

LEARNING SOFT-ATTENTION MODELS FOR TEMPO-INVARIANT AUDIO-SHEET MUSIC RETRIEVAL

Stefan Balke¹ Matthias Dorfer¹ Luis Carvalho¹ Andreas Arzt¹ Gerhard Widmer^{1,2}

¹ Institute of Computational Perception, Johannes Kepler University Linz, Austria

² The Austrian Research Institute for Artificial Intelligence (OFAI), Austria

stefan.balke@jku.at

ABSTRACT

Connecting large libraries of digitized audio recordings to their corresponding sheet music images has long been a motivation for researchers to develop new cross-modal retrieval systems. In recent years, retrieval systems based on embedding space learning with deep neural networks got a step closer to fulfilling this vision. However, global and local tempo deviations in the music recordings still require careful tuning of the amount of temporal context given to the system. In this paper, we address this problem by introducing an additional soft-attention mechanism on the audio input. Quantitative and qualitative results on synthesized piano data indicate that this attention increases the robustness of the retrieval system by focusing on different parts of the input representation based on the tempo of the audio. Encouraged by these results, we argue for the potential of attention models as a very general tool for many MIR tasks.

1. INTRODUCTION

Algorithms for content-based search and retrieval play an important role in many applications that are based on large collections of music data. In this paper, we re-visit a challenging cross-modal retrieval problem, namely audio-sheet music retrieval: given a short audio excerpt, we are trying to retrieve the corresponding score from a database of sheet music (stored as images).

Traditional methods for audio-sheet retrieval are based on common mid-level representations that allow for an easy comparison of time points in the audio and positions in the sheet music, see for instance [2, 14, 17, 21]. Typical examples are spectral features like pitch class profiles (chroma-features) or symbolic representations. Unfortunately, deriving such mid-level representations is an error-prone process, as this may involve preprocessing steps such as music transcription [6, 9, 16, 18, 19, 26], Optical Music Recognition [7, 15, 24, 25, 28], and sometimes both.

For a recent overview of different cross-modal retrieval approaches, see [22].

In [11], an alternative approach has been proposed that circumvents these problematic preprocessing steps by learning embeddings directly from the multi-modal data (see Figure 1 for a schematic overview). Given short snippets of audio and their respective excerpts of sheet music images, a cross-modal neural network is trained to learn an embedding space in which both modalities are represented as 32-dimensional vectors (Figure 1a). The vectors can then be easily compared in the embedding space by means of a distance function, e. g., the cosine distance (Figure 1b). The retrieval results are then selected by sorting the candidates by the obtained distances (Figure 1c). A conceptually similar retrieval approach for videos was presented in [1].

In essence, the neural network replaces the step of extracting mid-level features by directly learning a transformation from the audio and from the sheet music image data to a common vector space. A limitation of this approach is that the temporal context (or field of view) of the data in both modalities is of fixed size. For audio data, this implies that the actual musical content of the window depends on the tempo a piece is played in. If it is played in a fast tempo, the audio excerpt will contain a larger amount of musical content (i. e., more notes) than if the same piece is played slowly. Obviously, this can lead to large discrepancies between what the model has seen during training, and the data it will see during test time.

A possible solution to this problem is to let the network decide the appropriate temporal context for a given audio query by using a separate attention mechanism [3,8,23,27]. In a recent workshop paper [12], the concept of using a soft-attention mechanism in a cross-modal retrieval scenario was presented. This allowed the network to deal with a much larger temporal context at the input than without attention. In this paper, we substantially extend this idea with systematic and quantitative experiments with respect to tempo robustness, as well as giving further details about the used architecture and the training data. The remainder of this paper is structured as follows. In Section 2, we introduce the retrieval task and explain the necessary steps to approach it with end-to-end cross-modal embedding learning. Section 3 reports on systematic experiments that show the benefits of using a dedicated attention network on audio spectrograms, to improve retrieval results compared to



© Stefan Balke, Matthias Dorfer, Luis Carvalho, Andreas Arzt, Gerhard Widmer. Licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). **Attribution:** Stefan Balke, Matthias Dorfer, Luis Carvalho, Andreas Arzt, Gerhard Widmer. "Learning Soft-Attention Models for Tempo-invariant Audio-Sheet Music Retrieval", 20th International Society for Music Information Retrieval Conference, Delft, The Netherlands, 2019.

