

AlignNet: A Unifying Approach to Audio-Visual Alignment

Jianren Wang*
CMU

jianrenw@andrew.cmu.edu

Zhaoyuan Fang*
University of Notre Dame

zfang@nd.edu

Hang Zhao
MIT

hangzhao@csail.mit.edu

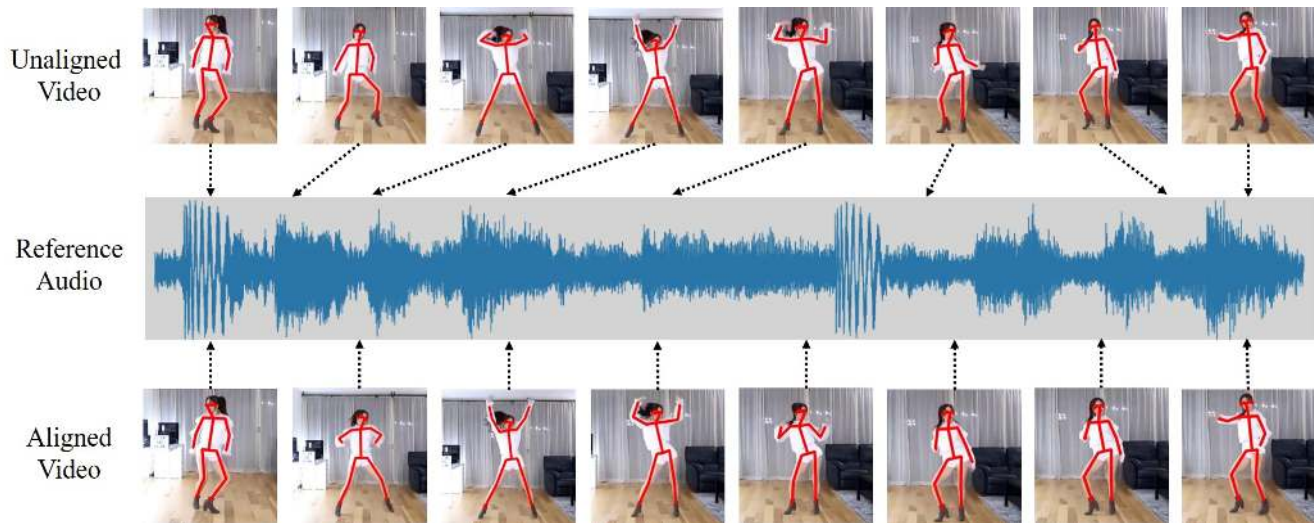


Figure 1: Given a pair of unaligned audio and video, our model aligns the video to the audio according to the predicted dense correspondence.

Abstract

We present *AlignNet*, a model that synchronizes videos with reference audios under non-uniform and irregular misalignments. *AlignNet* learns the end-to-end dense correspondence between each frame of a video and an audio. Our method is designed according to simple and well-established principles: attention, pyramidal processing, warping, and affinity function. Together with the model, we release a dancing dataset *Dance50* for training and evaluation. Qualitative, quantitative and subjective evaluation results on dance-music alignment and speech-lip alignment demonstrate that our method far outperforms the state-of-the-art methods. Code, dataset and sample videos are available at our project page¹.

1. Introduction

Dancers move their bodies with music, speakers talk with lip motions and are often accompanied by hand and arm gestures. The synchrony between visual dynamics and audio rhythms poses perfect performances. However, recorded videos and audios are not always temporally aligned. *e.g.* dancers may not be experienced enough to follow the music beats precisely; “Automated Dialogue Replacement (ADR)” is used in film making instead of simultaneous sounds for lip synchronization. It is not hard to imagine the humongous amount of efforts and time required to temporally synchronize videos and audios by aligning visual dynamics and audio rhythms. This alignment problem is especially difficult for humans since the misalignment is often non-uniform and irregular.

There are several previous attempts to address this problem. [8] extracted visual beats analogous to musical beats, and applied audio-video alignment to tasks such as video dancification and dance retargeting. Their method requires feature engineering on visual beats and only performs alignment on beat-level, which makes the generated dances

* indicates equal contribution

¹<https://jianrenw.github.io/AlignNet/>

sometimes exaggerated. SyncNet [6] proposed to learn the alignment features with a contrastive loss that discriminates matching pairs from non-matching pairs. However, they assume a global temporal offset between the audio and video clips when performing alignment. [14] further leveraged the pre-trained visual-audio features of SyncNet [6] to find an optimal alignment using dynamic time warping (DTW) to assemble a new, temporally aligned speech video. Their method can therefore stretch and compress the signals dynamically. However, the minimum temporal resolution of their method is limited (0.2 second). In a realistic setting, *e.g.* ADR or off-beat dancing, misalignment between audio and video can happen at any moment on arbitrary temporal scale. Therefore, a crucial property for audio-visual alignment solutions is the ability to deal with arbitrary temporal distortions. Another nice property to have is end-to-end training, as it significantly simplifies the solution to such a hard task.

