

AUDIO TO SCORE MATCHING BY COMBINING PHONETIC AND DURATION INFORMATION

Rong Gong, Jordi Pons and Xavier Serra

Music Technology Group, Universitat Pompeu Fabra, Barcelona

rong.gong@upf.edu, jordi.pons@upf.edu, xavier.serra@upf.edu

ABSTRACT

We approach the singing phrase audio to score matching problem by using phonetic and duration information – with a focus on studying the jingju a cappella singing case. We argue that, due to the existence of a basic melodic contour for each mode in jingju music, only using melodic information (such as pitch contour) will result in an ambiguous matching. This leads us to propose a matching approach based on the use of phonetic and duration information. Phonetic information is extracted with an acoustic model shaped with our data, and duration information is considered with the Hidden Markov Models (HMMs) variants we investigate. We build a model for each lyric path in our scores and we achieve the matching by ranking the posterior probabilities of the decoded most likely state sequences. Three acoustic models are investigated: (i) convolutional neural networks (CNNs), (ii) deep neural networks (DNNs) and (iii) Gaussian mixture models (GMMs). Also, two duration models are compared: (i) hidden semi-Markov model (HSMM) and (ii) post-processor duration model. Results show that CNNs perform better in our (small) audio dataset and also that HSMM outperforms the post-processor duration model.

1. INTRODUCTION

The ultimate goal of our research project is to automatically evaluate the jingju a cappella singing of a student in the scenario of jingju singing education – see Figure 1. Jingju, a traditional Chinese performing art also known as Peking or Beijing opera, is extremely demanding in the clear pronunciation and accurate intonation of each syllabic or phonetic singing unit. To this end, during the initial learning stages, students are required to completely imitate tutor’s singing. Therefore, the automatic jingju singing evaluation system we envision is based on this training principle and measures the intonation and pronunciation similarities between the student’s and the tutor’s singings. Before measuring similarities, the singing phrase should be automatically segmented into syllabic or phonetic units in




Figure 1. System design framework of the entire research project. The modules with bold border lines are addressed in this paper.

order to capture the temporal details. Jingju music scores, which contain the phonetic and duration information for each singing syllable, will be beneficial for this segmentation. In the application scenario, the score of a query audio could be selected from the database by the user itself. However, to avoid manual intervention and improve the user experience, we tackle the problem of automatically finding the corresponding music score for a given query audio (bold in Figure 1). Note that achieving successful methods for audio to score matching might be beneficial for several music informatics research (MIR) tasks, such as score-informed automatic syllable/phoneme segmentation [5] or score-informed source separation [14]. The objective of this research task is to find the corresponding score for a given singing audio query. We restrict this research to the “matching” scope by pre-segmenting both the singing audios and the music scores into the phrase units.

Xipi and *erhuang* are the main modes in jingju music. Each has two basic melodic contours – an opening phrase and a closing phrase. Each basic melodic contour is constructed upon characteristic pitch progressions for each mode [22]. Therefore, singing phrases from different arias sharing the same mode are likely to have a similar melodic contour. Figure 2 shows an example of this fact. However, melodic information tends to be intuitively used for such matching tasks. For example in Query-by-Singing/Humming (QBSH) [15], melodic similarities can be obtained by comparing the distance between the F0 contour of the query audio and those synthesized from the candidate scores. Then, the best matched music score can be retrieved by selecting the most similar melody. But note that using this approach for jingju music would bring matching errors since the melodic contours of the same mode are similar in this sense. In this case, it is more ap-



© Rong Gong, Jordi Pons and Xavier Serra. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** Rong Gong, Jordi Pons and Xavier Serra. “Audio to score matching by combining phonetic and duration information”, 18th International Society for Music Information Retrieval Conference, Suzhou, China, 2017.




Figure 2. An example of different phrases having a similar *Xipi* melodic contour in our score dataset. The lengths of these contours are normalized to a fixed sample size.

appropriate to use another notion of similarity. We propose using the lyrics since the stories narrated in different jingju arias are distinctive and lyrics tend to change through different jingju arias – even when they share the same mode. Therefore, phonetic information might be useful to identify a similar score given a query audio.