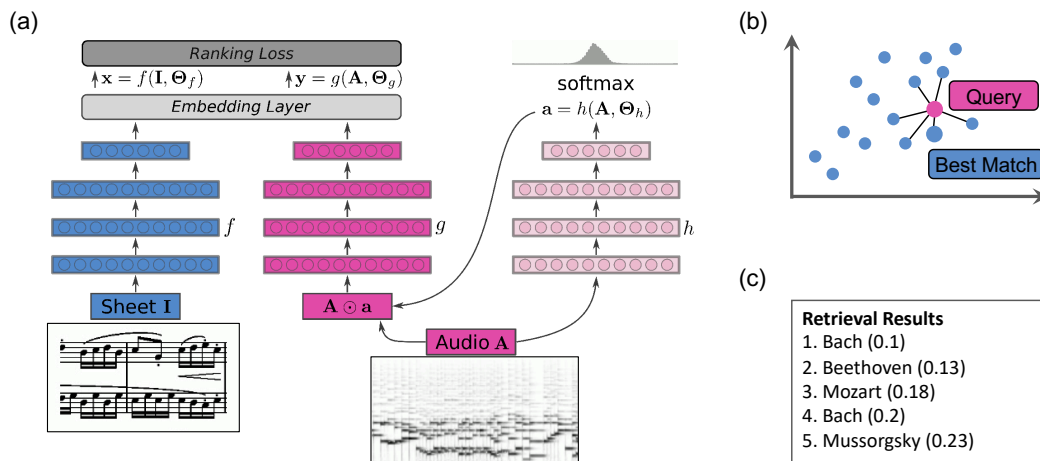


Figure 1: Illustration of the cross-modal retrieval application. (a) Audio snippet serves as input query to the embedding network. The attention network selects the “relevant” region of the audio spectrogram. (b) 2-dimensional visualization of the embedding space. The nearest neighbors are selected as candidates. (c) Ranked list of candidates sorted by their distances to the query embedding.

the current state of the art. Furthermore, we present qualitative examples to highlight the intuitive behavior of the attention network and conclude the paper in Section 4.

2. AUDIO-SHEET RETRIEVAL

We consider a cross-modal retrieval scenario (Figure 1): given an audio excerpt as a search query, we wish to retrieve the corresponding snippet of sheet music of the respective piece. We approach this retrieval problem by learning a low-dimensional multimodal embedding space (32 dimensions) for both snippets of sheet music and excerpts of music audio (Figure 1a). We desire for each modality a projection into a shared space where semantically similar items of the two modalities are projected close together, and dissimilar items far apart (Figure 1b). Once the inputs are embedded in such a space, cross-modal retrieval is performed using simple distance measures and nearest-neighbor search. Finally, the retrieval results are obtained by means of a ranked list through sorting the distances in an ascending order (Figure 1c).

The embedding space is trained with convolutional neural networks (CNN). Figure 1a sketches the network architecture. The baseline model (without attention) consists of two convolutional pathways: one is responsible for embedding the sheet music, and the other for embedding the audio excerpt. The key part of the network is the canonically correlated (CCA) embedding layer [13]. This layer forces the two pathways to learn representations that can be projected into a shared space. The desired properties of this multimodal embedding space are enforced by training with pairwise ranking loss (also known as contrastive loss) [20]. This is the basic structure of the model recently described and evaluated in [11]. This attention-less model serves as a baseline in our experiments, i. e., the input audio (or sheet music) has to be sliced into excerpts of a given size (e. g., 2 s). However, when processing performed music, the temporal context captured by the fixed-size input