In this work, we present AlignNet, an end-to-end trainable model that learns the mapping between visual dynamics and audio rhythms, without the need of any hand-crafted features or post-processing. Our method is designed according to simple and well-established principles: attention, pyramidal processing, warping, and affinity function. First, attention modules highlight the important spatial and temporal regions in the input. Then, casting in two learnable feature pyramids (one for video and one for audio), AlignNet uses the current level correspondence estimation to warp the features of the reference modality (video or audio). Warped features and reference signal features are used to construct an affinity map, which is further processed to estimate denser correspondence.

Together with the model, we introduce a dancing dataset *Dance50* for dance-music alignment. The dataset is cleaned and annotated with human keypoints using a fully automated pipeline.

We demonstrate the effectiveness of AlignNet on both dance-music alignment and lip-speech alignment tasks, with qualitative and quantitative results. Finally we conduct dance and speech retargeting experiments, and show the generalization capabilities of our approach with subjective user studies.

2. Related Works

Audio-video alignment Audio-video alignment refers to the adjustment of the relative timing between audio and visual tracks of a video. Automatic audio-video alignment has been studied over decades in computer vision. Early works like [2] and [23] used canonical correlation analysis (CCA) for synchronization prediction. Later methods tried to align video and audio based on handcrafted features. Lewis [16] proposed to detect phonemes (short units of speech) and subsequently associate them with mouth po-

sitions to synchronize the two modalities. Conversely, [11] classified parameters on the face into visemes (short units of visual speech), and used a viseme-to-phoneme mapping to perform synchronization. [8] split videos according to visual beats and applied video-audio alignment on beat-level.

Learning multimodal features is a recent trend. SyncNet [6] learned a joint embedding of visual face sequences and corresponding speech signals in a video by predicting whether a given pair of face sequence and speech track are synchronized. Similarly, [20] proposed a self-supervised method to predict the alignment of motion and sound within a certain time shift. These works attempted to detect and correct a global error, which is a common problem in TV broadcasting. However, they cannot address non-uniform misalignment, *e.g.* dancers do not only make mistakes at musical beats. In other words, misalignment in videos and audios is oftentimes completely unconscious and irregular. In these scenarios, the closest method to our work is proposed by [14], which can stretch and compress small units of unaligned video signal to match audio signal. However, their method can only adjust audio-video misalignment on a coarser granularity, since they assume the consistency of information within every 0.2 second.

Time Warping Given two time series, $X = [x_1, x_2, \dots, x_{n_x}] \in R^{d \times n_x}$ and $Y = [y_1, y_2, \dots, y_{n_y}] \in R^{d \times n_y}$, dynamic time warping (DTW) [22] is a technique to optimally align the samples of X and Y such that the following sum-of-squares cost is minimized: $C(P) = \sum_{t=1}^m \|x_{p_t^x} - y_{p_t^y}\|^2$, where m is the number of indices (or steps) needed to align both signals. Although the number of possible ways to align X and Y is exponential in n_x and n_y , dynamic programming [1] offers an efficient way to minimize C using Bellman’s equation. The main limitation of DTW lies in the inherent inability to handle sequences of varying feature dimensionality, which is essential in multimodal data, like video and audio. Furthermore, DTW is prone to failure when one or more sequences are perturbed by arbitrary affine transformations. To this end, the Canonical Time Warping (CTW) [31] is proposed, which elegantly combines the least-squares formulations of DTW and Canonical Correlation Analysis (CCA) [9], thus facilitating the utilization of sequences with varying dimensionality, while simultaneously performing feature selection and temporal alignment. Deep Canonical Time Warping (DCTW) [26] further extends CTW through the usage of deep features. These methods are widely used in audio-audio alignment [15, 18], video-audio alignment [14] and video-video alignment [25].

Audio-visual Datasets Most existing audio-visual datasets focus on speech and face motion. Some early datasets are obtained under controlled conditions: forensic

data intercepted by police officials [27], speech recorded from mobile devices [28], *etc.* In contrast, VoxCeleb [19, 5] and AVSpeech [10] collected a large amount of audio-visual data of human speeches in the wild. MUSIC [30, 29] and FAIR-Play [12] are datasets of music and instruments. In this work, we introduce a new dancing dataset to the community, focusing on the synchronization between music and body motions.

3. Approach

3.1. Formulation

AlignNet is an end-to-end trainable model that learns the implicit mapping between visual dynamics and audio rhythms at multiple temporal scales, without the need of any post-processing. It takes in a video feature sequence $\mathbf{F}_v^t, t \in (1, \dots, n)$ and an audio feature sequence $\mathbf{F}_a^t, t \in (1, \dots, m)$, to predict a dense correspondence \mathbf{d} , which is the temporal displacement for each video frame (n, m might not be the same). To generate unaligned training data, we apply random and non-uniform distortion temporally (speed-ups and slow-downs) to the aligned audio-video data of different lengths.

Different from previous methods [6, 7, 14] that use raw image frames as inputs, which are prone to large variations when there are changes in outfits, makeups, lighting, shadows, *etc.*, we instead use more robust pose/lip keypoint features provided by a keypoint detector OpenPose [3]. Furthermore, as suggested by [4], operating in the velocity space improves the results of human pose forecasting. For our task, at time t , we input velocity $\mathbf{v}_v(t)$ and the acceleration $\mathbf{a}_v(t)$ of the keypoints. For audio, we use normalized log-mel spectrogram as input. Note that the input video features and audio features do not have the same shape and granularity, but our proposed network learns to extract information from different modalities and align them on the same scale.

