

SyncUp: Vision-based Practice Support for Synchronized Dancing

ZHONGYI ZHOU, Interactive Intelligent Systems Lab., The University of Tokyo, Japan

ANRAN XU, Interactive Intelligent Systems Lab., The University of Tokyo, Japan

KOJI YATANI, Interactive Intelligent Systems Lab., The University of Tokyo, Japan

The beauty of synchronized dancing lies in the synchronization of body movements among multiple dancers. While dancers utilize camera recordings for their practice, standard video interfaces do not efficiently support their activities of identifying segments where they are not well synchronized. This thus fails to close a tight loop of an iterative practice process (i.e., capturing a practice, reviewing the video, and practicing again). We present SyncUp, a system that provides multiple interactive visualizations to support the practice of synchronized dancing and liberate users from manual inspection of recorded practice videos. By analyzing videos uploaded by users, SyncUp quantifies two aspects of synchronization in dancing: pose similarity among multiple dancers and temporal alignment of their movements. The system then highlights which body parts and which portions of the dance routine require further practice to achieve better synchronization. The results of our system evaluations show that our pose similarity estimation and temporal alignment predictions were correlated well with human ratings. Participants in our qualitative user evaluation expressed the benefits and its potential use of SyncUp, confirming that it would enable quick iterative practice.

CCS Concepts: • **Human-centered computing** → **Interactive systems and tools**.

Additional Key Words and Phrases: Synchronized dancing, practice support, computer vision, visualization

ACM Reference Format:

Zhongyi Zhou, Anran Xu, and Koji Yatani. 2021. SyncUp: Vision-based Practice Support for Synchronized Dancing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 5, 3, Article 143 (September 2021), 25 pages. <https://doi.org/10.1145/3478120>

1 INTRODUCTION

Synchronized dancing consists of a series of synchronized poses and/or temporally aligned movements by multiple dancers that create visual aesthetics. Many amateur dancers engage in such dances by replicating choreographies performed by popular idols and anime characters. They also share their performances via online social network services, often generating a huge popularity spike. It is thus an increasingly common practice for amateur dancers to record their dance performances.

Besides dissemination on social networks, dancers also use these recordings for their practice. They use cameras on smartphones or tablets to capture their practices and share the recordings with other members to review the portions of their dance that require improvements (e.g., moments where poses of the dancers are different or the timing of their motions are not aligned). This is an important activity because it offers a good common ground for dancers to understand what needs to be improved. However, the most common approach for reviewing videos uses a naïve video player with simple navigation functionalities (e.g., on the smartphone). Even experienced

Authors' addresses: Zhongyi Zhou, Interactive Intelligent Systems Lab., The University of Tokyo, Tokyo, Japan, zhongyi@iis-lab.org; Anran Xu, Interactive Intelligent Systems Lab., The University of Tokyo, Tokyo, Japan, anran@iis-lab.org; Koji Yatani, Interactive Intelligent Systems Lab., The University of Tokyo, Tokyo, Japan, koji@iis-lab.org.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2021 Association for Computing Machinery.

2474-9567/2021/9-ART143 \$15.00

<https://doi.org/10.1145/3478120>

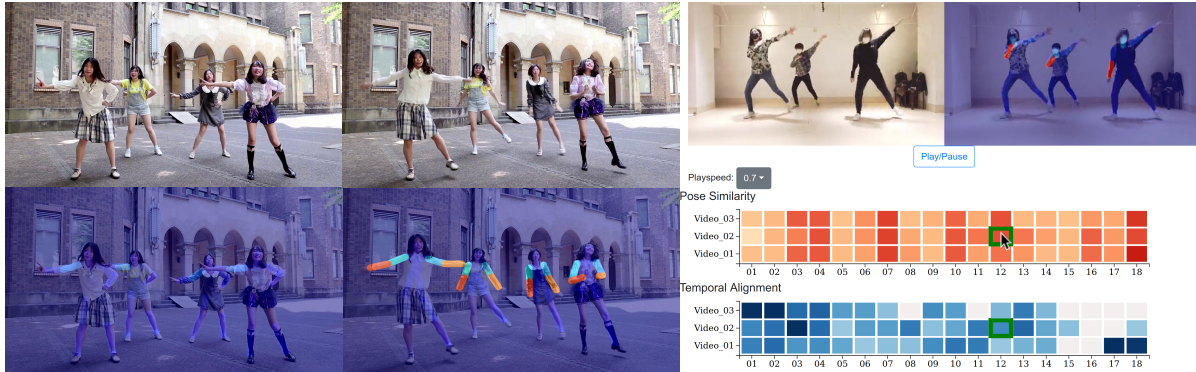


Fig. 1. Left: SyncUp heatmap overlays on dancers highlight which body parts break synchronization in their poses. Colors closer to red and blue represent higher or lower degrees of position discrepancy of each body part, respectively. Right: The SyncUp web interface. After dancers upload their practice videos, the system quantifies the degree of synchronization of their dance and offers feedback using three visualizations (i.e., heatmap overlays on dancers' poses and two 1D heatmaps). This helps users quickly look for and replay a portion of their dances where there is an issue in synchronization.

dancers require time and effort to identify moments where dancers are not well-synchronized. Thus, dancers typically review their recordings after they finish all practices. As a result, practices and their review activities are isolated, which can lead to less satisfactory user experience.

Existing research has examined interactive systems to support dance practices by recognizing dancers' poses and offering correction feedback [1, 13]. These systems typically require specialized hardware and infrastructure, making them less accessible to a general user population [1, 10]. To allow use by dancers without technical skills or ownership of such specialized hardware, a support system should offer quick access to moments where their dances are not well-synchronized while utilizing commodity video recording devices. Recent work [27] has demonstrated a mobile system to offer feedback on dances performed by a single person. However, the system only considers the differences in poses, and lacks the consideration of the temporal alignments of dancers' movements, which is another critical component that creates visual aesthetics seen in synchronized dancing.

We present SyncUp, which provides support for synchronized dance practices using computer vision technology and interactive visualization. The main advantage of SyncUp is to offer quick access to the portions of recordings where the system considers that dances are not well-synchronized in an in-situ manner. Unlike existing work [1, 6, 27], SyncUp considers both pose similarity among dancers and the temporal alignment of motions, another key element of synchronized dancing for inferring the degree of synchronization. SyncUp then visualizes them as 1D heatmaps (Figure 1 right). In this manner, dancers are presented with an overview of the performance evaluation estimated by the system. They can quickly navigate to moments where the degree of synchronization is low, and further investigate how to fix the issues. We envision that SyncUp can mitigate the existing isolation of practices from reviews using recordings.

This work offers the following contributions:

- A formative study that summarizes existing practice procedures and features desired in support systems for synchronized dancing;
- Development of two methods to quantify the degree of synchronization from synchronized dance videos captured with a normal RGB camera in terms of both pose similarity and the temporal alignment of dancers' movements;
- Design of an interactive system for synchronized dancing practice support;

- System evaluations on the accuracy and the robustness of our two quantification methods; and
- A qualitative user evaluation to confirm the benefits of SyncUp and uncover its potential uses in actual practices.

In this paper, we present an exploration of how we designed, built and evaluated SyncUp. We first conducted formative studies to understand current practices and problems in synchronized dance practices and derive desired features for SyncUp. We developed vision-based approaches to identify pose similarity and temporal alignment of motions using videos captured with a normal RGB camera. This paper reports the results of our system and user evaluations, and discusses our findings as well as future directions for improvements.

2 RELATED WORK

2.1 Dance Practice Systems

Prior research uncovered the complexity of dance learning [3, 20], stimulating the development of practice systems for dancers [4, 35]. Many such systems support practice with perception enhancement [6, 25]. Drobny et al. [10] found that a common challenge for dance learners was synchronizing their movements with the underlying rhythm of the music. To address this issue, they created Saltate, which provides beginners with drum-sound feedback at beat onsets. Nakamura et al. [34] employed a movable robot to address the depth ambiguity problem of dance learning videos that only convey 2D information to dancers. The robot moves forward or backward according to motion-depth information. Their research allowed dancers to improve their understanding of their movements in a 3D space. YouMove [1] is an augmented-reality mirror that supports dance learning by visually comparing the poses of a student and instructor at several key dance moments. The system also provides quantitative performance evaluations with detailed notes from the instructor to help the student understand how to improve their dancing.

Existing research has also revealed design guidelines for interactive dance practice support systems [8, 19, 36]. From comparative studies, Trajkova et al. [43] found that the efficacy of feedback modalities (e.g., verbal or visual) in augmented mirror systems depends on the dancer's expertise level. Riviere et al. [37] compared two dance decomposition methods (i.e., led by dancers or instructors) for dance learning. They found that decomposition performed by instructors is appropriate for introductory-level students whereas, decomposition by dancers is preferable for experts.

These projects only used pose information for dance performance evaluations, and the temporal alignment of dance movements were not fully considered. An interactive system that utilizes commodity devices (i.e., smartphones or tablets) and computer vision methods to support synchronized dancing practice can liberate users from specialized equipment (e.g., a depth camera and an augmented-reality mirror), and expand use cases.

2.2 Vision-based Human Motion Analysis

Accurate human pose detection is an active research area in the computer vision field. Researchers have developed real-time robust pose detection methods using depth cameras [24, 40, 41]. Additionally, recent breakthroughs in computer vision [23] have established robust algorithms to locate human skeletons in images, including those captured by a typical RGB camera [5, 16, 46, 48]. The key idea behind these methods is to train convolutional neural networks (CNNs) using large-scale fine-labeled datasets [2, 29]. OpenPose [5], one of the earliest CNN-based real-time multi-person pose estimators, predicts the pixel positions of 18 different human body parts (*keypoints*) in a given frame. Thus, developers can easily construct body skeletons using these predictions. SyncUp employs AlphaPose [16], another well-known multi-person pose estimator, because of its superior performance vs. OpenPose.