Audio (Spectrogram)	Sheet-Image
$92 \times \{42, 84, 168\}$	160×200
2x Conv(3, pad-1)-24 - BN MaxPooling(2)	} Attention Network
2x Conv(3, pad-1)-48 - BN MaxPooling(2)	
2x Conv(3, pad-1)-96 - BN MaxPooling(2)	
2x Conv(3, pad-1)-96 - BN MaxPooling(2)	
Conv(1, pad-0)-32 - Linear GlobalAvgPooling + Softmax Mult(Spectrogram, Attention)	
2x Conv(3, pad-1)-24 - BN MaxPooling(2)	2x Conv(3, pad-1)-24 - BN MaxPooling(2)
2x Conv(3, pad-1)-48 - BN MaxPooling(2)	2x Conv(3, pad-1)-48 - BN MaxPooling(2)
2x Conv(3, pad-1)-96 - BN MaxPooling(2)	2x Conv(3, pad-1)-96 - BN MaxPooling(2)
2x Conv(3, pad-1)-96 - BN MaxPooling(2)	2x Conv(3, pad-1)-96 - BN MaxPooling(2)
Conv(1, pad-0)-32 - Linear - BN Dense(32) + Linear	Conv(1, pad-0)-32 - Linear - BN Dense(32) + Linear
Embedding Layer + Ranking Loss	

Table 1: Overview of the network architecture. The upper part describes the attention network, the lower part the embedding part. Conv(3, pad-1)-24: 3×3 convolution, 24 feature maps and zero-padding of 1. We use ELU activation functions on all layers if not stated otherwise [10].

excerpts varies, based on the current tempo of the piece.

This attention-less model trains and operates on fixed-size input windows from both modalities. In other words, the musical content provided to the CNN may contain significantly less or more information—especially note onsets—than excerpts it has been trained on. One may of course compensate this with data augmentation, but a more general solution is to simply let the model decide how much information is needed from the audio.

For this purpose, we explore a soft-attention mecha-

nism. First, we substantially increase the audio field of view (number of spectrogram frames), up to a factor of four. Next, we add to the model the attention pathway h , which should learn to restrict the audio input again by focusing only at those parts that appear relevant for an efficient search query. As detailed in Table 1, this attention mechanism is implemented as a separate CNN. The output of this CNN is a probability density function a which has the same number of frames as the audio spectrogram. Before feeding the spectrogram to the audio embedding network g , we multiply each frame with its attention weight. This enables the model to cancel out irrelevant parts of the query and focus on the important information. In the following experiments, we show that adding this attention network in combination with the longer temporal context, substantially improves the results in the considered audio-sheet retrieval task.

3. EXPERIMENTS

This section reports on the conducted retrieval experiments. We start by describing the data used for training and testing the models, and the data augmentation steps applied during training. Afterwards, we present the results for the two main experiments, both dealing with audio-sheet retrieval: given an audio excerpt, retrieve the corresponding snippet of sheet music of the respective piece. Finally, we look at the attention-layer’s behavior for five examples and discuss benefits and limitations of the approach.

3.1 Data Description and Training

We use a dataset of synthesized classical piano music, called *MSMD* [11]. In version 1.1, MSMD comprises 467 pieces by 53 composers, including Bach, Mozart, Beethoven and Chopin, totalling in over a thousand pages of sheet music and 15+ hours of audio, with fine-grained cross-modal alignment between note onsets and noteheads. The main changes from version 1.0 (as used in [11]) to version 1.1 are that we cleaned the test set from broken pieces and set all note velocities to a value of 64. The scores and audio are both synthesized using the score engraving software LilyPond¹ and FluidSynth.²

During training, we augment the sheet music by resizing and shifting the image. For augmenting the audio, we vary the tempo between 95 and 110 % and sample from a pool of three different piano soundfonts. For details of these augmentation steps, we kindly refer the reader to the explanation given in [11]. After training convergence, we refine the used CCA embedding layer on the whole training set. The reason for this is that during training, the covariance estimates used in the CCA projection are only based on the number of samples contained in the mini-batch.

For testing, we select 10,000 audio-sheet music pairs from an unseen testset, where the synthesized audio is rendered with a separate, hold-out piano soundfont. The rendering procedure for the sheet music remains the same.

¹<http://www.lilypond.org>

²<http://www.fluidsynth.org>

Model	R@1	R@5	R@25	MRR	MR
BL1-SC [11]	13.67	34.62	57.44	0.24	15
BL2-SC	34.25	54.68	70.62	0.44	4
BL2-MC	36.59	60.52	77.21	0.48	3
BL2-LC	27.66	51.79	70.34	0.39	5
BL2-SC + AT	38.41	59.95	74.36	0.48	3
BL2-MC + AT	46.54	68.45	81.10	0.57	2
BL2-LC + AT	53.16	74.42	85.35	0.63	1

(a) Un-refined CCA layer.