QBSH is the most related research task to our study, which retrieves a song by singing a portion of itself. Most of the studies use melody information as the only cue. The typical process of such systems was introduced by Molina *et al.* [15]: firstly, the F0 contour and/or a note-level transcription for a given singing query are extracted; and then, a set of candidate songs are retrieved from a large database using a melodic matcher module. The most successful QBSH system, which obtained the best results in MIREX 2016 contest, is based on the method of multiple similarity measurements fusion [21]. This system proposed a melodic matcher which combines several similarities that are note-based and frame-based. The authors claim that the fusion mechanism improves the query performance because no similarity measurement is perfect. Therefore, information sources that are complementary to each other might be beneficial for this approach. Very few studies have explored the capability of the phonetic information for QBSH. Guo *et al.* [8] and Wang *et al.* [20] both used a lyric recognizer based on Hidden Markov Models (HMMs). Their recognition networks¹ were constructed with the phonetic information from the query candidates database. They used frame-based MFCCs to create the acoustic models with GMMs. Then, the Viterbi algorithm was executed over the recognition networks to either obtain the most likely phonetic state sequence (for Wang *et al.* [20]) or the posterior probability of each possible decoding path (for Guo *et al.* [8]). The final score of a query candidate is either based on semantic similarity [20] or based on the posterior probability of its corresponding lyrics [8].

Another research task related to our study is singing keyword spotting. The main goal of this task is to search for one or more keywords in a singing query. The system proposed by Kruspe [12] searches for a specific singing keyword on the resulting phoneme observations.

¹ The topology of the HMM is defined by the recognition network.

A keyword-filler HMMs is employed for this purpose. She used two phoneme duration models: the HSMM and the post-processor duration model.

Finally, both phonetic and duration information extracted from the score have been extensively used in alignment-related tasks, such as audio-to-score alignment and audio-to-lyrics alignment. For example, Gong *et al.* [4] construct a left-to-right HSMM using phonetic and duration information. Or Dzhambazov *et al.* [3] use a similar approach for aligning polyphonic audio. Analogously, the proposed approach explores the use of both phonetic and duration information (available in scores) to tackle the matching ambiguity problem existing in jingju music.

The remainder of this paper is organized as follows: the used dataset is introduced in section 2, section 3 explains the modules of the proposed approach – detailing how to incorporate phonetic and duration information. Experiments and results are reported in section 4, and section 5 concludes and points out future work.

2. DATASET

The jingju a cappella singing dataset is composed of two overlapping parts: (i) audio and (ii) score datasets.

The audio dataset [1] used for this study consists of two role-types singing: *dan* (young woman) and *laosheng* (old man). The *dan* part of this dataset has 42 recordings sung by 7 singers and the *laosheng* part has 23 recordings sung by 7 *laosheng* singers. The boundary annotations of the audio dataset have been done in Praat format (textgrid) considering a hierarchy of three levels: phrase, syllable and phoneme – using Chinese characters, pinyin notations and X-SAMPA notations, respectively. 32 phoneme classes are used in the phoneme-level annotation. Two Mandarin native speakers and a jingju musicologist have been devoted to this annotating work. Annotations and more detailed information can be found online². Some statistics about the dataset are reported in Table 2. The average phrase, syllable and voiced phoneme length of *dan* singing are ostensibly greater than those of *laosheng* singing (bold numbers in Table 2), which might indicate that *dan* singing tends to have more pitch variation and ornamentation – as we could observe empirically by listening to the data.

	Num.	Avg. len (s)	Std. len (s)
Phrases	325	16.42	14.11
Syllables	2933	1.58	2.82
Voiced phonemes	7198	0.61	0.97
Unvoiced phonemes	2014	0.10	0.67
Phrases	247	9.47	8.14
Syllables	2289	0.88	1.48
Voiced phonemes	4948	0.39	0.78
Unvoiced phonemes	1454	0.07	0.05

Table 1. Detailed information of the jingju a cappella singing audio dataset: *dan* (top), *laosheng* (bottom).

² <http://doi.org/10.5281/zenodo.344932>

The audio dataset, along with their boundary annotations, is split into three parts: training set, development (dev) set and test set. We define the training set to be the non-overlapping part with the score dataset, see Figure 3. The training set will be used for calculating the phonetic duration (duration information) and training acoustic models (phonetic information). After taking the training set out, we define the development set to be the half of the remaining phrases in the audio dataset (randomly selected) – it will be used for parameters optimization. The test set consists on the remaining phrases of the audio dataset – it will be used for testing the acoustic models performance and the matching performance.