2.3 Vision-based Pose Similarity Estimation

Pose similarity estimation plays a critical role in various applications, including action recognition [12, 30], motor skill learning [1, 6], and motion retrieval [38, 42]. One basic pose similarity estimation approach uses the Euclidean distances between body parts in two poses. For example, Chan et al. [6] created a virtual reality dance practice system that employed a simple threshold-based method using the Euclidean distance of each body part between the reference dancer and learner. However, Chen et al. [7] noted that such features were insufficient to evaluate more complex actions. Therefore, they introduced a collection of high-dimensional pose features (1,683 dimensions) to describe human motions and an algorithm to learn the similarity metric based on Mahalanobis distances, which was found robust when using these high-dimensional features. Recently, researchers [39, 47] started to encode the semantics of pose similarities using an implicit model (i.e., a graphic CNN) instead of manually crafted features. However, these methods require a large dataset, and their direct application to synchronized dancing is not yet feasible.

Although pose similarity estimation is a critical component in dance practice support systems, most prior work relies on the summation of the Euclidean distances of body parts. These methods implicitly assume uniform weights on body parts, and their results may not match well with human perception. For instance, the angles of arms may contribute to the perceived pose similarity of dancers more strongly than those of legs. Thus, further explorations are necessary to build a pose similarity quantification method for synchronized dancing.

2.4 Visual Beats

Existing computer vision research has explored methods for intelligent systems to understand the tempo of human motions. One application of these methods is to synthesize dance-like motions [28, 44, 50]. Davis and Agrawala [9] developed the concept of a visual beat, which involves the rhythmic patterns of object movements in a video. Their system synthesizes dance-like motions by matching visual beats with the beats of the background music. Lee et al. [26] developed a system that generates various types of artificial human-dancing videos based on the background music style, such as “Ballet” or “Zumba”. Their method uses a generative adversarial network generator [18] that includes a feature representing the music genre to recurrently create a sequence of video frames.

Our work applies visual beats to synchronized dance performance evaluations. Instead of synthesizing dance-like motions, the primary objective of this work lies in quantifying the temporal alignment of dancer motions. In particular, SyncUp extracts the visual beat features from dancers’ movements and computes how their timings are aligned (i.e., whether a movement is ahead of or behind the reference motion).

3 FORMATIVE STUDIES

To inform our interface design, we conducted focus groups to understand how dancers practice synchronized dancing. Additionally, we conducted an online survey to prioritize features to be implemented in SyncUp.

3.1 Focus Groups

We first held focus groups to understand the current practices of synchronized dancers and derive potential system features. We recruited three dance groups at our university to participate in the focus groups. Each dancer group consisted of 4–10 people who regularly perform synchronized dancing at university events (e.g., school festivals) and external occasions (e.g., dance competitions). We began our focus group by asking ice-breaker questions (e.g., what type of dancing they perform and how often they practice), and then we asked our participants to share their practice approaches and current challenges.

We extracted the common practices and problems observed in our focus groups and derived the following major findings:

Table 1. Online survey results. Features are sorted by their average scores from highest to lowest. Six features (F1, F3, F4, F5, F7 and F8) were implemented in the current SyncUp system. Note that we excluded F2 and F6 in this work because they are already well supported in existing systems.

Feature	Description	Mean (SD)
F1. Pose correction feedback	Provide detailed textual and/or visual instructions on how to fix dancer out-of-sync motions and poses.	4.66 (0.55)
F2. Slow practice mode	Allow dancers to practice a portion of or the entire dance at a slower tempo.	4.44 (1.06)
F3. Mistake summary	Automatically summarize the moments where dancers made mistakes.	4.34 (0.81)
F4. Temporal alignment feedback	Visualize the synchronization degree between dance and music.	4.31 (0.81)
F5. Leader mode	Allow a team leader to record the reference video, and other dancers to compare it to their own dancing. Dancers can also visually compare their dancing and ground truth side by side using two video streams.	4.21 (1.01)
F6. Focus mode	Allow dancers to repeatedly practice a specific portion of the dance.	4.17 (1.14)
F7. Comparison against previous practices	Allow dancers to compare their current dancing to past data.	4.10 (1.08)
F8. Detection of pose differences	Detect and highlight differences in poses among multiple dancers.	3.97 (0.87)
Hand pose tracking	Track dancer hand poses during dancing and provide feedback.	3.65 (1.20)
Dance scoring	Scores represent the degree of pose similarity, motion similarity and facial expressions.	3.55 (1.12)
Detection of differences in motions between poses	Detect and highlight differences in transition between two consecutive poses among multiple dancers.	3.55 (1.12)
Facial expression detection	Detect dancer facial expressions and provide feedback (e.g., “smile more” or “show your face”).	3.38 (1.12)
Detection of motion stability	Detect dancer stability and provide feedback.	3.34 (1.01)

- All groups used video recording (mostly smartphones) for their practices. These videos were useful for finding small but important mistakes in their dancing. In addition, the participants commented that recorded videos allow them to review their dancing from the audience perspective.
- All groups preferred checking their recording only after finishing their entire practices (e.g., after going back home). One group had never checked their videos when they were together because it was time consuming.
- In all groups, dancers first practiced individually; then, they practiced as a group. In this manner, they could focus on improving their synchronization instead of learning movements and poses during group practices.
- All groups typically selected a leader responsible for providing feedback to the other members (i.e., followers) in their group practices. However, they noted that the leader could not always identify all mistakes.
- One group manually created a sheet that summarizes which parts of their dances required additional practice after repeatedly reviewing their videos.

We found that video recording is a common approach for synchronized dancing practice. However, focus group participants tend to rely on the manual inspection of these recorded videos to identify dance segments requiring further practice. These results are in line with our motivation to design an interactive system that supports synchronized dance practices using commodity devices.

3.2 Online Survey

We later conducted an online survey to understand which functionalities are most desired by synchronized dancers. We generated 13 potential features (Table 1) based on the results of our focus groups to address participant concerns regarding the time-consuming aspects of their practices. We conducted an online survey to examine the importance of each feature from the user perspective. The survey asked participants to rate each feature on a 5-Point Likert scale concerning the potential usefulness to their practice (5: they would definitely use the feature; 1: they would never use the feature). We recruited 29 individuals (six males and 23 females) with different dancing engagement levels from casual dancers to professionals.

3.3 Survey Results and Discussion

Table 1 summarizes our results, including the mean rating and standard deviation of each feature. All features were considered useful; the lowest average score was greater than 3.0. However, implementing all proposed features would overwhelm our target users and hinder the core value of our system. In addition, certain features may be error-prone due to the limitations of existing computer vision technology (e.g., hand pose tracking). Furthermore, we decided to exclude features that are already well supported by existing systems (e.g., F2 and F6) and focused on designing quantitative feedback on dance performances. We thus narrowed down the features to those that are critical and sufficiently robust with existing technologies for our prototype. We excluded features with scores of 3.65 or below as there was a substantial gap (0.32) there. This led us to focus on the six features in Table 1 (i.e., F1, F3, F4, F5, F7, and F8).

4 SYNCUP INTERFACE

SyncUp is designed as a web-based interface that can run on laptops, smartphones, and tablets. All features except the spotlight view (explained in Section 4.4) are contained in one web page (F3).


We envision that dancers will capture their multiple practice sessions consecutively at one time, and SyncUp will then offer feedback when they take short breaks between practices. When each practice session finishes, the system will upload the video for analysis while a mobile device captures the next session. In our current implementation, SyncUp requires approximately 85 s to analyze a 1-min video with three dancers at 30 fps and at 360p resolution. In addition, providing feedback in a post-hoc manner (i.e., after finishing a practice session rather than real time during the practice) is in line with the findings by Trajkova et al. [43]. SyncUp highlights segments where synchronization is deemed low by the system, and dancers can discuss whether they want to add practice in corresponding portions of their dance program. In this manner, SyncUp alleviates the isolation of practices from reviews using video recordings, as observed in our formative studies.

4.1 Side-by-Side Video Replay

The main feature in SyncUp is the side-by-side video replay with heatmap overlays. The contents in the video replay view differ in the two practice modes offered by SyncUp: group and individual practice modes.

4.1.1 Group Practice Mode. The group practice mode is utilized when dancers are physically together practicing a synchronized dance. In the group practice mode, SyncUp offers a side-by-side view of the original practice video and the one with heatmap overlays that visualize the degree of similarity of dancer poses at a particular time (F1). In the example of Figure 2a, the heatmap overlays highlight dancers' arms in red, suggesting that this body part is not well synchronized. This visualization helps dancers quickly identify which body parts require adjustment in future practices.


4.1.2 Individual Practice Mode. The individual practice mode is used when dancers wish to practice separately. Our focus groups suggested that this is common during the early stages of practices. In this practice mode, one



(a) Side-by-side video replay with heatmap overlays in the group practice mode.


(b) Side-by-side video replay with heatmap overlays in the individual practice mode.

Pose Similarity



(c) Pose similarity visualization. Dance segments with low pose similarity are indicated by a dark red color.

Temporal Alignment



(d) Temporal alignment visualization. Dance segments where dancers' movements are not aligned are represented by a dark blue color.

Fig. 2. Interactive visualizations in SyncUp. SyncUp presents a side-by-side video (either 2a or 2b) at the top of the interface. The interface also includes 1D heatmaps of multiple practice videos. The values in the horizontal axis in both 2c and 2d show the indices of video segments (8 beats long). Green rectangles highlight the corresponding segment played in the side-by-side video replay view. Both 2c and 2d show the analysis results for 2a.

dancer is typically responsible for teaching the others. In this paper, we refer to this dancer as the “leader” and the others as “followers”. The side-by-side video replay shows the video of the leader on the left and that of the follower with heatmap overlays on the right (Figure 2b). Alignment between the two videos is performed through the background music alignment [14]. This mode allows the leader to record their dances as the reference data to support followers' practices (F5).