Model	R@1	R@5	R@25	MRR	MR
BL1-SC [11]	19.12	44.16	66.63	0.31	8
BL2-SC	48.91	67.22	78.27	0.57	2
BL2-MC	47.08	68.19	80.82	0.57	2
BL2-LC	43.46	68.38	82.84	0.55	2
BL2-SC + AT	55.43	72.64	81.05	0.63	1
BL2-MC + AT	58.14	76.50	84.60	0.67	1
BL2-LC + AT	66.71	84.43	91.19	0.75	1

(b) Refined CCA layer.

Table 2: Overview of the experimental results. (a) Lists the results without and (b) with a refinement of the CCA layer. All experiments used 10,000 candidates and were conducted on MSMD-v1.1 (R@k = Recall@k, MRR = Mean Reciprocal Rank, MR = Median Rank).

The above described augmentation steps are disabled during testing. In terms of performance, the refinement step of the CCA layer yields an increase in performance of around 0.07 up to 0.15 in terms of mean reciprocal rank (MRR) on the test set (see Table 2 for details). The used dataset,³ as well as the implementation along with trained models,⁴ are publicly available.

3.2 Experiment 1: Attention

In the first experiment, we investigate the influence of the additional attention network. We systematically increase the temporal context of the audio representation from a short context (SC, 42 frames = 2.1 s), over a medium (MC, 84 frames = 4.2 s), to a long context (LC, 168 frames = 8.4 s). The results are summarized in Table 2. As evaluation metrics we use different Recalls (R@1, R@5, R@25 \in [0, 100], higher is better), the mean reciprocal rank (MRR \in [0, 1], higher is better), as well as the median rank (MR \in [1, 10000], lower is better). In the following discussion, we focus on the results of the refined CCA layer, as given in Table 2b.

As a baseline (BL1-SC), we use the network architecture as described in [11], which uses the short context (SC). For this model, approximately 20% of the queries are on rank 1 (R@1 = 19.12) and in almost 70% of the audio

³<https://github.com/CPJKU/msmd/tree/v1.1>

⁴https://github.com/CPJKU/audio_sheet_retrieval/tree/ismir-2019

Model	$\rho = 0.5$	$\rho = 0.66$	$\rho = 1$	$\rho = 1.33$	$\rho = 2$
BL1-SC [11]	0.20	0.27	0.31	0.30	0.22
BL2-SC	0.44	0.52	0.57	0.56	0.46
BL2-MC	0.46	0.53	0.57	0.55	0.43
BL2-LC	0.44	0.50	0.55	0.51	0.35
BL2-SC + AT	0.55	0.63	0.63	0.64	0.56
BL2-MC + AT	0.54	0.61	0.67	0.67	0.62
BL2-LC + AT	0.64	0.69	0.75	0.73	0.64

Table 3: Mean Reciprocal Rank (MRR) for different models and different tempo ratios $\rho \in \{0.5, 0.66, 1, 1.33, 2\}$. For example, $\rho = 0.5$ stands for half of the original tempo and $\rho = 2$ for doubling the tempo. The listed models correspond to the refined CCA models as listed in Table 2b with the same test set size of 10,000 candidates.

queries, the relevant sheet image is within the first 25 ranks ($R@25 = 66.63$). As a second baseline (BL2-SC), we slightly adapt the original architecture by exchanging the global average pooling layer (before the embedding layer) by a dense layer for each modality (see the non-attention part of Table 1 for details). With this adaptation, the results improve significantly to $R@1 = 48.91$, $R@25 = 78.27$, and a median rank $MR = 2$, instead of $MR = 8$ for BL1.