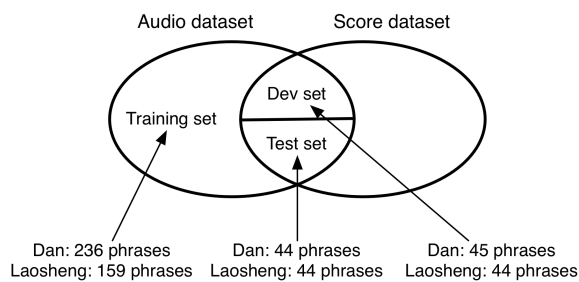


Figure 3. The intersection between the audio and the score datasets. The partition of the audio dataset.

On the other hand, the score dataset contains 435 *dan* phrases and 481 *laosheng* phrases. The scores have been typed in stave notation (including lyrics in Chinese characters) using MuseScore from different printed sources in jianpu notation. Since tempo is usually not clearly noted in the printed score, we do not include this information in the dataset. The relative syllabic durations are indicated by the note durations corresponding to the lyrics, which will be used to calculate the phonetic duration (duration information) and the matching network. The whole score dataset will be used as candidates for testing the matching performance and for parameter optimization.

3. APPROACH

The proposed approach aims to match the query audio to its score by using phonetic and duration information. During the training process (red boxes in Figure 4): the acoustic models of each phoneme are shaped by using the audio training set and its phonetic boundary annotations; the score dataset is used to construct a matching network; and phoneme duration distributions are estimated by using both audio training set and scores. During the matching process (green boxes in Figure 4): two duration models –HSMM and post-processor– are explored for the Viterbi decoding step. Finally, the best-matched phrase is found by ranking the decoded state sequence probabilities.

3.1 Acoustic models

Here presented acoustic models aim to represent the relationship between an audio signal and the 32 phoneme




Figure 4. Diagram of the proposed approach.

classes present in our dataset. The output of these models yield probability scores for each phoneme class.

The most popular way to approach acoustic modeling is by using GMMs and MFCCs features [8, 20]. For that reason, we set as baseline a 40-component GMM with the following input vector: 13 MFCCs, their deltas and delta-deltas. Moreover, DNNs have been found very useful for acoustic modeling [9, 13]. Therefore, we propose an additional baseline: a DNN with 2 hidden layers followed by the 32-way softmax output layer – the input is set to be a log-mel spectrogram.

However, DNNs are very prone to over-fitting and the available dataset is relatively small. For that reason we propose using CNNs since these are more robust against over-fitting – note that CNNs allow parameter sharing. Additionally, Pons *et al.* [17] have successfully used spectrograms-based CNNs for learning music timbre representations from small datasets. Given that timbre is an important feature for acoustic modeling, we propose using the same architecture: a single convolutional layer with filters of various sizes [16, 17]. The input is set to be a log-mel spectrogram. We use 128 filters of sizes 50×1 and 70×1 , 64 filters of sizes 50×5 and 70×5 , and 32 filters of sizes 50×10 and 70×10 – where the first and second numbers denote the frequential and temporal size of the filter, respectively. A max-pool layer of $2 \times N'$ is followed by a 32-way softmax output layer with 30% dropout – where N' denotes the *temporal* dimension of the resulting feature map. $2 \times N'$ max-pool layer was chosen to achieve time-invariant representations while keeping the frequency resolution. And *same* padding is used to preserve the dimensions of the feature maps so that these are concatenable. Filter shapes are designed so that filters can capture the relevant time-frequency contexts for learning timbre representations – according to the design strategy proposed by Pons *et al.* [17]

Log-mel spectrograms are of size 80×21 – the network takes a decision for a frame given its context: ± 10 frames, 21 frames in total. Activation functions are ELUs [2] for all deep learning models and these are optimized with

stochastic gradient descent (batch size: 64), using ADAM [11] and early stopping – when validation loss (categorical cross-entropy) does not decrease after 10 epochs.

Spectrograms are computed from audio recordings sampled at 44.1 kHz. STFT is performed using a window length of 25ms (2048 samples with zero-padding) with a hop size of 10ms. The 80 log-mel bands energies are calculated on frequencies between 0Hz and 11000Hz and these are standardized to have zero mean and unit variance.

The acoustic models are trained separately for each role-type and their performance is reported in section 4.2.

3.2 Matching network

The matching network defines the topology of the hidden Markov model. By using each candidate phrase in the score dataset as an isolated unit, isolated-phrase matching networks can be constructed. Figure 5 shows the structure of this matching network, which has $K = 916$ lyric paths.