The full pipeline of AlignNet is shown in Figure 2. It is designed based on the following principles: (1) spatial and temporal attention modules highlight the important regions; (2) learnable feature pyramids extract features from visual and audio inputs at multiple levels (temporal scales); (3) inspired by [24], warping layers warp the lower-level reference modality feature based on the correspondence estimation on the current level to better predict distortion on larger scales; (4) correlation layers models the affinity between video features with audio features; (5) final dense correspondence estimation is predicted from the affinity map. We will explain each module in details in the following paragraphs.

3.2. Spatial and Temporal Attention Modules

Intuitively, certain keypoints are more reflective of the visual rhythms than others, so are certain sound fragments of the auditory beats. Hence, we use attention modules to highlight these important parts. Applying spatio-temporal attention jointly requires excessive parameters, so we propose to decouple it, using a spatial attention module and a temporal attention module.

For the spatial attention, we assign a weight to each keypoint, indicating the importance of this spatial location. To make the learning more effective and meaningful, we force the attention of symmetric keypoints (*e.g.* left and right elbow) to be the same. Note that we assume keypoint attentions are the same over the whole dataset, so they are learned as model parameters (the number of parameters is the same as the number of keypoints).

For the temporal attention, we train a self-attention module conditioned on the input audio, and re-weight audio features at each time step accordingly. Different from keypoint attention, temporal attention varies with the input data, so the attention weights are the intermediate outputs of the model.

3.3. Multimodal Feature Pyramid Extraction

For each modality, we generate an N -level feature pyramid, with the 0-th level being the input features, *i.e.*, $\mathbf{F}_p^0 = \mathbf{F}_{v/a}^t, t \in (1, \dots, n/m)$. Convolutional layers then extract the feature representation at the l -th level, \mathbf{F}_p^l , from the feature representation at the $(l - 1)$ -th level, \mathbf{F}_p^{l-1} .

3.4. Time Warping

The concept of warping is common to flow-based methods and goes back to [17]. We introduce time warping layer to warp the video features given the estimated correspondences. The warping operation assesses previous errors and computes an incremental update, helping the prediction module to focus on the residual misalignment between video and audio. Formally, at the l -th level, we warp reference audio features toward video features using the upsampled prediction from the $(l + 1)$ -th level, \mathbf{d}^{l+1} :

$$\mathbf{F}_w^l(\mathbf{x}) = \mathbf{F}_a^l(\text{up}_\gamma(\mathbf{d}^{l+1})(\mathbf{x})) \quad (1)$$

where \mathbf{x} is the frame indices and up_γ is the upsampling function with scale factor γ .

3.5. Affinity Function

In order to better model the correspondence, we employ an affinity function A at each feature pyramid level. A provides a measure of similarity between video features $\mathbf{F}_v^l, t \in (1, \dots, n^l)$ and audio feature $\mathbf{F}_a^l, t \in (1, \dots, m^l)$, denoted as $A(\mathbf{F}_{p_v}^l, \mathbf{F}_{p_a}^l)$. The affinity is the dot product

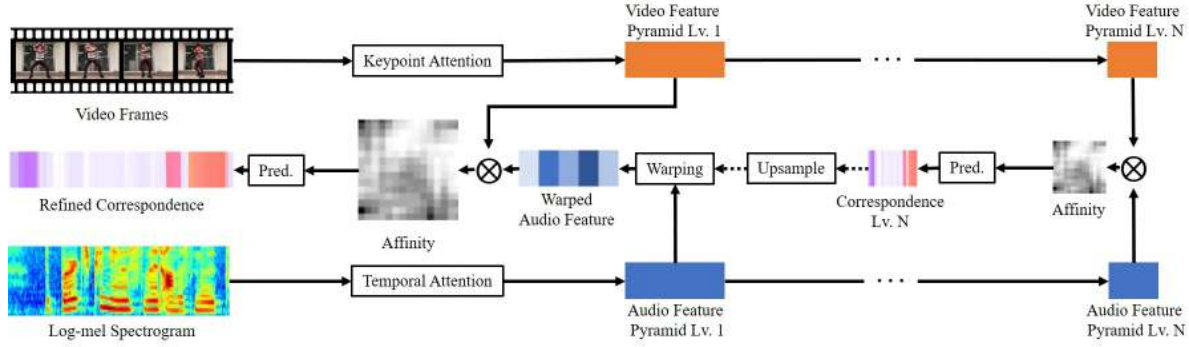


Figure 2: Our model consists of attention, pyramidal processing, time warping, affinity and correspondence prediction.

between embeddings: for $\mathbf{F}_v^{l_i}$ and $\mathbf{F}_a^{l_j}$,

$$A(j, i) = \frac{\left(\mathbf{F}_a^{l_j \top} \mathbf{F}_v^{l_i} \right)}{\sum_j \left(\mathbf{F}_a^{l_j \top} \mathbf{F}_v^{l_i} \right)} \quad (2)$$

3.6. Multimodal Dense Correspondence Estimation

Correspondences are normalized to $(-1, 1)$. For a video with N frames, we first include a correspondence loss, namely an L1 regression loss of the predicted correspondence at all pyramid levels:

$$L_{fs} = \sum_{l=0}^N \lambda_l \sum_{i=0}^{n^l} \left| \mathbf{d}^l(i) - \tilde{\mathbf{d}}^l(i) \right| \quad (3)$$

where $\mathbf{d}^l(i)$ is the ground truth correspondence and $\tilde{\mathbf{d}}^l(i)$ is the predicted correspondence of frame i at level l , and λ_l is the weighting factor of level l .