By increasing the temporal context on BL2 to medium sized context (BL2-MC), mean reciprocal rank ($MRR = 0.57$) and median rank ($MR = 2$) stay unchanged. When increasing the temporal context to the long context (BL2-LC), the model degrades in performance, e. g., $R@1$ drops from 48.91% for SC to 43.46% for LC and the MRR from 0.57 to 0.55. Adding the attention network to the audio input (BL2-SC + AT) improves the results by 7, 5, and 3% for the recalls, as well as 0.05 for the MRR compared to BL2-SC. The more context is given to the network, the better the performance metrics get, e. g., $R@1$ improves from 58.14 (BL2-MC + AT) to 66.71 (BL2-LC + AT). The MRR, improves from 0.63 (BL2-SC + AT), over 0.67 (BL2-MC + AT), up to 0.75 (BL2-LC + AT).

We derive two main observations from these results. First, optimizing the network architecture is important; dropping the global average pooling in favour of a fully-connected dense layer lifted the results to another level. The reason could be that the fully-connected layer better retains the structure of the input spectrogram than the average pooling and in addition can be more selective on relevant input parts, e. g., by setting weights to zero. Second, the attention network enables the network to deal with larger temporal context sizes. From a signal processing perspective, one would expect that more context (and thus longer queries) would always help since it increases the specificity of the query. However, since we squash this information into a 32-dimensional embedding vector, it seems that too much information (e. g., too many onsets in the audio), actually harms the retrieval quality when not using attention.




Figure 2: Example of tempo variations within a piece by Johann André - Sonatine (Op. 34, I.). The three boxes below the sheet music show the attention output, and the corresponding audio spectrogram for the respective excerpt.

3.3 Experiment 2: Tempo Robustness

In a second experiment, we explicitly test the system’s robustness to global tempo changes. For this purpose, we re-rendered the MSMD test dataset with various tempo ratios $\rho \in \{0.5, 0.66, 1, 1.33, 2\}$, where $\rho = 0.5$ stands for halving and $\rho = 2$ for doubling the original tempo. The results are given in Table 3. For the sake of brevity, we only show the mean reciprocal rank (MRR) for the refined CCA models.

In general, we observe a similar trend as in the first experiment. While the original baseline approach (BL1-SC) performs rather poorly (0.20 to 0.31), exchanging the global average pooling layer (BL2-SC) helps to improve the performance (0.44 to 0.57). The best attention model (BL2-LC + AT) yields values ranging from 0.64 to 0.75. In this experiment, we would have expected that the improvement holds true for all testsets rendered at different global tempi. However, the numbers tell a slightly different story. To understand this, we had to go back to the generation of the MSMD dataset. Recall that the dataset is generated from LilyPond files obtained from the Mutopia Project.⁵ These files contain tempo specifications which are retained in the synthesis pipeline. The specified tempi vary in a range between 30 and 182 bpm (mean tempo = 106 bpm,