Figure 5. The structure of the K paths isolated-phrase matching network. Path 3 shows an example of the left-to-right state chain structure of an individual lyric path.

The matching network uses HMMs or HSMMs, depending on how the internal duration is modeled. Each path is a left-to-right state chain which represents the phoneme transcription of its lyrics. In order to construct the lyric path, pinyin lyrics are segmented into phonetic units and transcribed into X-SAMPA notations by using a predefined dictionary. For example, a path which has the lyrics *yan jian de hong ri* in pinyin is a chain consisting of 12 states: j, En, c, j, En, c, ʔ, x, UN, N, r\', 1 in X-SAMPA notation. When the decoding process has finished, each lyric path can get a posterior probability which will be used as the similarity measure between the query phrase and the candidate phrase.

3.3 Phonetic duration distributions

Phonetic duration information comes from two sources: the boundary annotations of audio training dataset and the score dataset. The phonetic duration is not directly indicated in the score. However, it is indispensable for modeling the phonetic duration distribution for each state in the matching network. The syllable, of which duration can be deduced by the corresponding note(s), is used to restrict the durations of the phonemes.

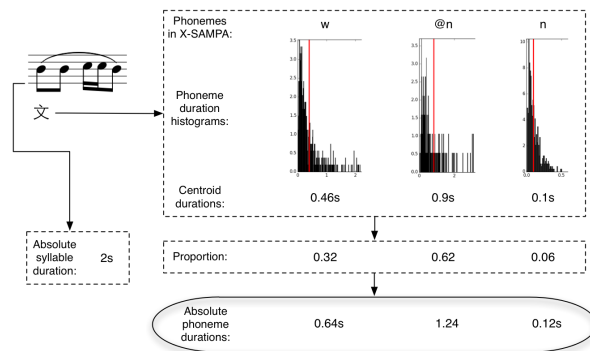


Figure 6. Flowchart example of estimating the phonetic durations of a syllable.

In the following, we propose a method for estimating the absolute phonetic duration given: (i) the score, and (ii) the phonemes duration histograms computed from the audio dataset annotations. First, we omit silence parts in the query audio (with a simple voice activity detection method [19]) and also in the score by removing the rest notes. Second, we compute the duration histogram and its duration centroid for each phoneme class – by aggregating the phonetic durations indicated in the boundary annotations of the audio training dataset. Then, we segment each syllabic duration in the score dataset into phonetic durations according to the proportion of their duration centroids. Finally, as the scores do not contain tempo, we normalize the phonetic durations of each phrase such that their summation is equal to the duration of the query audio. See Figure 6 for an equivalent graphic explanation. In Figure 6, the centroid durations of these three phonemes are: 0.46s, 0.9s and 0.1s, summing: 1.46s – alternatively, these can be expressed as a proportion of 1.46s: 0.32, 0.62 and 0.06. With these proportions and the absolute syllable duration (2s), we can compute the absolute phoneme durations: $0.32 \cdot 2s = 0.64s$, $0.62 \cdot 2s = 1.24s$ and $0.06 \cdot 2s = 0.12s$.

The phonetic duration distribution needs to be calculated for each state in the matching network in order to incorporate the *a priori* phonetic duration information. We model it by Gaussian distributions:

$$\mathcal{N}(x; \mu_l, \sigma_l^2) = \frac{1}{\sqrt{2\pi}\sigma_l} \exp\left(-\frac{(x - \mu_l)^2}{2\sigma_l^2}\right). \quad (1)$$

where μ_l is the duration of the phoneme l deduced by the above method and the standard deviation σ_l is proportional to μ_l : $\sigma_l = \gamma\mu_l$. The proportionality constant γ will be optimized in section 4.3 for each role-type.

3.4 Duration modeling

Standard Markovian state does not impose explicitly duration distribution, instead, imposing an implicit state occupancy distribution which corresponds to a “1-shifted” geometric distribution [6]:

$$d_j(u) = (1 - \tilde{p}_{jj})\tilde{p}_{jj}^{u-1} \quad (2)$$

where u denotes the occupancy or sojourn time in a Markovian state j and \tilde{p}_{jj} denotes the self-transition probability of the state j . Because of the implicitness of the Markovian state occupancy, the phonetic duration distribution introduced in section 3.3 can not be imposed. Kruspe [12] presents two duration modeling techniques for HMMs: Hidden semi-markov model (HSMM) and post-processor duration model.