4.2 Pose Similarity Visualization

Pose similarity quantifies how well the poses of multiple dancers are synchronized in a particular video frame. In addition to the heatmap overlays, SyncUp also offers a 1D heatmap to visualize the average score of pose similarity (F8) in each segment of the dance video derived by our algorithm (explained later). A darker red color indicates a lower degree of synchronization. We define the length of a segment as 8 beats long because it is commonly used in dance practices. Multiple rows of 1D heatmaps present a comparison of the synchronization scores among multiple practice videos. This helps dancers review how their pose similarity improves over multiple rounds of practice (F7, Figure 1 right).

When the user clicks a block on the heatmap, it activates a replay of the corresponding segment in the side-by-side video replay at a play speed specified by the user. In addition, when a user replays the video, the green highlight in the heatmaps is also updated to inform the location of the corresponding segment.

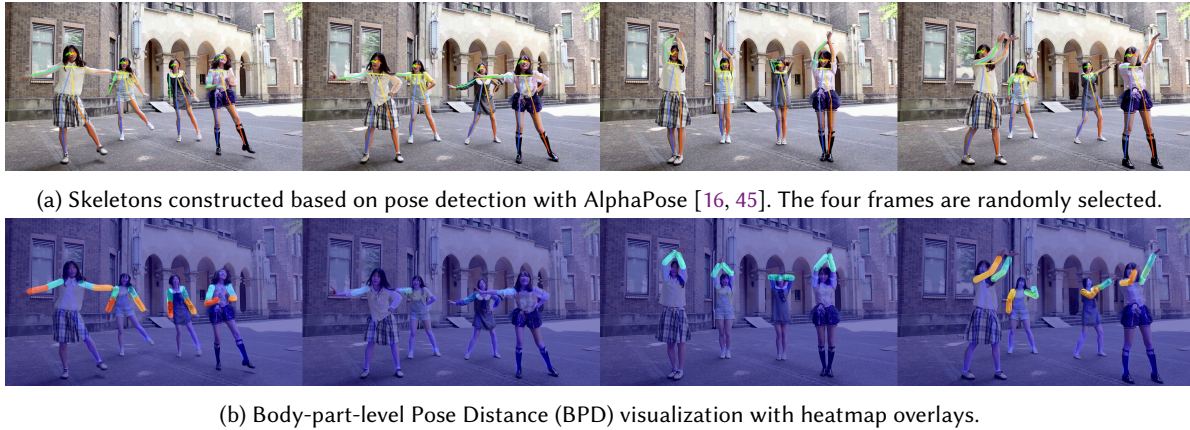


Fig. 3. Example results of our pose similarity quantification method.

4.3 Temporal Alignment Visualization

Pose similarity considers the degree of synchronization at a video frame level. Another aspect that contributes to synchronization is the alignment of the movement timing. For example, if one dancer raises his/her right hand at the same time that another dancer raises his/her left leg, the poses of two dancers differ, but the timing of their motions is synchronized. We thus need to also consider whether dancers' motions are aligned in the temporal domain.

It is challenging to directly detect how the movements of the dancers are aligned with the music without human annotations. Such annotations are very tedious. Our system therefore determines how the timing of the followers' movements (i.e., the two dancers on the side of Figure 2a) is synchronized with that of the leader (i.e., the central dancer in Figure 2a). Dancers identify the leader by choosing a skeleton from the menu. SyncUp then visualizes the degree of the temporal alignment of their movements in a 1D heatmap, as shown in Figure 2d (F4). A darker blue color suggests a larger difference in terms of the movement timing. Similar to the pose similarity visualization, users can click a block of the heatmap to navigate to the corresponding segment.

4.4 Spotlight View

In addition to the default view shown in Figure 2, SyncUp offers another interface: the spotlight view. This view offers a collection of segments that are sorted by the synchronization score calculated from the pose similarity and temporal alignment in ascending order. The segment atop this view is the portion of the video in which the system considers that a dance had the worst synchronization. This view is intended to offer direct access to moments where additional practices may be necessary by liberating users from inspecting with the full-length footage.

5 DANCE PERFORMANCE QUANTIFICATION

SyncUp executes two kinds of analyses to quantify the degree of synchronization in dances: pose similarity and temporal alignment. Pose similarity analysis examines how the poses of multiple dancers are synchronized with each other in a given frame. Temporal alignment analysis quantifies the timing differences of motions among multiple dancers, indicating how well synchronized a dancer's motion is with others at a specific time interval. Our interface ultimately visualizes the results of these analyses, as shown in Figure 2c.

5.1 Video Pre-processing

5.1.1 Video Segmentation by Music Beats. Counting eight beats in a loop is a common dance practice technique. Referring to this technique, SyncUp divides a practice video into a series of segments, each of which is 8 beats long. We use Ellis’s method [15] to estimate the tempo of the music and segment the video.

5.1.2 Dancer Pose Detection. We use AlphaPose [16], a multi-person pose estimator, for identifying dancers’ poses. In AlphaPose, a person’s pose is represented by a list of pixel positions having 18 keypoints (p) [5, 29], similar to OpenPose [5]. These keypoints are the 2D pixel locations of a person’s eyes, ears, nose, neck, shoulders, elbows, wrists, hips, knees, and ankles. Figure 3a displays the overlays of skeletons created from the AlphaPose keypoints on example recorded dancing videos. We did not formally experiment with the accuracy of pose detection because it is out of the scope of this work. However, we found that AlphaPose successfully detected dancers’ poses except in the case where a person was largely occluded.

5.1.3 Dancer Tracking. After pose detection, the system performs Xiu et al.’s method [45], which tracks each dancer between frames. We hypothesize that the poses of the same dancer at two consecutive frames will be very similar. Hence, the body part locations will also be similar between the frames. The system first calculates the distances of all skeletons between frames. We define the skeleton distance between frames as the total distances of the 18 keypoints. The system then finds the combination that minimizes the overall distances, resulting in the mapping of the dancers between frames.

In rare cases, AlphaPose fails to identify the same number of dancers as in the previous frame (e.g., three dancers in the previous frame, but two in the current frame). The system still performs the same procedure, and it carries over the skeleton information of the unmatched dancer(s) to the current frame (i.e., the system regards that these dancers did not move at all between the two frames).

5.2 Pose Similarity Analysis

The algorithm first computes the degree of pose similarity among dancers at the body part level (*Body-part-level Pose Distance*, or BPD in short), which the system uses to display the heatmap overlay (Figure 2a & 2b). The system also computes a metric called *Overall Pose Similarity* (OPS), which describes the overall similarity of poses at a given frame with the BPD values of all dancers. SyncUp then averages the OPS values across all frames in each segment and presents the result as 1D heatmaps (Figure 2c).

5.2.1 Body-part-level Pose Distance. SyncUp uses 14 keypoints (all except eyes and ears) to compute 13 body-part feature vectors. We exclude the four features of eyes and ears because we wanted to focus on body poses rather than face directions. We normalize the feature vectors into their corresponding unit vectors to preserve only the directional information of each body part. This avoids a scale-difference problem caused by different human heights and camera distances. We use these 13 unit vectors as the source input to our algorithm to calculate the BPD value ($BPD(i, t)$) for each body part i at a given frame t .

The system calculates the accumulated absolute difference of each body part across all dancers ($d(i, t)$) according to the formula:

$$d(i, t) = \sum_j^J |\vec{v}_j(i, t) - \vec{v}_R(i, t)| \quad (1)$$

where $\vec{v}_j(i, t)$ is the i -th unit vector ($i \in \{1, 2, \dots, 13\}$) of the j -th person ($j \in \{1, 2, \dots, J\}$), and $\vec{v}_R(i, t)$ is the i -th unit reference vector, equal to the average of the unit vectors of all dancers (i.e., $\vec{v}_R(i, t) = \frac{1}{J} \sum_j \vec{v}_j(i, t)$).




Fig. 4. Heatmap overlay results with different λ values.

Note that $d(i, t)$ is an unbound variable that is strongly influenced by the number of people in the analysis. We thus use the following formula for normalization to derive BPD :

$$BPD(i, t) = \left(\frac{d(i, t)}{J} \right)^\lambda = \left(\frac{\sum_J |\vec{v}_j(i, t) - \vec{v}_R(i, t)|}{J} \right)^\lambda \quad (2)$$

A higher value of $BPD(i, t)$ means a larger discrepancy in the poses of the i -th body part across dancers, implying low synchronization performance. In our heatmap overlays, we linearly convert the value of $BPD(i, t)$ to a color in the spectrum implemented in the COLORMAP_JET class¹.

λ is a trade-off parameter used to control the sensitivity of the heatmap overlays to the discrepancy of poses. Figure 4 shows how this parameter can affect the heatmap overlays of SyncUp. A small λ highlights small differences which advanced dancers may use to achieve perfect synchronization. A large λ is useful to suppress such minor discrepancies and only highlight major differences.

5.2.2 Overall Pose Similarity. Given the $BPD(i, t)$ values for all body parts in a frame, we then compute an *Overall Pose Similarity* ($OPS(t)$) value. We used a Support Vector Regression (SVR) model [11] as our study found that it was the most robust (see Section 6.3). In our current implementation, our model takes all 13 $BPD(i, t)$ values as input and predicts an $OPS(t)$ value that ranges from 0 (totally ill-synchronized) to 1 (perfectly synchronized). The system uses the predicted $OPS(t)$ values for the 1D heatmap (Figure 2c).