⁵ <https://www.mutopia-project.org>

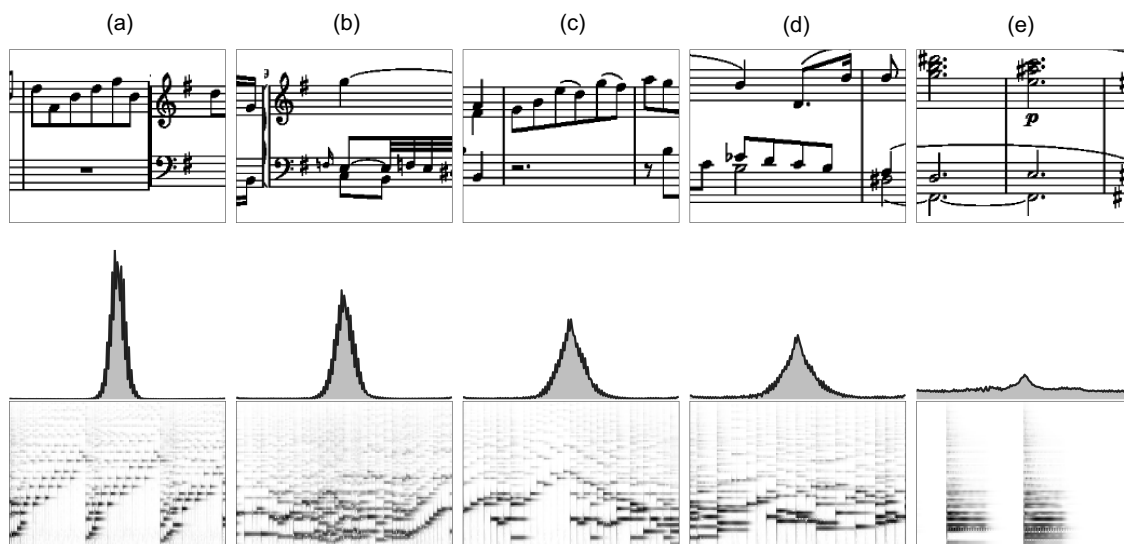


Figure 3: Illustration of the sheet music, input attention (normalized), and spectrogram for five examples from the MSMD test set. (a) L. v. Beethoven - Piano Sonata (Op. 79, 1st Mvt.), (b) J. S. Bach - Goldberg Variations: Variatio 12 (BWV 988), (c) J. S. Bach - French Suite VI: Menuet (BWV 817), (d) R. Schumann - Album for the Youth: Untitled (Op. 68, Nr. 26), and (e) M. Mussorgsky - Pictures at an Exhibition VIII: Catacombae.

std. dev. = 25 bpm). This distribution of tempi implies that all experiments we perform on the original MSMD testset ($\rho = 1$) already test for a variety of tempi. Synthesizing this dataset with different tempo factors—as done in our second experiment—shifts and stretches this distribution. In the edge cases ($\rho = 0.5$ and $\rho = 2$), this leads to unrealistic tempi, e. g., 15 bpm or 364 bpm, producing input audio windows with absurdly low or high onset density.

In summary, all of the tested models have in common that they work best for the original testset tempo ($\rho = 1$) and have similar relative performance drops for the tempo variations. However, the attention model is able to keep the retrieval results at a much higher level than all the other models. In the following section, we take a closer look at some examples from the MSMD test set to get a better intuition of the soft-attention mechanism and how it reacts to tempo deviations.

3.4 Examples

In the experiments above, we have seen that the attention models improve the retrieval results by a considerable margin. Another positive aspect of the soft-attention mechanism is that its behavior is directly interpretable—a rare case in the deep learning landscape. For all the presented examples, we provide further videos with audio on an accompanying website, along with detailed instructions for reproduction.⁶

Figure 2 shows the sheet music for Johann André’s Sonatine, Op.34. The piece starts with a calm part with mainly legato quarter notes at 103 bpm. The first (blue) box shows the corresponding attention output and the au-

dio spectrogram for measure five. The attention output is relatively flat with a small peak in the middle of the audio excerpt. In bar 17, the note density changes, with eighth notes in the left hand entering the scene. As shown by the second (red) box, the attention layer reacts to this by building up a more distinct focus around the center, placing less weight on the outer parts of the spectrogram. In the third part beginning with measure 24, the “tempo” (more precisely: the perceived “speed”, in terms of events per time unit) essentially doubles, which is reflected in a high peak and narrow distribution in the attention output (green right box). From these three examples, it also becomes clear that tempo and note density are essentially two sides of the same coin for the attention network.

In Figure 3, we show examples from the testset with different global performance tempi, along with the sheet music excerpt, the attention weights, and the spectrogram. The fastest piece, with around 250 bpm, is shown in Figure 3a. The corresponding attention output is very focused on the middle part of the audio spectrogram, trying to concentrate the attention to the notes that actually appear in the sheet music snippet. The second example (b) is rather slow with 95 bpm. However, through the use of 16th and 32nd notes, the main melody gets a double-time feel, thus the actual tempo is perceived at around 190 bpm. In Figure 3c, the tempo is at 115 bpm. The attention output starts to widen up, allowing more temporal context to reach the actual embedding network. This trend goes in Figure 3d when the tempo 78 bpm. Figure 3e shows an extreme example where the piece mainly consists of chords with long, dotted half notes. Here, the attention has to use the complete temporal context of the audio spectrogram to match the score information with the audio.

