. 540–543.
- Huang, Y.-S.; Chou, S.-Y.; and Yang, Y.-H. 2017a. DJnet: A dream for making an automatic DJ. In *Proc. ISMIR, late-breaking demo paper*.
- Huang, Y.-S.; Chou, S.-Y.; and Yang, Y.-H. 2017b. Music thumbnailing via neural attention modeling of music emotion. In *Proc. APSIPA*.
- Hudson, N. J. 2011. Musical beauty and information compression: Complex to the ear but simple to the mind? *BMC research notes* 4(1):9.
- Kim, A.; Park, S.; Park, J.; Ha, J.-W.; Kwon, T.; and Nam, J. 2017. Automatic DJ mix generation using highlight detection. In *Proc. ISMIR, late-breaking demo paper*.
- Lee, H.-Y.; Huang, J.-B.; Singh, M.; and Yang, M.-H. 2017. Unsupervised representation learning by sorting sequences. *arXiv preprint arXiv:1708.01246*.
- Lin, Y.-T.; Liu, I.-T.; Jang, J.-S. R.; and Wu, J.-L. 2015. Audio musical dice game: A user-preference-aware medley generating system. *ACM Trans. Multimedia Comput. Commun. Appl.* 11(4):52:1–52:24.
- Liu, J.-Y., and Yang, Y.-H. 2016. Event localization in music auto-tagging. In *Proc. ACM MM*, 1048–1057.
- Lotter, W.; Kreiman, G.; and Cox, D. 2017. Deep predictive coding networks for video prediction and unsupervised learning. In *Proc. Int. Conf. Learning Representations*.
- Luo, W.; Schwing, A. G.; and Urtasun, R. 2016. Efficient deep learning for stereo matching. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 5695–5703.
- Maaten, L. v. d., and Hinton, G. 2008. Visualizing data using t-SNE. *J. Machine Learning Research* 9:2579–2605.
- Marwan, N.; Romano, M. C.; Thiel, M.; and Kurths, J. 2007. Recurrence plots for the analysis of complex systems. *Physics reports* 438(5):237–329.
- Misra, I.; Zitnick, C. L.; Hebert, M.; and . 2016. Shuffle and learn: Unsupervised learning using temporal order verification. In *Proc. European Conf. Computer Vision*, 527–544.
- Mueller, J., and Thyagarajan, A. 2016. Siamese recurrent architectures for learning sentence similarity. In *Proc. AAAI*, 2786–2792.
- Nieto, O., and Bello, J. P. 2016. Systematic exploration of computational music structure research. In *Proc. Int. Soc. Music Information Retrieval Conf.*, 547–553.
- Noroozi, M., and Favaro, P. 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. In *Proc. European Conf. Computer Vision*, 69–84. Springer.
- Paulus, J.; Miller, M.; and Klapuri, A. 2010. State of the art report: Audio-based music structure analysis. In *Proc. Int. Soc. Music Information Retrieval Conf.*, 625–636.
- Serrà, J.; Müller, M.; Grosche, P.; and Arcos, J. L. I. 2012. Unsupervised detection of music boundaries by time series structure features. In *Proc. AAAI*, 1613–1619.
- Smith, J. B.; Kato, J.; Fukayama, S.; Percival, G.; and Goto, M. 2017. The CrossSong Puzzle: Developing a logic puzzle for musical thinking. *J. New Music Research* 1–16.
- Upham, F., and Farbood, M. 2013. Coordination in musical tension and liking ratings of scrambled music. In *Proc. Meeting of the Society for Music Perception and Cognition*.
- Wang, F.; Zuo, W.; Lin, L.; Zhang, D.; and Zhang, L. 2016. Joint learning of single-image and cross-image representations for person re-identification. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 1288–1296.
- Widmer, G. 2016. Getting closer to the essence of music: The Con Espressione manifesto. *ACM Trans. Intell. Syst. Technol.* 8(2):19:1–19:13.