5.3 Temporal Alignment Analysis

The temporal alignment of movements is another aspect of synchronization considered in SyncUp. We extend Davis and Agrawala's [9] notion of directogram to quantify motion rhythms. This notion is similar to a spectrogram, which factors volume changes into frequencies, except that a directogram factors motions of a video into different angles. We further modify their method to compute directogram-equivalent metrics for body parts of multiple dancers.

We first define the *pose flow* of one dancer $\vec{f}_j(t)$, which describes the movement of dancer j at frame t :

$$\vec{f}_j(t) = \{p_j(i, t+1) - p_j(i, t)\}, i \in \{1, 2, \dots, 18\} \quad (3)$$

where $p_j(i, t)$ represents the i -th AlphaPose keypoint at frame t . We use all 18 AlphaPose keypoints to compute $\vec{f}_j(t)$ in our current implementation.

We then compute a *posegram* ($\mathcal{P}(t, \theta)$) for each dancer in a similar manner to the directogram:

$$\mathcal{P}_j(t, \theta) = \sum_i |\vec{f}_j(t)| \mathbb{1}_\theta(\angle \vec{f}_j(t)), \quad \mathbb{1}_\theta(\phi) := \begin{cases} 1, & \text{if } |\theta - \phi| \leq \frac{2\pi}{N_{bins}} \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

¹<https://docs.opencv.org/2.4/modules/contrib/doc/facerec/colormaps.html>

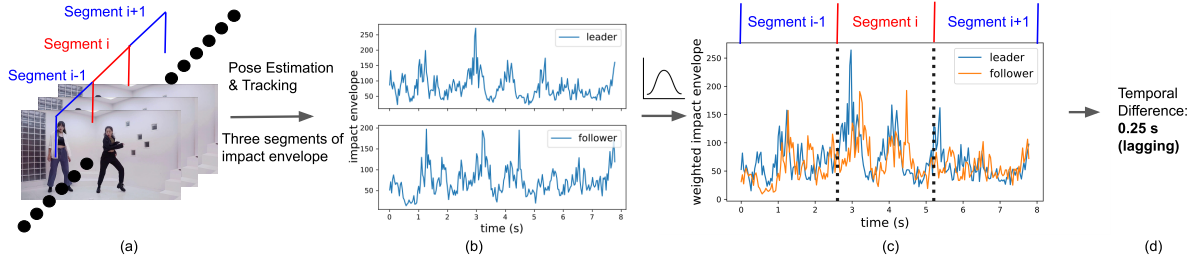


Fig. 5. The Workflow of our temporal alignment analysis. (a)->(b): SyncUp first computes human poses and performs dancer tracking. (b)->(c): To analyze the temporal alignment at the video segment s , SyncUp samples impact envelopes (Equ. 5) at the segment $s - 1$, s , and $s + 1$. The system then applies a Gaussian weight distribution on impact envelopes over these three segments to weigh the contributions by the segment s . The system calculates the impact envelopes for the leader dancer and one follower. (c)->(d): The system then computes the cross-correlation between the two envelopes. In this example, the calculation result is 0.25 second, suggesting that the follower's movements are behind those of the leader.

The *posegram* factors body-part motions of a dancer into N_{bins} different angles. For example, when a dancer stretches only the right hand horizontally, the *posegram* will show a large value in the corresponding direction while the other values are nearly zero.

Similar to Davis and Agarawala's work, we further define a flux $\mathcal{P}_{F,j}(t, \theta)$ and an impact envelope $u_j(t)$ of a *posegram* as follows:

$$\mathcal{P}_{F,j}(t, \theta) = \mathcal{P}_j(t, \theta) - \mathcal{P}_j(t - 1, \theta), \quad u_j(t) = \sum_{\theta=\theta_i}^{\theta_{N_{bins}}} |\mathcal{P}_{F,j}(t, \theta)| \quad (5)$$

The flux is the temporal differentiation of consecutive *posegrams*, which is analogous to acceleration. The impact envelope sums up all fluxes within a dancer along the angular bin axes. This highlights how much one person's motion changes at one time. Figure 5b shows an example of impact envelopes of two different dancers.

The impact envelopes identifies the temporal differences of visual beats between two dancers. We use three consecutive segments to secure a sufficient number of samples for reliable calculation of cross-correlation. We also apply Gaussian weights to impact envelopes in three consecutive segments to highlight the degree of the temporal alignment at the center segment. The Gaussian weights also help reduce artificial noise caused by the cutoff of the impact envelopes. Cross-correlation analysis identifies how much time shift will offer the best match between the two impact envelopes. The time shift exhibiting the highest cross-correlation value is considered as the degree of the temporal alignment against the reference dance, denoted as $\tau_j(s)$ (s is a segment). SyncUp uses the accumulation of the absolute values of the temporal alignment scores calculated between all leader-follower pairs (i.e., $\tau(s) = \sum_j |\tau_j(s)|$) in the 1D heatmap visualization (Figure 2d).

Figure 5 summarizes the workflow of how SyncUp computes the temporal alignment at each segment. Figure 5c shows two examples of impact envelope plots after the Gaussian weights are applied. The dotted lines are the boundaries of the segments. By looking at the peaks and valleys of the two plots, the figure suggests that the follower's dance (orange line) is behind the leader's dance (blue line). $\tau_j(s)$ is 0.25 s in this example, confirming the delay.

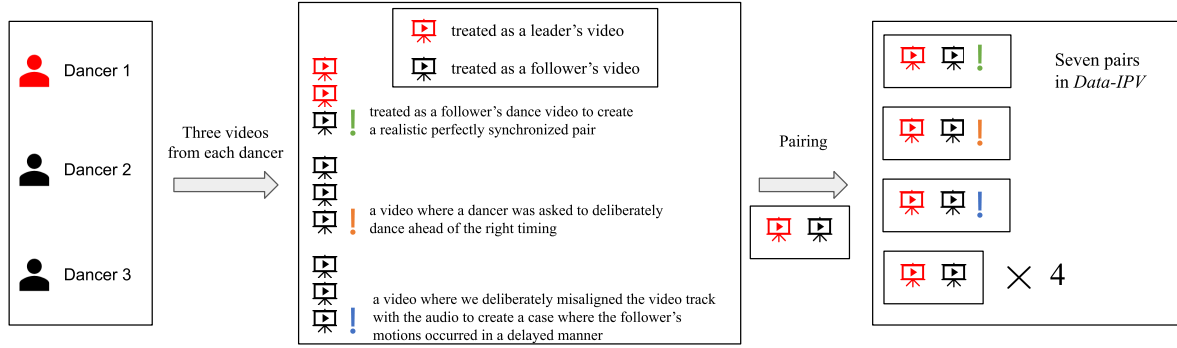


Fig. 6. Composition of *Data-IPV*. Each dancer in one group provided 3 practice videos, resulting 9 videos in total. Two of the 3 leader's videos were then treated as the references (🔴). To cover different degrees of synchronization, we deliberately created pairs to represent extreme cases where synchronization was close to or far from perfection (i.e., pairs with 🟡 with exclamation marks). In this manner, we created seven pairs from each dance group.

6 SYSTEM EVALUATION

We conducted system evaluations on the two dance performance quantification methods. For our evaluations and the training of the neural networks for pose similarity analysis, we created a collection of dance videos and the corresponding labels of the degree of synchronization.

6.1 Dance Video Data Collection

Because we have two practice modes (individual and group), we collected two types of dance videos for our evaluations as well as training of learning-based approaches.

6.1.1 Individual Practice Video Pairs (*Data-IPV*). We collected videos where individual dancers practiced separately. We recruited one volunteer group, consisting of three dancers, for this part of data collection. All of them practiced synchronized dancing that Asian pop stars or anime characters originally performed. We asked each volunteer to perform practices of three dance routines individually and capture their practices using their smartphones or webcams. This led to nine videos from one group, and we used them to create seven pairs.

Figure 6 shows how we created these pairs. Our objective was to create a wide range of in- and out-of-synchronization samples from the limited number of dance data collected. To create a realistic perfectly-synchronized pair, we paired two videos from the leader (the pair that included 🟢 in Figure 6). When collecting the dancing data, we asked one follower to deliberately dance ahead of the correct timing (🟡 in Figure 6). By pairing this with the reference video, we created a pair that represented a case where the follower's motions occurred ahead of those of the leader. Furthermore, to create a pair that represents a case where the follower's motions occurred in a delayed manner, we deliberately introduced a delay of 100 ms only to the video footage (but not the audio track) in another follower's practice videos (🟡 in Figure 6). We confirmed that this caused sufficient lagging effects throughout the video for our ML approaches. These three pairs were created to cover extreme cases where synchronization was close to or far from perfection. We then paired the remaining four videos (🔴 without the exclamation mark) to the reference videos (🔴). As a result, we created seven pairs of videos from each dance group.

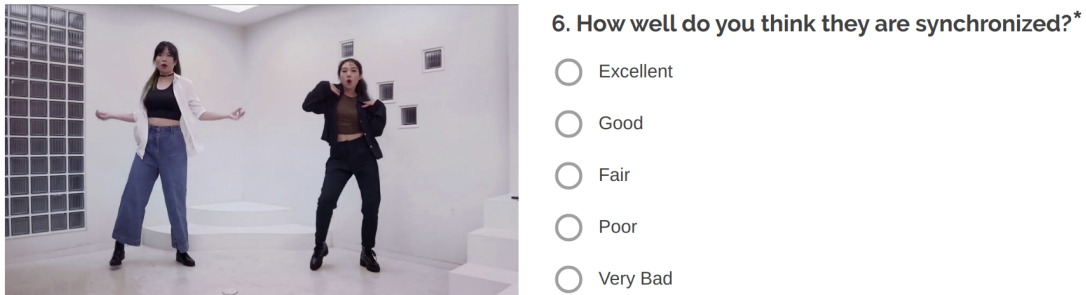


Fig. 7. A web interface was used for collecting human ratings on pose similarity. We used a similar system for human assessment collection on the temporal alignment evaluation.