To force the network to predict realistic correspondence that are temporally monotonic, we further incorporate a monotonic loss:

$$L_{mono} = \sum_{l=0}^N \lambda_l \sum_{i=0}^{n^l-2} \max \left(0, 1 - \tilde{\mathbf{d}}^l(i) + \tilde{\mathbf{d}}^l(i+1) \right) \quad (4)$$

With μ being a weighting hyperparameter, our full objective to minimize is therefore:

$$loss = L_{fs} + \mu L_{mono} \quad (5)$$

3.7. Model Architecture Details

AlignNet takes in normalized log-mel spectrograms with 128 frequency bins as audio input and pose/lip keypoints as visual input with shape $(2 \times \#keypoints, \#frames)$, where the 2 accounts for two dimensional (x,y) coordinates. Note that all convolutional layers in our network are 1D convolutions, which speed up the prediction at both training and testing time. We use a 4-level feature pyramid, with channels 128, 64, 32, 16 and the temporal downscale factors 1/3,

1/2, 1/2, 1/2, respectively. The final output is obtained by upsampling the 1st-level correspondence to match with the input dimension. Note that our model can take in videos of arbitrary lengths.

4. Experiments

We evaluate AlignNet on two tasks: dance-music alignment, and lip-speech alignment. We test our method, both quantitatively and qualitatively on these tasks, and compare our results with two benchmark methods. Finally, subjective user studies are performed for dance and speech re-targeting.

Baselines In order to compare our methods with other state-of-the-art audio-video synchronization techniques, we implemented two baselines: SyncNet [6] and SyncNet+DTW [14]. For SyncNet, we follow the exact steps described in [6] except that we replace raw frame inputs with pose/lip keypoints and replace MFCC feature inputs with normalized log-mel spectrograms. Since (1) the relevant features to dancing and speaking are essentially body and lip motion and (2) MFCC features are handcrafted and low-dimensional (*i.e.* contain less information than spectrograms), these changes should reach similar performance of the original method, if not better. SyncNet+DTW uses the pre-trained SyncNet as their feature extractor, so we also replace the SyncNet here with our implementation. Then, DTW is applied on the extracted video and audio features.

Training During training, we use online sample generation for both our method and the baselines. In our method, the video features are randomly distorted and scaled temporally each time to prevent the network from overfitting finite misalignment patterns. To better mimic real world situations, the distortions are always temporally monotonic, and we linearly interpolate the keypoints from adjacent frames. For the baselines, SyncNet is trained with random pairs of matching / non-matching pairs with the same clip length and maximum temporal difference as the original paper [6]. We

Methods	Performance	
	AFE	Accuracy
SyncNet [6]	6.58	25.88
SyncNet+DTW [14]	4.27	38.29
AlignNet (Ours)	0.94	89.60

Table 1: Performance of music-dance alignment on *Dance50* dataset.

further augment the training data by horizontally flipping all poses in a video clip, which is a natural choice because visual dynamics is mirror-invariant. For the audio inputs, we employ the time masking and frequency masking suggested by SpecAugment [21]. *Adam* is used to optimize the network parameters with a learning rate of 3×10^{-4} .

Evaluation Quantitative evaluation was performed using a human perception-inspired metric, based on the maximum acceptable audio-visual asynchrony used in the broadcasting industry. According to the International Telecommunications Union (ITU), the auditory signal should not lag by more than 125 ms or lead by more than 45 ms. Therefore, the accuracy metric we use is the percentage of frames in the aligned signal which fall inside of the above acceptable (human undetectable) range, compared to the ground truth alignment.

We also defined Average Frame Error (AFE) as the average difference between the reconstructed frame indices and the original undistorted frame indices. This gives a more direct measure of how close the reconstructed video is compared to the original video.

4.1. Dance-music Alignment

Dance50 Dataset We introduce *Dance50* dataset for dance-music alignment. The dataset contains 50 hours of dancing clips from over 200 dancers. There are around 10,000 12-second videos in the training set, 2,000 in the validation set, and 2,000 in the testing set. Our training, validation, and testing sets contain disjoint videos, such that there is no overlap between the videos from any two sets. All dancing clips are collected from K-pop dance cover videos from *YouTube.com* and *Bilibili.com*. Similar to [13], we represent our annotations of the poses with a temporal sequence of 2D skeleton keypoints, obtained using OpenPose [3] BODY_25 model, discarding 6 noisy keypoints from the feet and keeping the rest. We refer the readers to supplementary material for more details.

Settings We use *Dance50* to evaluate the performance of method on dance synchronization. We compare our method against the state-of-the-art baselines using the aforementioned evaluation metrics.