The examples demonstrate that depending on the spec-

⁶ http://www.cp.jku.at/resources/2019_ASR-TempoInv_ISMIR

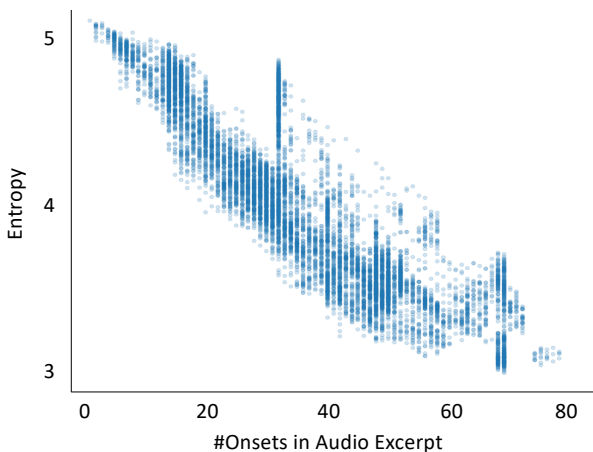


Figure 4: Entropy of the input attention vs. the number of onsets in the respective audio frame.

rogram content, the model indeed attends to whatever it believes is a representative counterpart of the target sheet music snippet. Since the fixed-size sheet snippets contain roughly similar numbers of notes, as the density of note heads on the printed page tends to be independent of the tempo of the piece, attention is sharply peaked when the density of onsets in the audio is high, and conversely is distributed more evenly when there are fewer notes in the audio excerpt.

4. SUMMARY

In this paper, we have described a soft-attention mechanism that helps to overcome the fixed window sizes uses in Convolutional Neural Networks. In our end-to-end audio-sheet music retrieval application, the results improved substantially compared to the state of the art. By looking at a number of examples from the retrieval results, the soft-attention mechanism showed an intuitive and interpretable behavior.

This appealing and intuitive behavior is summarized in Figure 4, which shows the entropy of the attention distribution in relation to the number of onsets contained in the audio excerpt, for all 10,000 test samples. The entropy is a measure of flatness of the attention output: a flat function gives high entropy, a very narrow function low entropy. The downward trend in the figure confirms our observations from above: the more onsets in the respective audio spectrogram, the narrower the attention distribution.

Given the improved retrieval performance and the intuitive behavior of the attention model, we think this is a promising line of research for reducing the sensitivity of cross-modal music retrieval models to the audio input window size. To this end, our experiments were conducted on synthesized piano recordings. However, results in [11] indicate that the embedding models trained on this data generalize to real scores and performances. A possible next step would be to investigate whether the attention mechanism reacts to local tempo changes as occurring frequently in real performances (e. g., ritardandi and accelerandi). Furthermore, it would be interesting to

leave the piano music domain and extend the model to cope with differences in timbre, (e. g., orchestral music) as done in [4] for the challenging Barlow–Morgenstern scenario [5].

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement number 670035, project CON ESPRESSIONE, and the Marie Skłodowska-Curie grant agreement number 765068, MIP-Frontiers).

5. REFERENCES

- [1] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 609–617, Venice, Italy, 2017.
- [2] Andreas Arzt, Sebastian Böck, and Gerhard Widmer. Fast identification of piece and score position via symbolic fingerprinting. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 433–438, Porto, Portugal, 2012.
- [3] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*, pages 1–15, San Diego, CA, USA, 2015.
- [4] Stefan Balke, Vlora Arifi-Müller, Lukas Lamprecht, and Meinard Müller. Retrieving audio recordings using musical themes. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 281–285, Shanghai, China, 2016.
- [5] Harold Barlow and Sam Morgenstern. *A Dictionary of Musical Themes*. Crown Publishers, Inc., revised edition edition, 1975.
- [6] Sebastian Böck and Markus Schedl. Polyphonic piano note transcription with recurrent neural networks. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 121–124, Kyoto, Japan, March 2012.
- [7] Donald Byrd and Jakob G. Simonsen. Towards a standard testbed for optical music recognition: Definitions, metrics, and page images. *Journal of New Music Research*, 44(3):169–195, 2015.
- [8] William Chan, Navdeep Jaitly, Quoc V. Le, and Oriol Vinyals. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4960–4964, Shanghai, China, 2016.