6.1.2 Group Practice Videos (Data-GPV). To evaluate our dance performance quantification methods in the scenario of group practices, we also gathered videos where multiple dancers danced together. We recruited another four dance groups (2–4 dancers in each group) and received nine videos of group practices from them in total. Their dance style was similar to that of the individual practice video data.

6.2 Dance Performance Assessments with Human Raters

6.2.1 Pose Similarity. Because the pose similarity algorithm uses frames of a dance video, we created a set of frames and asked our volunteers to rate how well the poses of the two dancers were synchronized. We randomly-sampled 100 frames from each of the two data sources (i.e., *Data-IPV* and *Data-GPV*), resulting in 200 frames in total. The number of random samples chosen from each practice video was set to be proportional to its video length. We also confirmed that the seed for each data source was set to be different in our random sampling so that we did not include frames of identical dance portions in both training and testing data. We then created an online form where our volunteers shared their perceived pose similarity of the two dancers with according to a 5-Point Likert scale (“Excellent”, “Good”, “Fair”, “Poor”, and “Very Bad”), which was internally mapped to numerical scores of 1, 0.75, 0.5, 0.25, and 0, respectively.

We recruited 161 volunteers (98 males, 60 females, and 3 prefer not to say) for this part of the rating collection. Each volunteer was assigned to randomly selected 25 frames. As a result, we obtained 20 ratings on average for each frame. We removed outliers that were beyond three standard deviations from the mean rating of each sample, which is a common heuristic for outlier removal. Our algorithm regards the mean of the remaining ratings as the overall human rating of the pose similarity. We treated this overall human rating as a continuous value and used it as the training label for our following learning-based regression models.

6.2.2 Temporal Alignment. For evaluating the temporal alignment algorithm, we extracted 20 frames (randomly-selected 10 scenes from each of *Data-IPV* and *Data-GPV*). We then created another online form similar to that for pose similarity for collecting ratings from our volunteers. To clarify our focus in this part of the rating collection, we explicitly instructed our volunteers to judge only the temporal alignment of two dancers in the given frame. We deliberately slowed down the playback speed of each frame (70% of original) so that the volunteers could closely examine the temporal alignment. We also clarified which dancer was the reference in the instruction. The volunteers were asked to rate their perceived temporal alignment according to a 5-Point Likert scale (“A lot faster”, “A little faster”, “At the same time”, “A little slower”, and “A lot slower”), which corresponded to the numerical scores of 1, 0.5, 0, -0.5, and -1, respectively.

We recruited 15 volunteers (eight males and seven females) for this rating collection. All volunteers rated the whole set of 20 frames. We employed the same outlier exclusion as in the pose similarity, but we did not find any

Table 2. Comparison of OPS prediction results. The results are shown in the format of “RMSE (Pearson’s r value)”. All p values for Pearson’s r values were smaller than .05.

Method	λ									
	0.333	0.426	0.543	0.693	0.885	1.13	1.44	1.84	2.35	3
Simple Addition	0.168 (0.65)	0.170 (0.67)	0.174 (0.69)	0.181 (0.71)	0.191 (0.73)	0.205 (0.72)	0.237 (0.69)	0.278 (0.64)	0.317 (0.64)	0.348 (0.59)
SVR	0.163 (0.63)	0.159 (0.65)	0.156 (0.66)	0.154 (0.68)	0.151 (0.69)	0.153 (0.69)	0.158 (0.66)	0.161 (0.64)	0.169 (0.59)	0.180 (0.51)
Short_NN	0.204 (0.23)	0.216 (0.14)	0.199 (0.34)	0.174 (0.55)	0.165 (0.61)	0.164 (0.62)	0.170 (0.58)	0.177 (0.55)	0.184 (0.49)	0.190 (0.45)
Long_NN	0.172 (0.57)	0.173 (0.60)	0.166 (0.61)	0.167 (0.61)	0.167 (0.61)	0.165 (0.62)	0.170 (0.61)	0.170 (0.60)	0.175 (0.58)	0.176 (0.54)

outlier. Therefore, we directly used the mean value of the ratings for each frame as the ground truth temporal alignment score.

6.3 Comparison for Pose Similarity

To evaluate both the accuracy and robustness of our prediction of *OPS*, we experimented four different methods, including the SVR-based approach used in our current SyncUp implementation.

6.3.1 Simple Addition. The most straightforward approach to computing $OPS(t)$ uses the naïve sum of all $BPD(i, t)$ values at every body part. We first normalized the sum of $BPD(i, t)$ values into $[0, 1]$ (denoted as $S_{BPD}(i, t)$). To match our human ratings, we applied the reverse conversion on $S_{BPD}(i, t)$ (i.e., $OPS(t) = 1 - S_{BPD}(i, t)$). The method based on the simple addition is commonly used in existing research [7, 49]. One underlying assumption is that all features (in our case, *BPD*) should weigh equally for the final overall score. This assumption does not necessarily fit human perception of pose similarity [7, 33]. We used this as a baseline method in our evaluation.

6.3.2 Neural Network. Neural networks (NNs) are often appropriate for complicated perceptual tasks, such as object detection and face recognition [17, 23]. We included two types of NNs in our performance comparison: short-NN and long-NN. Short-NN consisted of two layers: an input layer (i.e., *BPS* with 13 dimensions) and an output layer (*OPS*, one dimension). Long-NN consisted of four layers: input, two hidden (10-dimension and 5-dimension) and output. More complicated network structures may outperform, but we decided to experiment rather simple networks due to the limited number of samples for training. We used *LeakyReLU* as the activation function after each linear transformation. For training, we used the Root Mean Squared Error (RMSE) loss function and Adam [21] as the optimizer with the learning rate of 0.01. The whole training session took 50 epochs.

6.4 Results

6.4.1 Pose Similarity. We performed predictions using the four *OPS* computation methods (i.e., Simple Addition, SVR, Short-NN, and Long-NN) with different λ values. We chose 10 values uniformly distributed between 0.333 and 3 in a logarithmic scale. An overly low or high λ was excluded because it was not practical (i.e., too many differences or none were highlighted in the visualization, respectively). We employed cross validation where one of the 16 data sources used for creating static frames (Section 6.2.1) was reserved for testing, and the remaining data were used for training to evaluate our SVR and NN-based approaches. This led to 16-fold cross validation (seven pairs from *Data-IPV* and nine dances from *Data-GPV*).

Table 3. The prediction results with *Data-GPV* and *Data-IPV*. The results are shown in the format of “RMSE (Pearson’s r value)”. We highlight the lowest RMSE in each dataset condition. All p values were smaller than .001.

	<i>Data-GPV</i>	<i>Data-IPV</i>	Both
Addition	0.193 (0.54)	0.189 (0.88)	0.191 (0.73)
SVR	0.174 (0.45)	0.124 (0.87)	0.151 (0.69)
Short_NN	0.180 (0.34)	0.149 (0.80)	0.165 (0.61)
Long_NN	0.193 (0.33)	0.137 (0.82)	0.167 (0.61)

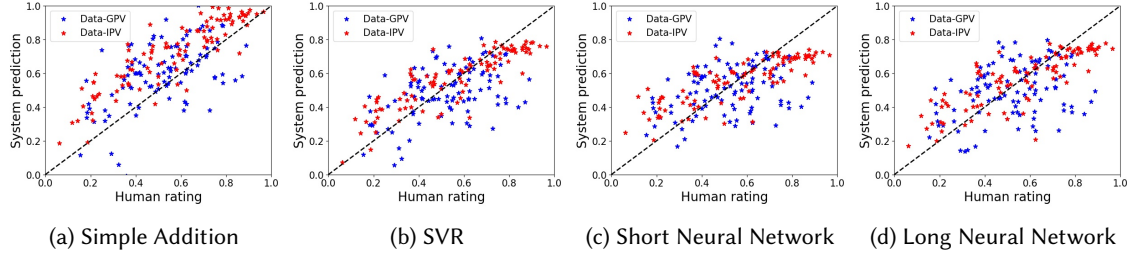


Fig. 8. Scatter plots of system predictions and average human ratings for the four OPS computation methods. λ was set to 0.885 in these results. The black dotted line represents the perfect prediction.

Table 2 shows the RMSE and Pearson’s r values between the human rating averages and predictions by the four methods under all λ conditions. The p values for all r were smaller than .05. With all λ values, we found that the SVR-based method demonstrated the smallest RMSE. This suggests that the SVR-based method was the most accurate and robust in our examination. We also found that correlation values of Simple Addition and SVR were close while the values of the other two methods were lower. Figure 8 further visualizes the detailed results of all 200 samples with $\lambda = 0.885$. The plots compare their human ratings (x-axis) and predictions (y-axis). The red and blue dots represent the results of *Data-IPV* and *Data-GPV*, respectively. We found that the red samples tended to be positioned more closely to the black dotted line (perfect prediction). Table 3 summarizes the accuracy of different methods using *Data-GPV*, *Data-IPV*, and both under the condition of $\lambda = 0.885$. The results show that the overall accuracy in *Data-IPV* was higher than that in *Data-GPV*.

6.4.2 Temporal Alignment. To calculate the temporal alignment with our algorithm, we extracted the previous and next segments of each segment given to the raters from the original video data. Figure 9 presents the comparison of the temporal alignment results between human ratings and our prediction. Note that human ratings were collected according to a 5-Point Likert scale (from -1 to 1), whereas our prediction was given using the estimated time difference (ms). There are seven samples within a region of $(\pm 0.2, \pm 20$ [ms]).