3.4.1 Hidden semi-markov model

Guédon [6] defined a semi-Markov chain S_t with finite state space $0, \dots, J - 1$ by the following parameters:

- initial probabilities $\pi_j = P(S_0 = j)$ with $\sum_j \pi_j = 1$
- transition probability of semi-Markovian state j : for each $k \neq j, p_{jk} = P(S_{t+1} = k | S_{t+1} \neq j, S_t = j)$ with $\sum_{k \neq j} p_{jk} = 1$ and $p_{jj} = 0$

An explicit occupancy distribution is attached to each semi-Markovian state:

$$d_j(u) = P(S_{t+u+1} \neq j, S_{t+u-v} = j, v = 0, \dots, u - 2 | S_{t+1} = j, S_t \neq j), u = 1, \dots, M_j \quad (3)$$

where M_j denotes the upper bound to the time spent in state j . $d_j(u)$ defines the conditional probability of leaving state j at time $t + u + 1$ and entering state j at time $t + 1$.

To apply HSMMs to the matching network, we first use the matching network as the HSMMs topology. Thus the state occupancy distribution is set to its corresponding phonetic duration distribution. Then the probabilities of each left-to-right state transition are set to 1 because all self-transition probabilities in HSMMs are 0. The goal is to find the most likely sequence of hidden states for each lyric path and collect its posterior probability. The Viterbi algorithm meets this specific goal and its complete implementation is provided in [7].

3.4.2 Post-processor duration model

The post-processor duration model was first introduced by Juang *et al.* [10]. It was then experimentally proved in Kruspe’s paper [12] that this duration model works better than HSMMs for the keyword spotting task in English pop singing voice. The post-processor duration model uses the original HMMs Viterbi algorithm – therefore, during the decoding process, no explicit occupancy duration distribution is imposed.

The log posterior probability of the decoded most likely state sequence is augmented by the log duration probabilities:

$$\log \hat{f} = \log f + \alpha \sum_{j=1}^N \mathcal{N}(u_j; \mu_j, \sigma_j^2) \quad (4)$$

where f is the HMMs posterior probability, α is a weighting factor which will be optimized in section 4.3, $j = 1, \dots, N$ is the decoded state number in the most likely state sequence, and $\mathcal{N}(u_j; \mu_j, \sigma_j^2)$ is the occupancy probability of being in state j for the occupancy u_j .

4. EXPERIMENTS AND RESULTS

4.1 Performance metrics

Two experiments³ are performed: the first is to evaluate the performance of the acoustic models, and the second is to evaluate the proposed matching approaches. For the first task, we use one simple evaluation metric: the overall classification accuracy which is defined as the fraction of instances that are correctly classified. For the second task, our goal is to evaluate the ability to match the ground-truth phrase in the score dataset to the query one, which is almost identical to the goal of a QBSH system: “finding the ground-truth song in a song database from a given singing/humming query”. Therefore, we borrow the standard performance metrics used in QBSH task to evaluate our approaches: Top-M hit and Mean Reciprocal Rank (MRR) [8]. The Top-M hit rate is the proportion of queries for which $r_i \leq M$, where r_i denotes the rank of the ground-truth score phrase. MRR is the average of the reciprocal ranks across all queries, n is the number of queries, and $rank_i$ is the posterior probability rank of the ground-truth phrase corresponding to the i -th query.

$$MRR = \frac{1}{n} \sum_{i=1}^n \frac{1}{rank_i} \quad (5)$$

4.2 Acoustic models

CNN, DNN and GMM acoustic models yield probability scores for each phoneme class. In order to evaluate the classification accuracy, we choose the phoneme class with the maximum probability score as the prediction. Table 4.2 reports the performance of CNN, DNN and GMM acoustic models evaluated on the test set.

	<i>dan</i> (#parameters)	<i>laosheng</i> (#parameters)
CNNs	0.484(222k)	0.432(222k)
DNNs	0.284(481k)	0.282(430k)
GMMs	0.290(-)	0.322(-)

Table 2. Overall classification accuracies of CNN and baseline acoustical models for *dan* and *laosheng* datasets.