Results Table 1 compares the performance of our method and the two baselines on *Dance50* testing set. Our method significantly reduces the AFE, and achieves a gain of 51.31% in accuracy, which reaches 89.60%. Our network obtains a satisfactory performance even on video clips completely unseen in training, and this shows that the proposed method effectively learns the implicit mapping between video dynamics and audio features on multiple temporal scales.

To get a better idea of how well our method aligns the video to the audio compared to the baselines methods, we show some visualization of the skeletons reconstructed from the correspondence predictions in Figure 3. It can be seen that our method closely recovers the original aligned poses, while both SyncNet and DTW fails to align poses warped on different temporal scales.

We also show the mean motion error and location error of each human keypoint after alignment in Figure 4. The motion error reflects visual rhythm difference between aligned video and original video (Figure 4 Left), while the location error reflects the objective difference between aligned video and original video (Figure 4 Right). For motion error, we normalize the keypoint velocity in pixel space of both x and y coordinates to (-1,1). Similarly, we normalize the location of each keypoint for calculating location error. As shown in the figure, for both motion and location, our proposed method outperforms the baseline method consistently by a large margin under all human keypoints. It’s worth noticing that mid-hip locations are always subtracted during preprocessing. Thus, the mid-hip error of both motion and location is significantly smaller than other keypoints.

Furthermore, we plot examples of the ground truth distortion compared to the distortion predicted by SyncNet-DTW and our method in Figure 6 (left). It can be seen that SyncNet-DTW captures some of the mapping between video and audio very coarsely, while failing to align with more subtle changes. We believe that this is because SyncNet-DTW uses SyncNet as its backbone network, which is trained with blocks of matching / non-matching pairs, assuming that the frames within each block is temporally aligned. In our application scenario, however, the video could be distorted at any moment on any temporal scale, so their method cannot deal with such distortions very well. In contrast, the prediction from our method closely matches the ground truth distortion and is able to undistort & recover the input video. We believe that this is because (1) the hierarchical nature of our model helps the network extract visual and auditory features on multiple temporal scales; (2) the warping layer makes it easier to predict coarser trends on larger temporal scales; and (3) the affinity map helps to correlate the relationship between video features and audio features. We also show the effectiveness of each model in the ablation study later.

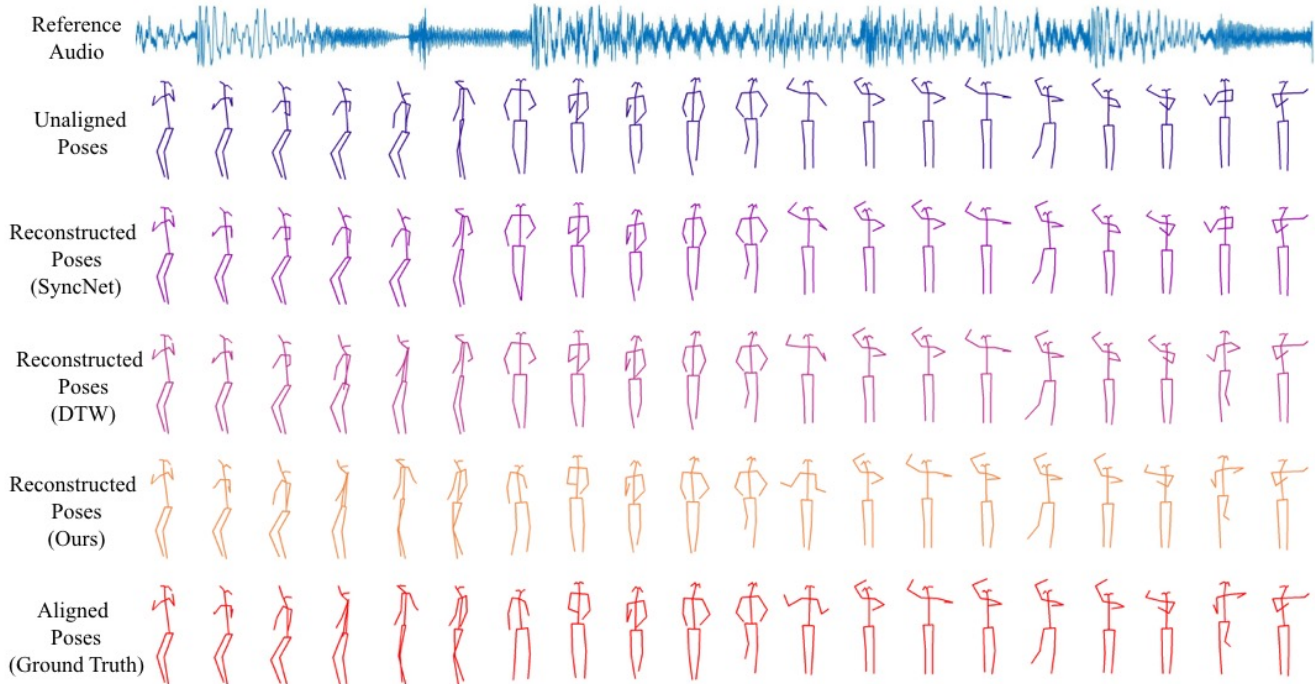


Figure 3: Skeleton visualizations of the proposed method and comparison with the baselines on *Dance50*. For better visualization, we evenly sample 20 frames from the original 180 frames.