The Pearson’s r was .712 ($p < .001$), showing a strong positive correlation. In the two cases where raters unanimously marked “at the same time (a score of zero)”, our predictions were both zero. There are three notable cases where our predicted scores were quite far from the human ratings (highlighted with red squares in Figure 9). Excluding the three highlighted cases, the other 17 samples in Figure 9 shows that SyncUp’s temporal alignment predictions were highly correlated with human ratings (Pearson’s $r = .890$ ($p < .001$)).

6.5 Discussion

6.5.1 Pose Similarity. The evaluation results of pose similarity demonstrate a smaller RMSE by our prediction (0.151) than human raters (0.223). The results suggest that our quantification method can predict a closer score of




Fig. 9. A scatter plot for the human ratings and SyncUp temporal alignment predictions. We highlight three cases with a red rectangle where our predictions were largely different from the ground truth. The number annotations are included for the discussions in the paper.




Fig. 10. A scatter plot of the standard deviations vs. human ratings and absolute prediction errors. The Pearson's r value was $-.08$ ($p = .26$).

pose similarity to the ground-truth data (i.e., the average of human ratings) than raters. However, it does not necessarily mean that our method is more accurate than human raters in general. Quantifying pose similarity is an inherently challenging task as even human raters demonstrate certain degrees of disagreement. In one case, the standard deviation of human ratings was as large as .378. Human raters may have weighed different aspects of dancing, potentially leading to a diverse set of ratings. Future work should investigate what caused the large disagreements in some cases and how algorithms might account for such disagreement.

The Pearson's r values in the Simple Addition method were slightly higher than those in the SVR-based method. This implies that the data distribution in the Simple Addition method may demonstrate linearity alongside a different slope and intercept. Learning-based models can be expected to accommodate such biases, and the SVR-based method and NNs offer linear and non-linear mapping, respectively.

Our results show that the SVR-based method performed well over NNs. Neural network methods usually perform well when they are applied directly to raw data, (e.g., images and texts). They can extract high-level semantic features within these data. However, in our cases, the input of our ML models was a set of manually crafted feature vectors because we were not able to have a sufficient number of labeled data for training NNs with raw data. Furthermore, appropriate pre-trained models were not available for our purpose. This configuration may have impacted the performance of the two methods using NNs. Another possible explanation for the results in favor of the SVR-based method is that the intrinsic function that maps the input vectors to expected output values was not highly complex unlike image recognition and text translation. Future work is encouraged to further examine the capability of various machine learning approaches, but based on our RMSE results across different λ values (Table 2), we decided to employ the SVR-based method in our current SyncUp prototype.

One major cause of the lower accuracy in *Data-GPV* was the occlusion of body parts (refer to Appendix B for more details). Similar to other computer vision methods, occlusion leads to less reliable pose detections. We closely examined all the segments in *Data-GPV* and found that 14 included occlusions of a dancer. When excluding these 14 samples, we found the RMSE of *Data-GPV* improved to 0.174 from 0.146.

We examined the relationship between the standard deviation of human ratings and prediction errors. We expected that prediction would be less precise in cases where human ratings diverged. Figure 10 presents the results vs. our SVR-based approach with $\lambda = 0.885$ (showing the best performance in Table 2). The Pearson's r value was $-.08$ ($p = .26$). Therefore, we were not able to confirm a clear connection between the degree of rater agreement and prediction accuracy. This result reflects the challenging nature of predicting pose similarity using computer vision approaches, though our SVR-based method performs well in general.

6.5.2 Temporal Alignment. Our results showed that the predictions of temporal alignments correlated well with human ratings. We closely examined the three cases where we observed large errors. In the following section, we discuss these cases in detail and each case corresponds to annotations in Figure 9.

- Case 1. In this case, the average human rating was $.87$, suggesting that our raters considered that the movements of a follower in the video were largely delayed. However, the SyncUp prediction was 0 [ms]. By inspecting the corresponding video, we found that the follower skipped some movements to catch up with the lead dancer. Although our method does not consider such skips, the human raters took them into account for the rating, resulting in a large discrepancy in the temporal alignment score. However, our pose similarity analysis can identify such cases, and thus this observation confirms that SyncUp should present both pose similarity and temporal alignment results to users.
- Case 2. In this case, the average human rating was $-.10$, but the SyncUp prediction was -134 [ms]. We also confirmed that the dancers were well synchronized in the given segment. However, the follower dancer was largely ahead of the leader in the two adjacent segments which were not shown to our raters. We thus concluded that the performance in these two adjacent segments affected our temporal alignment prediction.
- Case 3. In this case, the average human rating was $-.53$, but the SyncUp prediction was 0 [ms]. In the corresponding video, we found that the poses of the two dancers were largely different although the timing of their movements was the same. We contacted the dancers who provided this video footage and they agreed that they were synchronized in terms of the temporal alignment by commenting: “*We truly look different, but the movements are on-beat together.*” We thus concluded that our prediction was correct, but the large pose difference between the dancers caused a bias among the human raters. This case suggests a benefit of SyncUp, in which the contributions of poses and movement timing can be distinguished to accurately identify segments that exhibit a low degree of synchronization.

7 QUALITATIVE USER EVALUATION

We next conducted a user study to understand the potential benefits of SyncUp in actual use. Due to the spread of COVID-19, we were unable to conduct an in-person user study. We instead conducted an online qualitative study to obtain feedback from our target users.

We recruited three new dance groups (nine dancers in total, eight females and one male, ages: 24–32). All dancers were amateur synchronized dancers practicing k-pop dances. We asked each group to perform the same dance routine (65–78 s long) 12 times and to share the recordings with us. We then invited the leader dancers of these three groups (DL 1–3) for our qualitative study. We first presented our SyncUp interface (including the spotlight view) alongside the videos they shared with us. After explaining the features, they were invited to use our system. To encourage active interaction, we asked them to perform tasks of identifying segments where dancers were not well synchronized (ill-synchronized segments) with our system as well as a standard Web video player interface. We also asked them to experience the spotlight view and perform the same task. After they agreed that they understood the functionalities of SyncUp, we conducted a semi-structured interview that included questions about how the system could be useful in their actual practices and how it could change

interactions and communications among group members. After the study, each participant, including those who did not attend our interviews, was offered approximately USD 70 in their local currency as compensation.

We also recruited another six dancer pairs (23–35 years old, and 11 female and one male) only for the interview part. The experimental procedure was the same, except that we used the videos offered by the three groups prior this interview. Hence, we reduced the participants' burden for participation. We recruited the participants as pairs so that they had opportunities to discuss their thoughts with each other during our interviews. We refer to the dancer pairs as DP 1–6. They were offered approximately 40 USD in a local currency as compensation at the completion of the study.

We transcribed all interviews and analyzed the information by categorizing quotes using the open coding approach [32]. Open coding is a qualitative analysis approach by categorizing observed quotes that describe the same or similar phenomena and creating names (or codes) that represent the corresponding categories. Because none of our participants (leaders or pairs) was fluent in English, we conducted interviews in their spoken language. We translated the quotes as faithfully as possible for the report in this paper.

7.1 Results

7.1.1 Enabling Quick Access to Ill-Synchronized Segments. All participants appreciated the design of visualizations in SyncUp. DP2 and DP6 explicitly rementioned the inefficiency of dragging the seek bar of a video player for locating ill-synchronized segments. On the other hand, they saw clear values in the visualizations offered by SyncUp.

“These (1D heatmaps) are intuitive. It helps me quickly locate these [ill-synchronized] segments, and I do not need to watch all the videos.” [DP1]

“This heatmap [overlays] is very clear. I can easily know where the problems are. Very useful.” [DP2]

“I like those vertical comparisons [in 1D heatmaps]. I can clearly know in which one (segment) dancers are not well synchronized in each practice. We can then focus on practicing these parts (segments).” [DL-3]

Our participants also enjoyed the spotlight view and commented its potential benefits.

“I really like this function. It can intelligently recommend us segments in which we are very synchronized or out of synchronization.” [DP5]

7.1.2 Offering an Objective Assessment on Dance Performances. In five of the interviews (DL2, DP1, DP3, DP4, and DP6), we received explicit comments about how SyncUp supported the identification of ill-synchronized segments more accurately than a standard video player interface. One reason mentioned by our participants is that people may have biases when they review practice videos.

“If a person dances very well, then you will unconsciously focus your attention on that person, and ignore the others. When you look back again, you may find they are synchronized quite well, but in fact, they are not.” [DP4]

“Usually there are lots of people dancing together, and I might not be able to notice whether they are synchronized or not.” [DL1]

They appreciated that our system could avoid these situations by “treating everyone equally” (DP4). Our participants also appreciated that SyncUp considers temporal alignments. This feature is helpful for those who are not well-trained with a sense of tempo.

“In fact, there truly exist people who are not sensitive to the tempo ... If there exists an app that can help this, it can effectively help those people, like me.” [DP4]

One interesting anecdote was that an objective assessment offered by SyncUp can be more socially acceptable and could encourage communication among dancers.

"[Other] dancers are usually my friends ... If you directly say you are wrong here and there, it would hurt our friendship ... In the past, we actually noticed lots of places that we all were wrong. Because we were shy, we just kept silent." [DL3]

"It (SyncUp) is objective. For those highlighted places (segments), we can click one, and check it together. This encourages some proper communication." [DL3]

7.1.3 Enabling a Quick Iteration of Review and Practice. Our participants confirmed the benefits of video recordings for their practices. However, they also explicitly mentioned that they were unable to review their videos at the place due to time constraints imposed by external factors:

"Mostly we watch these videos after we go back home because it is time-consuming. Also, there is a time limit in renting a dance studio." [DP4]

"Although we might occasionally watch [practice videos] on the spot, because we pay for renting the dance studio and the time is limited, it is impossible to examine videos in detail." [DP3]

Participants agreed that SyncUp could help them quickly review practice videos without viewing the full footage closely. In particular, they appreciated the "high efficiency" (DL1, DL2, DL3, DP1, DP2, DP3, and DP6) of our system.