The relatively low classification accuracies for all three models show that modeling the phonetic characteristics of jingju singing voice is a challenging problem. Our best results are achieved with CNNs – and GMMs perform better than DNNs. Interestingly, these results contrast with the literature where Hinton *et al.* [9] describe that DNNs acoustic models largely outperform GMMs for automatic speech recognition, and Maas *et al.* [13] showed that CNNs

³ Code: <https://goo.gl/1XB6j1>

perform worse than DNNs for building speech acoustic models. First, we argue that in our case DNNs perform worse than GMMs and CNNs because a small amount of training data is available. DNNs require a lot of training data to achieve good performance and note that large amounts of training data are typically not available for most MIR tasks. And second, note the CNNs used here are specifically designed to efficiently learn timbre representations [17] while Maas *et al.* [13] used small square filters, which proved successful in computer vision tasks. These results show that using CNN architectures designed for the task at hand is especially beneficial in small data scenarios. A CNN model is used in the following experiments.

4.3 Parameters optimization

The parameters which need to be optimized for *dan* and *laosheng* role-types are: the weighting factor α for the post-processor duration model, and the proportionality constant γ for both models: HSMMs and post-processor duration model. Table 4.3 reports the optimal values we obtained by doing grid search on the development set – MRR metric was maximized.

Parameters	Search bounds	Optimal values
α	[0.25, 2] with step 0.25	1.0 / 1.0
γ HSMMs	[0.1, 2] with step 0.1	0.1 / 0.1
γ post-processor	[0.1, 2] with step 0.1	0.7 / 1.5

Table 3. Parameters to be optimized, search bounds and resulting optimal values (*dan* / *laosheng*).

4.4 Duration modeling

To highlight the advantage of using duration modeling methods for audio to score matching, a standard HMM without explicitly imposing the occupancy distribution is used as a baseline. Results in Figure 7 show that its performance is inferior to the HSMM duration model.




Figure 7. Phrase matching performance of HSMM and post-processor duration model with CNN acoustic model: *dan* (top), *laosheng* (bottom).

One can also observe in Figure 7 that HSMM performs the best, improving the baseline MRR metric performance by 13.2% for *dan* role-type and 15.1% for *laosheng* role-type. This means that HSMMs explicit duration modeling

can help achieve a better audio to score matching by using phonetic information.

The post-processor duration model does not significantly improve the baseline performance. This result contrasts with the literature, where the post-processor duration model worked better than HSMMs for singing voice keyword spotting [12]. This inconsistency might result from (i) the length difference of the matching unit (singing-words vs. singing-phrases), and (ii) the large standard deviation of the jingju singing phonemes length. First, in Kruspe’s work [12], the matching unit is the singing keyword – which usually contains fewer phonemes than a singing phrase (as in our case). And second, the vowel length standard deviation of the a cappella dataset used by Kruspe [12] (around 0.3s) is much short than in our dataset (*dan*: 0.97s, *laosheng*: 0.78s) – denoting less vowel duration variance than in our study case. Moreover, a significant deficiency of the post-processor duration model is that it does not provide the most likely state sequence by internally considering the durations, but it computes a *new* weighted likelihood given the obtained sequence [12]. If the most likely state sequence is decoded poorly, it can’t be restored by the post-processor duration model.

5. CONCLUSIONS AND FUTURE WORK

In this paper, we presented an audio to score matching approach that uses phonetic and duration information.

We explored two duration models: HSMM and post-processor duration model. HSMMs achieved better results than the post-processor duration model – probably due to (i) the matching units length, (ii) the large standard deviation of the considered phonemes, and (iii) because for the post-processor duration model it is hard to recover a decoding mistake. Moreover, HSMMs achieved a better matching performance than the baseline-HMMs approach, which only took into account phonetic information, denoting the utility of using duration information.

We also compared CNN, DNN and GMM acoustic models, and CNNs have shown to be superior in our small singing voice audio dataset. The used CNN architecture was specifically designed to learn timbral representations efficiently [17] – this being the key factor for enabling CNNs (a deep learning method requiring large amounts of data) to perform so well on such a small dataset.

There are many possibilities to improve our approach. It has been shown in the speech research field that LSTM RNNs achieved the best acoustic modeling performance [18]. However, this method requires a large training dataset in order prevent from over-fitting. Another possibility to improve our acoustic model is to go deeper with the current single-layer CNN architecture, but this will also require more training data. We plan to collect more jingju a cappella singing recordings and perform data augmentation to leverage the capability of the acoustic models. Furthermore, in order to take advantage of the melodic information existing in both audio and score datasets, we also plan to investigate methods which can fuse melodic, phonetic and duration information.