Figure 4: Comparison of the performance (mean motion & location error) of our method and SyncNet + DTW (SN+DTW) on the dance-music alignment task.

4.2. Lip-speech Alignment

Settings To evaluate our model on the lip-speech alignment task, we assemble a subset from the *VoxCeleb2* dataset [5], which is an audio-visual dataset consisting of short clips of human speech extracted from YouTube videos. We select and cut 22000 video clips of length 12 seconds and with 25 fps, mounting to 73 hours in total. The dataset contains 333 different identities and a wide range

of ethnicity and languages. The dataset split is as follows: the training set has 18000 clips, and the validation and testing sets each have 2000 clips. The sets are exclusively video-disjoint, and there is no overlap between any pair of videos from any two sets. Similar to *Dance50*, to deal with the noisy frame-wise results, we first remove the outlier keypoints by performing median filtering and then the missing values are linearly interpolated and smoothed with a Savitzky-Golay filter. Again, we compare our method against the state-of-the-art baselines using the evaluation metrics mentioned above, and all methods are trained with online sample generation and data augmentation, for 500 epochs, and the best performing model on the validation set is selected for testing.

Results Similar to Section 4.1, we show the mean motion error and location error of each lip keypoint after alignment in Figure 5. Keypoints labels are the same as used in OpenPose [3]. As shown in the figure, for both motion and location, our proposed method outperforms the baseline method consistently by a large margin for all the lip keypoints.

Table 2 compares the performance of our method and the two baselines on subset of *VoxCeleb2* testing set. Our method significantly reduces the AFE, and achieves a gain of 45.11% in accuracy, which achieves 81.05%. Note that here the AFE is similar to the AFE on *Dance50*, but the




Figure 5: Comparison of the performance (mean motion & location error) of our method and SyncNet + DTW (SN+DTW) on the lip-speech alignment task.

Methods	Performance	
	AFE	Accuracy
SyncNet [6]	6.33	27.41
SyncNet+DTW [14]	3.96	35.94
AlignNet (Ours)	1.03	81.05

Table 2: Performance of speech-lip alignment on *VoxCeleb2* dataset.

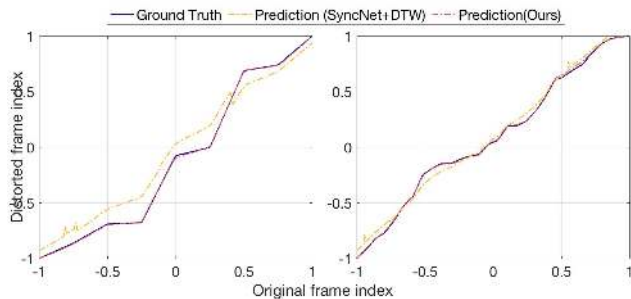


Figure 6: Qualitative results of the ground truth distortion compared to the predicted distortions by SyncNet-DTW and our method. Left: dance-music alignment, Right: lip-speech alignment.

accuracy on *Dance50* is much higher. This is because the frame rates of the two datasets are different (30 fps for *Dance50* and 25 fps for *VoxCeleb2*), and one frame in *VoxCeleb2* corresponds to a longer duration in ms. Therefore, these results do not suggest the relative difficulty of dance-music alignment and speech-lip alignment, but they together indicate that our proposed method for audio-video alignment works well in different application scenarios.

Methods	Performance	
	AFE	Accuracy
Base	2.87	43.57
FP	2.45	56.88
FP+MI	1.81	67.49
FP+MI+SA	1.32	75.33
FP+MI+SA+TA	1.14	79.65
FP+MI+SA+KA	0.97	88.20
FP+MI+SA+KA+TA	0.94	89.60

Table 3: Ablation study on dance-music alignment by adding these modules sequentially: Base (base model), FP (feature pyramid), MI (motion inputs, velocity and acceleration), SA (spectrogram augmentation), KA (keypoint attention) and TA (temporal attention).

Similar to the dance-music alignment task, we plot an example of the ground truth distortion compared to the distortion predicted by DTW and our method in the bottom part of Figure 6 (Right). This result clearly shows that our method recover the distorted video at a very high accuracy.

4.3. Ablation Study

To show the effectiveness of each module in our method, we conduct an ablation study on the *Dance50* dataset, where we add the modules sequentially to the pipeline and evaluate the performance. The training setup is the same as before. We start with a basic network, denoted as base, where direct pose keypoints are fed in as inputs and no feature pyramid is used, predicting the correspondence directly from an affinity map. Then, the following modules are added sequentially to the network: FP (feature pyramid), MI (motion inputs, velocity and acceleration), SA (spectrogram augmentation from [21]), KA (keypoint attention) and TA (temporal attention). Table 3 shows the results of the ablation experiments. It is clear that both feature pyramid and motion inputs boost the performance of the network by a large margin, and spectrogram augmentation further helps the network to learn meaningful feature even when some temporal and frequency information are missing and helps alleviate overfitting.