"If I use the system, it will automatically recognize those ill-synchronized segments. In this way, we can focus on these parts [during the practice], saving lots of time." [DL2]

"The system can save time in many aspects of the practice. For example, we can record a video, and the system outputs the analysis, showing several segments. Then we can just quickly practice these parts (segments). Compared with slowly watching these videos and checking which parts (segments) we can improve after we go home, this (SyncUp) is very efficient." [DP6]

7.1.4 Other Potential Use of SyncUp. Participants also shared their potential creative uses of SyncUp for their practices. DP6 explained their thoughts on how differently SyncUp could be useful depending on the dancer's skill level.

"For those dancers who are already familiar with their dance routine, they would use it at a later stage. They could study some tiny issues of the pose similarity. However, for beginners, they might not be able to reach that stage, and would probably use this to help evaluate their tempo. Maybe from the early stage." [DP6]

DP4 expressed their interests in using SyncUp to appreciate their good performances. They further mentioned that automatic extraction of good performance segments would be useful for sharing with online social networks.

"We record videos not only to find where we are out-of-synchronization, but also want to enjoy watching the movements that we are perfectly synchronized ... If this system can have these moments there (in the spotlight view), we can just download these videos and share them in our social app. We do not need to do video editing any more." [DP4]

7.2 Discussion

Our qualitative study with target users confirmed that SyncUp has a strong potential to support their practices. In particular, they agreed that it would allow for a quick iterative practice process, which aligns with our goal. Our results also imply that one major advantage of SyncUp is its objective assessment of dance performances. As DL3 commented, direct critical feedback on others' mistakes can be strongly discouraging for people from certain cultural backgrounds. A system like SyncUp could serve as a mediator in such cases, and help group dancers maintain their relationships while allowing them to exchange honest feedback on their dance performances.

Our participants suggested several improvements to make SyncUp more practical. These suggestions included a customizable synchronization score calculation method to reflect their preferences; an intelligent segmentation method that considers thematic changes of dances; and customizable discretization in overlay coloring (e.g., showing overlays only when the synchronization score becomes below a pre-defined threshold).

8 LIMITATIONS

This work has limitations that we mentioned herein. The current interface design and functionality are tailored toward synchronized dancing, choreography consisting of a series of synchronized poses and/or temporally-aligned movements among multiple dancers. Although some of the features may be useful for other kinds of dancing or motions executed by multiple people, future work should examine whether such extensions would match user requirements. For example, a Mexican wave (or a stadium wave) is a type of collective motions by multiple people which is beyond the scope of the current SyncUp implementation. Supporting such a metachronal rhythm would be an interesting direction for future work.

In our system evaluations, our data only covered particular types of synchronized dancing (i.e., dances originally performed by Asian pop stars or anime characters). Although our performance quantification methods do not assume domain knowledge of dances, future work should study how applicable they would be to different types of synchronized dancing.

In our current implementation, the system uses 2D human skeletons. Recent technology [22] can achieve 3D skeletons with a normal RGB camera, but the computation cost is much higher. Our main contributions lie in the intergration of computer vision technology to enable quick reviews of synchronized dancing practice videos rather than improving the accuracy of pose similarity analysis in synchronized dancing.

Our temporal alignment analysis assumes that ground truth movements (i.e., dances by the leader) are given. Another approach is to identify the beats of the background music and quantify how well dancers' motions are aligned with them. We experimented with this, but decided not to incorporate it into the current SyncUp implementation. We observed that dancers' movements are not always on the beats of the background music even though their dancing can be considered well synchronized. This issue can be intrinsically challenging, and further explorations on human perception of temporal alignment in dances are encouraged.

In our current implementation, video analysis is performed using a remote computing resource in a post-hoc manner (i.e., after one practice session was complete.) instead of on a mobile device in real time. Thus, users would have a small amount of waiting time. However, we do envision SyncUp would be able to provide feedback while dancers are taking a short break after running multiple practice sessions. In this manner, SyncUp can contribute to mitigating large isolations between practices and reviews with recordings that dancers currently experience. Future implementations may include more computationally efficient pose detection and temporal alignment analysis, which may be executable for a smaller amount of time and/or on a mobile device. Recent computer vision technology has achieved real-time human pose detection on mobile devices [31], however, the extension to multi-person pose detection is still under development. We note that our scope of this work is the development of dance analysis methods tailored toward synchronized dancing and both system and user evaluations on the proposed system instead of improving the computational costs of the analysis process. Our work encourages further research and development of such computer vision technologies by demonstrating applications and interface designs that fit users' needs and workflow.

9 CONCLUSION

Synchronized dancing attracts amateur dancers, but interactive support for its practice is still insufficient. This paper presents SyncUp, a vision-based interactive system that includes multiple visualizations and offers quick access to segments where dances are not well synchronized. The system integrates two approaches to quantify

the performance of synchronized dancing: pose similarity and temporal alignment. Our system evaluations confirm that performance predictions by SyncUp are highly correlated with human ratings. Participants in our qualitative user study shared their positive opinions on the features in SyncUp, and they commented several potential uses of the system. Future work should conduct an in-depth user study through the deployment of SyncUp into actual practice sessions to validate the results reported in this paper.

ACKNOWLEDGMENTS

We would like to thank Arissa J. Sato, Carla F. Griggio, Zefan Sramek, and all the anonymous reviewers for their valuable feedback on our manuscript. We also appreciate Minghui Chen for his support to create the demonstration video of this project. We finally want to show our gratitude to all participants in our user studies, interviews, and surveys for their invaluable comments that are inspiring for both our work and future work. This research received support from the NII CRIS collaborative research program jointly managed by NII CRIS and LINE Corporation.

REFERENCES

- [1] Fraser Anderson, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2013. YouMove: Enhancing Movement Training with an Augmented Reality Mirror. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (St. Andrews, Scotland, United Kingdom) (*UIST '13*). ACM, New York, NY, USA, 311–320. <https://doi.org/10.1145/2501988.2502045>
- [2] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2014. 2D Human Pose Estimation: New Benchmark and State of the Art Analysis. In *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR '14)*. IEEE Computer Society, USA, 3686–3693. <https://doi.org/10.1109/CVPR.2014.471>
- [3] Bettina Bläsing, Beatriz Calvo-Merino, Emily S. Cross, Corinne Jola, Julianne Honisch, and Catherine J. Stevens. 2012. Neurocognitive control in dance perception and performance. *Acta Psychologica* 139, 2 (2012), 300 – 308. <https://doi.org/10.1016/j.actpsy.2011.12.005>
- [4] Antonio Camurri, Corrado Canepa, Nicola Ferrari, Maurizio Mancini, Radosław Niewiadomski, Stefano Piana, Gualtiero Volpe, Jean-Marc Matos, Pablo Palacio, and Muriel Romero. 2016. A System to Support the Learning of Movement Qualities in Dance: A Case Study on Dynamic Symmetry. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct* (Heidelberg, Germany) (*UbiComp '16*). Association for Computing Machinery, New York, NY, USA, 973–976. <https://doi.org/10.1145/2968219.2968261>
- [5] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2019. OpenPose: Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields. arXiv:1812.08008
- [6] Jacky C. P. Chan, Howard Leung, Jeff K. T. Tang, and Taku Komura. 2011. A Virtual Reality Dance Training System Using Motion Capture Technology. *IEEE Trans. Learn. Technol.* 4, 2 (2011), 187–195. <https://doi.org/10.1109/TLT.2010.27>
- [7] Cheng Chen, Yueting Zhuang, Feiping Nie, Yi Yang, Fei Wu, and Jun Xiao. 2010. Learning a 3D human pose distance metric from geometric pose descriptor. *IEEE Transactions on Visualization and Computer Graphics* 17, 11 (2010), 1676–1689.
- [8] Marianela Cioffi Felice, Sarah Fdili Alaoui, and Wendy E. Mackay. 2018. Knotation: Exploring and Documenting Choreographic Processes. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (Montreal QC, Canada) (*CHI '18*). ACM, New York, NY, USA, Article 448, 12 pages. <https://doi.org/10.1145/3173574.3174022>
- [9] Abe Davis and Maneesh Agrawala. 2018. Visual Rhythm and Beat. *ACM Trans. Graph.* 37, 4, Article 122 (July 2018), 11 pages. <https://doi.org/10.1145/3197517.3201371>
- [10] Dieter Drobny, Malte Weiss, and Jan Borchers. 2009. Saltate!: A Sensor-based System to Support Dance Beginners. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems* (Boston, MA, USA) (*CHI EA '09*). ACM, New York, NY, USA, 3943–3948. <https://doi.org/10.1145/1520340.1520598>
- [11] Harris Drucker, Chris JC Burges, Linda Kaufman, Alex Smola, Vladimir Vapnik, et al. 1997. Support vector regression machines. *Advances in neural information processing systems* 9 (1997), 155–161.
- [12] Yong Du, Wei Wang, and Liang Wang. 2015. Hierarchical recurrent neural network for skeleton based action recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*. IEEE Computer Society, 1110–1118. <https://doi.org/10.1109/CVPR.2015.7298714>
- [13] Katerina El Raheb, Sarah Whatley, and Antonio Camurri. 2018. A Conceptual Framework for Creating and Analyzing Dance Learning Digital Content. In *Proceedings of the 5th International Conference on Movement and Computing* (Genoa, Italy) (*MOCO '18*). Association for Computing Machinery, New York, NY, USA, Article 2, 8 pages. <https://doi.org/10.1145/3212721.3212837>
- [14] Daniel Ellis. 2014. The 2014 labrosa audio fingerprint system.
- [15] Daniel PW Ellis. 2007. Beat tracking by dynamic programming. *Journal of New Music Research* 36, 1 (2007), 51–60.