- [9] Tian Cheng, Matthias Mauch, Emmanouil Benetos, and Simon Dixon. An attack/decay model for piano transcription. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 584–590, New York City, United States, 2016.
- [10] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (ELUs). In *Proceedings of the International Conference on Learning Representations (ICLR) (arXiv:1511.07289)*, 2016.
- [11] Matthias Dorfer, Jan Hajič jr., Andreas Arzt, Harald Frostel, and Gerhard Widmer. Learning audio–sheet music correspondences for cross-modal retrieval and piece identification. *Transactions of the International Society for Music Information Retrieval*, 1(1), 2018.
- [12] Matthias Dorfer, Jan Hajič jr., and Gerhard Widmer. Attention as a perspective for learning tempo-invariant audio queries. In *Proceedings of the ICML Joint Workshop on Machine Learning for Music*, Stockholm, Sweden, 2018.
- [13] Matthias Dorfer, Jan Schlüter, Andreu Vall, Filip Korzeniowski, and Gerhard Widmer. End-to-end cross-modality retrieval with CCA projections and pairwise ranking loss. *International Journal of Multimedia Information Retrieval*, 7(2):117–128, Jun 2018.
- [14] Christian Fremerey, Michael Clausen, Sebastian Ewert, and Meinard Müller. Sheet music-audio identification. In *Proceedings of the International Conference on Music Information Retrieval (ISMIR)*, pages 645–650, Kobe, Japan, October 2009.
- [15] Jan Hajič jr., Jiri Novotný, Pavel Pecina, and Jaroslav Pokorný. Further steps towards a standard testbed for Optical Music Recognition. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 157–163, New York City, United States, 2016.
- [16] Curtis Hawthorne, Erich Elsen, Jialin Song, Adam Roberts, Ian Simon, Colin Raffel, Jesse Engel, Sageev Oore, and Douglas Eck. Onsets and frames: Dual-objective piano transcription. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 50–57, Paris, France, 2018.
- [17] Özgür Izmirli and Gyanendra Sharma. Bridging printed music and audio through alignment using a mid-level score representation. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 61–66, Porto, Portugal, 2012.
- [18] Rainer Kelz, Sebastian Böck, and Gerhard Widmer. Deep polyphonic ADSR piano note transcription. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Brighton, United Kingdom, May 2019.
- [19] Rainer Kelz, Matthias Dorfer, Filip Korzeniowski, Sebastian Böck, Andreas Arzt, and Gerhard Widmer. On the potential of simple framewise approaches to piano transcription. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 475–481, New York City, USA, 2016.
- [20] Ryan Kiros, Ruslan Salakhutdinov, and Richard S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint (arXiv:1411.2539)*, 2014.
- [21] Frank Kurth, Meinard Müller, Christian Fremerey, Yoon ha Chang, and Michael Clausen. Automated synchronization of scanned sheet music with audio recordings. In *Proceedings of the International Conference on Music Information Retrieval (ISMIR)*, pages 261–266, Vienna, Austria, September 2007.
- [22] Meinard Müller, Andreas Arzt, Stefan Balke, Matthias Dorfer, and Gerhard Widmer. Cross-modal music retrieval and applications: An overview of key methodologies. *IEEE Signal Processing Magazine*, 36(1):52–62, 2019.
- [23] Chris Olah and Shan Carter. Attention and augmented recurrent neural networks. *Distill*, 2016.
- [24] Christopher Raphael and Jingya Wang. New approaches to optical music recognition. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 305–310, Miami, Florida, USA, 2011.
- [25] Ana Rebelo, Ichiro Fujinaga, Filipe Paszkiewicz, Andre R. S. Marcal, Carlos Guedes, and Jaime S. Cardoso. Optical music recognition: state-of-the-art and open issues. *International Journal of Multimedia Information Retrieval*, 1(3):173–190, 2012.
- [26] Siddharth Sigtia, Emmanouil Benetos, and Simon Dixon. An end-to-end neural network for polyphonic piano music transcription. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(5):927–939, 2016.
- [27] Carl Southall, Ryan Stables, and Jason Hockman. Automatic drum transcription for polyphonic recordings using soft attention mechanisms and convolutional neural networks. In *Proceedings of the International Society for Music Information Retrieval Conference (ISMIR)*, pages 606–612, Suzhou, China, 2017.
- [28] Cuihong Wen, Ana Rebelo, Jing Zhang, and Jaime S. Cardoso. A new Optical Music Recognition system based on combined neural network. *Pattern Recognition Letters*, 58:1–7, 2015.