Both keypoints attention and temporal attention can improve the alignment performance. This means some keypoints and time steps are more important for alignment. And focusing on these spatial temporal regions will significantly improve the performance. However, temporal attention is not as effective as keypoints attention. One possible reason is that most music pieces have fixed and similar beats, which can be implicitly encoded in the network without temporal attention module. Thus, adding temporal might have limited improvement. On the other hand,

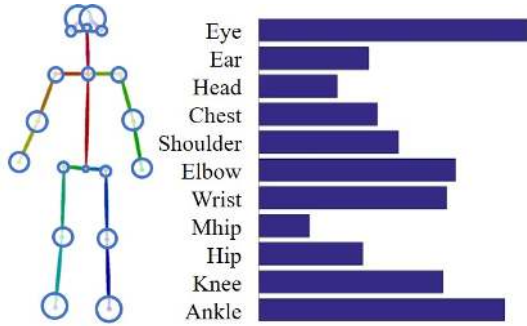


Figure 7: Keypoint attention visualization for the dance-music alignment task. Circle sizes and bar lengths represent the attention magnitudes on human keypoints.

motion rhythms are not always as clear as music beats, *i.e.* dancer might dance differently given the same music piece. Thus, focusing on specific keypoints can significantly improve the performance of our proposed module.

4.4. Attention Visualizations

Keypoint Attention We visualize the learned keypoint attentions by our model for the dance-music alignment task in Figure 7, where a skeleton visualization is presented alongside the histogram of the attention values after softmax. A circle is drawn around each keypoint, and the radii are larger for keypoints with larger attention values. Results show the attention for eyes and limbs (elbows, wrists, knees, ankles) are significantly larger than the attention for the bulk of the body (chest, shoulders, hip). The network has learnt to rely more on the limbs whose movements are more rapid, and less on the bulk parts that seldom moved very fast. We also note that the network assigns more attention to the eyes but very little to head or ears. A possible explanation for this is that the movements of head (like turns) is best represented by tracking two eyes, rather than ears that are sometimes occluded.

Temporal Attention Figure 8 shows the temporal attention on an example audio mel-spectrogram. As can be seen, the attention module tends to focus more on the onsets of a music piece which can be easily understood since dancers always switch motion pattern at these onsets. Thus, focusing more on these temporal regions can help with better alignment.

4.5. Dance and Speech Retargeting

By time-warping the visual beats of existing dance footage into alignment with new music, we can change the song that a performer is dancing to, which is known as dancing retargeting. We generate 7 evaluation pairs for dance retargeting evaluation. Each pair contains a dancing video directly combined with another music piece and the same

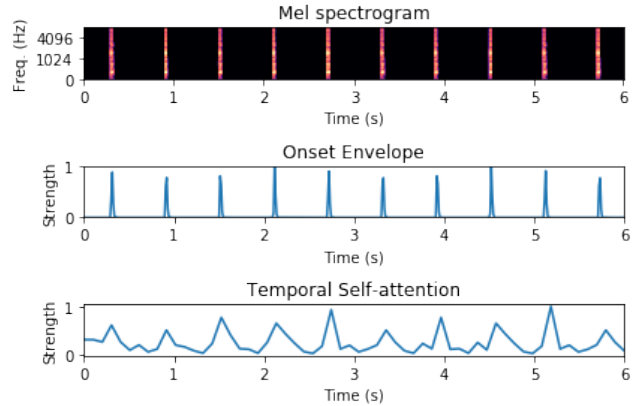


Figure 8: Mel-spectrogram (top), onset envelop (middle), and temporal attention (bottom) on a sample audio. Attention magnitudes agree with audio onsets.

Task	Percentage	
	Direct Combination	Ours
Dance Retargeting	31%	69%
Speech Retargeting	22%	78%

Table 4: Subjective user study for dance and speech retargeting.

video retargeted according to the audio using AlignNet. We ask 13 people for their preference of all the evaluation pairs, and the result is shown in Table 4. Similarly, we generate 7 evaluation pairs for speech retargeting evaluation. It’s worth noticing that in these speech pairs, both videos say the same sentence since it is not reasonable to align talking face with different sentences. As can be seen in Table 4, subjects prefer ours to direct combinations, which means that our method can generalize well to real-world data. However, our method performs slightly worse in dance retargeting. The main reason is most dancing musics have similar beats and direct combination can achieve decent performance. Demo videos of both the retargeting task and the synthetic re-alignment task can be found in the supplementary materials and on our webpage.

5. Conclusion

In conclusion, we proposed AlignNet, an end-to-end model to address the problem of audio-video alignment at arbitrary temporal scale. The approach adopts the following principles: attention, pyramidal processing, warping, and affinity function. AlignNet establishes new state-of-the-art performance on dance-music alignment and speech-lip alignment. We hope our work can inspire more studies in video and audio synchronization under non-uniform and irregular misalignment (*e.g.* dancing).