6. ACKNOWLEDGEMENTS

We are grateful for the GPUs donated by Nvidia. This work is partially supported by the Maria de Maeztu Programme (MDM-2015-0502) and by the European Research Council under the European Union's Seventh Framework Program, as part of the CompMusic project (ERC grant agreement 267583).

7. REFERENCES

- [1] D. A. A. Black, M. Li, and M. Tian. Automatic Identification of Emotional Cues in Chinese Opera Singing. In *ICMPC*, Seoul, South Korea, August 2014.
- [2] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv:1511.07289*, 2015.
- [3] G. B. Dzhabazov and X. Serra. Modeling of phoneme durations for alignment between polyphonic audio and lyrics. In *SMC*, Maynooth, Ireland, 2015.
- [4] R. Gong, P. Cuvillier, N. Obin, and A. Cont. Real-Time Audio-to-Score Alignment of Singing Voice Based on Melody and Lyric Information. In *Interspeech*, Dresden, Germany, September 2015.
- [5] Rong Gong, Nicolas Obin, Georgi Dzhabazov, and Xavier Serra. Score-informed syllable segmentation for jingju a cappella singing voice with mel-frequency intensity profiles. In *International Workshop on Folk Music Analysis*, Málaga, Spain, June 2017.
- [6] Y. Guédon. Hidden hybrid markov/semi-markov chains. *Computational Statistics & Data Analysis*, 49(3):663 – 688, 2005.
- [7] Yann Guédon. Exploring the state sequence space for hidden markov and semi-markov chains. *Computational Statistics & Data Analysis*, 51(5):2379–2409, 2007.
- [8] Z. Guo, Q. Wang, G. Liu, J. Guo, and Y. Lu. A music retrieval system using melody and lyric. In *2012 IEEE International Conference on Multimedia and Expo Workshops*, pages 343–348, July 2012.
- [9] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, et al. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *IEEE Signal Processing Magazine*, 29(6):82–97, 2012.
- [10] B. Juang, L. Rabiner, S. Levinson, and M. Sondhi. Recent developments in the application of hidden markov models to speaker-independent isolated word recognition. In *ICASSP*, volume 10, pages 9–12, Apr 1985.
- [11] D. Kingma and J. B. Adam. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014.
- [12] A. M. Kruspe. Keyword spotting in singing with duration-modeled hmms. In *EUSIPCO*, pages 1291–1295, August 2015.
- [13] Andrew L Maas, Peng Qi, Ziang Xie, Awni Y Hannun, Christopher T Lengerich, Daniel Jurafsky, and Andrew Y Ng. Building dnn acoustic models for large vocabulary speech recognition. *Computer Speech & Language*, 41:195–213, 2017.
- [14] Marius Miron, Julio J. Carabias-Orti, Juan J. Bosch, Emilia Gómez, and Jordi Janer. Score-informed source separation for multichannel orchestral recordings. *JECE*, 2016, December 2016.
- [15] E. Molina, L. J. Tardón, I. Barbancho, and A. M. Barbancho. The importance of f0 tracking in query-by-singing-humming. In *ISMIR*, Taipei, Taiwan, 2014.
- [16] Jordi Pons and Xavier Serra. Designing efficient architectures for modeling temporal features with convolutional neural networks. In *ICASSP*, New orleans, USA, 2017.
- [17] Jordi Pons, Olga Slizovskaia, Rong Gong, Emilia Gómez, and Xavier Serra. Timbre analysis of music audio signals with convolutional neural networks. *arxiv.org:1703.06697*, 2017.
- [18] Hasim Sak, Andrew W. Senior, Kanishka Rao, and Françoise Beaufays. Fast and accurate recurrent neural network acoustic models for speech recognition. *CoRR*, abs/1507.06947, 2015.
- [19] J. Sohn, N. S. Kim, and W. Sung. A statistical model-based voice activity detection. *IEEE signal processing letters*, 6(1):1–3, 1999.
- [20] C. C. Wang and J. S. R. Jang. Improving query-by-singing/humming by combining melody and lyric information. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(4):798–806, April 2015.
- [21] L. Wang, S. Huang, S. Hu, J. Liang, and B. Xu. An effective and efficient method for query by humming system based on multi-similarity measurement fusion. In *International Conference on Audio, Language and Image Processing*, pages 471–475, July 2008.
- [22] E. Wichmann. *Listening to Theatre: The Aural Dimension of Beijing Opera*. University of Hawaii Press, 1991.