References

- [1] D. P. Bertsekas, D. P. Bertsekas, D. P. Bertsekas, and D. P. Bertsekas. *Dynamic programming and optimal control*, volume 1. Athena scientific Belmont, MA, 1995.
- [2] H. Bredin and G. Chollet. Audiovisual speech synchrony measure: application to biometrics. *EURASIP Journal on Applied Signal Processing*, 2007(1):179–179, 2007.
- [3] Z. Cao, G. Hidalgo, T. Simon, S.-E. Wei, and Y. Sheikh. OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. In *arXiv preprint arXiv:1812.08008*, 2018.
- [4] H.-K. Chiu, E. Adeli, B. Wang, D.-A. Huang, and J. C. Niebles. Action-agnostic human pose forecasting. *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1423–1432, 2018.
- [5] J. S. Chung, A. Nagrani, and A. Zisserman. Voxceleb2: Deep speaker recognition. In *INTERSPEECH*, 2018.
- [6] J. S. Chung and A. Zisserman. Out of time: automated lip sync in the wild. In *Asian conference on computer vision*, pages 251–263. Springer, 2016.
- [7] S.-W. Chung, J. S. Chung, and H.-G. Kang. Perfect match: Improved cross-modal embeddings for audio-visual synchronisation. *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3965–3969, 2018.
- [8] A. Davis and M. Agrawala. Visual rhythm and beat. *ACM Transactions on Graphics (TOG)*, 37(4):122, 2018.
- [9] F. De la Torre. A least-squares framework for component analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(6):1041–1055, 2012.
- [10] A. Ephrat, I. Mosseri, O. Lang, T. Dekel, K. Wilson, A. Hassidim, W. T. Freeman, and M. Rubinstein. Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation. *arXiv preprint arXiv:1804.03619*, 2018.
- [11] S. Furukawa, T. Kato, P. Savkin, and S. Morishima. Video reshuffling: Automatic video dubbing without prior knowledge. In *ACM SIGGRAPH 2016 Posters*, SIGGRAPH '16, pages 19:1–19:2, New York, NY, USA, 2016. ACM.
- [12] R. Gao and K. Grauman. 2.5 d visual sound. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 324–333, 2019.
- [13] S. Ginosar, A. Bar, G. Kohavi, C. Chan, A. Owens, and J. Malik. Learning individual styles of conversational gesture. In *Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2019.
- [14] T. Halperin, A. Ephrat, and S. Peleg. Dynamic temporal alignment of speech to lips. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3980–3984. IEEE, 2019.
- [15] B. King, P. Smaragdis, and G. J. Mysore. Noise-robust dynamic time warping using plca features. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1973–1976. IEEE, 2012.
- [16] J. Lewis. Automated lip-sync: Background and techniques. *The Journal of Visualization and Computer Animation*, 2(4):118–122, 1991.
- [17] B. LUCAS. An iterative image registration technique with an application to stereo vision. In *Proceedings of 7th International Joint Conference on Artificial Intelligence*, 1981, pages 674–679, 1981.
- [18] Y.-J. Luo, M.-T. Chen, T.-S. Chi, and L. Su. Singing voice correction using canonical time warping. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 156–160. IEEE, 2018.
- [19] A. Nagrani, J. S. Chung, and A. Zisserman. Voxceleb: a large-scale speaker identification dataset. *Telephony*, 3:33–039, 2017.
- [20] A. Owens and A. A. Efros. Audio-visual scene analysis with self-supervised multisensory features. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 631–648, 2018.
- [21] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le. SpecAugment: A simple data augmentation method for automatic speech recognition. *ArXiv*, abs/1904.08779, 2019.
- [22] L. R. Rabiner and B.-H. Juang. *Fundamentals of speech recognition*. Prentice Hall, 1993.
- [23] M. E. Sargin, Y. Yemez, E. Erzin, and A. M. Tekalp. Audio-visual synchronization and fusion using canonical correlation analysis. *IEEE Transactions on Multimedia*, 9(7):1396–1403, 2007.
- [24] D. Sun, X. Yang, M.-Y. Liu, and J. Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [25] G. Trigeorgis, M. A. Nicolaou, B. W. Schuller, and S. Zafeiriou. Deep canonical time warping for simultaneous alignment and representation learning of sequences. *IEEE transactions on pattern analysis and machine intelligence*, 40(5):1128–1138, 2017.
- [26] G. Trigeorgis, M. A. Nicolaou, S. Zafeiriou, and B. W. Schuller. Deep canonical time warping. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5110–5118, 2016.
- [27] D. v. Vloed, J. Bouten, and D. A. van Leeuwen. Nfi-frits: A forensic speaker recognition database and some first experiments. *[SI]: ISCA Speaker and Language Characterization special interest group*, 2014.
- [28] R. H. Woo, A. Park, and T. J. Hazen. The mit mobile device speaker verification corpus: data collection and preliminary experiments. In *2006 IEEE Odyssey-The Speaker and Language Recognition Workshop*, pages 1–6. IEEE, 2006.
- [29] H. Zhao, C. Gan, W.-C. Ma, and A. Torralba. The sound of motions. *arXiv preprint arXiv:1904.05979*, 2019.
- [30] H. Zhao, C. Gan, A. Rouditchenko, C. Vondrick, J. McDermott, and A. Torralba. The sound of pixels. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- [31] F. Zhou and F. Torre. Canonical time warping for alignment of human behavior. In *Advances in neural information processing systems*, pages 2286–2294, 2009.