- [16] Haoshu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. 2017. RMPE: Regional Multi-person Pose Estimation. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017*. IEEE Computer Society, 2353–2362. <https://doi.org/10.1109/ICCV.2017.256>
- [17] Ross B. Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2014. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*. IEEE Computer Society, 580–587. <https://doi.org/10.1109/CVPR.2014.81>
- [18] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger (Eds.). 2672–2680. <https://proceedings.neurips.cc/paper/2014/hash/5ca3e9b122f61f8f06494c97b1afccf3-Abstract.html>
- [19] Stacy Hsueh, Sarah Fdili Alaoui, and Wendy E. Mackay. 2019. Understanding Kinaesthetic Creativity in Dance. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (Glasgow, Scotland UK) (CHI '19)*. ACM, New York, NY, USA, Article 511, 12 pages. <https://doi.org/10.1145/3290605.3300741>
- [20] Falisha J Karpati, Chiara Giacosa, Nicholas EV Foster, Virginia B Penhune, and Krista L Hyde. 2015. Dance and the brain: a review. *Annals of the New York Academy of Sciences* 1337, 1 (2015), 140–146.
- [21] Diederik P. Kingma and Jimmy Ba. 2017. Adam: A Method for Stochastic Optimization. arXiv:1412.6980
- [22] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. 2020. VIBE: Video Inference for Human Body Pose and Shape Estimation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*. IEEE, 5252–5262. <https://doi.org/10.1109/CVPR42600.2020.00530>
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 1097–1105.
- [24] Alexander Krull, Eric Brachmann, Frank Michel, Michael Ying Yang, Stefan Gumhold, and Carsten Rother. 2015. Learning Analysis-by-Synthesis for 6D Pose Estimation in RGB-D Images. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*. IEEE Computer Society, 954–962. <https://doi.org/10.1109/ICCV.2015.115>
- [25] Matthew Kyan, Guoyu Sun, Haiyan Li, Ling Zhong, Paisarn Muneesawang, Nan Dong, Bruce Elder, and Ling Guan. 2015. An Approach to Ballet Dance Training through MS Kinect and Visualization in a CAVE Virtual Reality Environment. *ACM Trans. Intell. Syst. Technol.* 6, 2, Article 23 (March 2015), 37 pages. <https://doi.org/10.1145/2735951>
- [26] Hsin-Ying Lee, Xiaodong Yang, Ming-Yu Liu, Ting-Chun Wang, Yu-Ding Lu, Ming-Hsuan Yang, and Jan Kautz. 2019. Dancing to Music. arXiv:1911.02001
- [27] Jae-Jun Lee, Jong-Hyeok Choi, Tserenpurev Chuluunsaikhan, and Aziz Nasridinov. 2020. Pose Evaluation for Dance Learning Application Using Joint Position and Angular Similarity. In *Adjunct Proceedings of the 2020 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2020 ACM International Symposium on Wearable Computers (Virtual Event, Mexico) (UbiComp-ISWC '20)*. Association for Computing Machinery, New York, NY, USA, 67–70. <https://doi.org/10.1145/3410530.3414402>
- [28] Ruilong Li, Shan Yang, David A. Ross, and Angjoo Kanazawa. 2021. Learn to Dance with AIST++: Music Conditioned 3D Dance Generation. arXiv:2101.08779
- [29] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common Objects in Context. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V (Lecture Notes in Computer Science, Vol. 8693)*, David J. Fleet, Tomás Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.). Springer, 740–755. https://doi.org/10.1007/978-3-319-10602-1_48
- [30] Jun Liu, Amir Shahroudy, Dong Xu, and Gang Wang. 2016. Spatio-Temporal LSTM with Trust Gates for 3D Human Action Recognition. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part III (Lecture Notes in Computer Science, Vol. 9907)*, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.). Springer, 816–833. https://doi.org/10.1007/978-3-319-46487-9_50
- [31] Camillo Lugaresi, Jiuqiang Tang, Hadon Nash, Chris McClanahan, Esha Uboweja, Michael Hays, Fan Zhang, Chuo-Ling Chang, Ming Guang Yong, Juhyun Lee, Wan-Teh Chang, Wei Hua, Manfred Georg, and Matthias Grundmann. 2019. MediaPipe: A Framework for Building Perception Pipelines. arXiv:1906.08172
- [32] Alireza Moghaddam. 2006. Coding issues in grounded theory. *Issues in educational research* 16, 1 (2006), 52–66.
- [33] Meinard Müller, Tido Röder, and Michael Clausen. 2005. Efficient Content-Based Retrieval of Motion Capture Data. In *ACM SIGGRAPH 2005 Papers (Los Angeles, California) (SIGGRAPH '05)*. Association for Computing Machinery, New York, NY, USA, 677–685. <https://doi.org/10.1145/1186822.1073247>
- [34] Akio Nakamura, Sou Tabata, Tomoya Ueda, Shinichiro Kiyofuji, and Yoshinori Kuno. 2005. Multimodal Presentation Method for a Dance Training System. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems (Portland, OR, USA) (CHI EA '05)*. ACM, New York, NY, USA, 1685–1688. <https://doi.org/10.1145/1056808.1056997>

- [35] Katerina El Raheb, Marina Stergiou, Akrivi Katifori, and Yannis Ioannidis. 2019. Dance Interactive Learning Systems: A Study on Interaction Workflow and Teaching Approaches. *ACM Comput. Surv.* 52, 3, Article 50 (June 2019), 37 pages. <https://doi.org/10.1145/3323335>
- [36] Jean-Philippe Rivière, Sarah Fdili Alaoui, Baptiste Caramiaux, and Wendy E. Mackay. 2018. How Do Dancers Learn To Dance? A First-Person Perspective of Dance Acquisition by Expert Contemporary Dancers. In *Proceedings of the 5th International Conference on Movement and Computing* (Genoa, Italy) (MOCO '18). Association for Computing Machinery, New York, NY, USA, Article 6, 7 pages. <https://doi.org/10.1145/3212721.3212723>
- [37] Jean-Philippe Rivière, Sarah Fdili Alaoui, Baptiste Caramiaux, and Wendy E. Mackay. 2019. Capturing Movement Decomposition to Support Learning and Teaching in Contemporary Dance. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 86 (Nov. 2019), 22 pages. <https://doi.org/10.1145/3359188>
- [38] Jan Sedmidubský, Jakub Valčík, and Pavel Zezula. 2013. A Key-Pose Similarity Algorithm for Motion Data Retrieval. In *Advanced Concepts for Intelligent Vision Systems - 15th International Conference, ACIVS 2013, Poznań, Poland, October 28-31, 2013. Proceedings (Lecture Notes in Computer Science, Vol. 8192)*, Jacques Blanc-Talon, Andrzej J. Kasinski, Wilfried Philips, Dan C. Popescu, and Paul Scheunders (Eds.). Springer, 669–681. https://doi.org/10.1007/978-3-319-02895-8_60
- [39] Lei Shi, Yifan Zhang, Jian Cheng, and Hanqing Lu. 2019. Two-Stream Adaptive Graph Convolutional Networks for Skeleton-Based Action Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*. Computer Vision Foundation / IEEE, 12026–12035. <https://doi.org/10.1109/CVPR.2019.01230>
- [40] Jamie Shotton, Andrew W. Fitzgibbon, Mat Cook, Toby Sharp, Mark Finocchio, Richard Moore, Alex Kipman, and Andrew Blake. 2011. Real-time human pose recognition in parts from single depth images. In *The 24th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2011, Colorado Springs, CO, USA, 20-25 June 2011*. IEEE Computer Society, 1297–1304. <https://doi.org/10.1109/CVPR.2011.5995316>
- [41] Jamie Shotton, Ross B. Girshick, Andrew W. Fitzgibbon, Toby Sharp, Mat Cook, Mark Finocchio, Richard Moore, Pushmeet Kohli, Antonio Criminisi, Alex Kipman, and Andrew Blake. 2013. Efficient Human Pose Estimation from Single Depth Images. *IEEE Trans. Pattern Anal. Mach. Intell.* 35, 12 (2013), 2821–2840. <https://doi.org/10.1109/TPAMI.2012.241>
- [42] Jennifer J. Sun, Jiaping Zhao, Liang-Chieh Chen, Florian Schroff, Hartwig Adam, and Ting Liu. 2020. View-Invariant Probabilistic Embedding for Human Pose. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part V (Lecture Notes in Computer Science, Vol. 12350)*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, 53–70. https://doi.org/10.1007/978-3-030-58558-7_4
- [43] Milka Trajkova and Francesco Cafaro. 2018. Takes Tutu to Ballet: Designing Visual and Verbal Feedback for Augmented Mirrors. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 1, Article 38 (March 2018), 30 pages. <https://doi.org/10.1145/3191770>
- [44] Baijun Xie and Chung Hyuk Park. 2020. Dance with a Robot: Encoder-Decoder Neural Network for Music-Dance Learning. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction* (Cambridge, United Kingdom) (HRI '20). Association for Computing Machinery, New York, NY, USA, 526–528. <https://doi.org/10.1145/3371382.3378372>
- [45] Yuliang Xiu, Jiefeng Li, Haoyu Wang, Yinghong Fang, and Cewu Lu. 2018. Pose Flow: Efficient Online Pose Tracking. In *British Machine Vision Conference 2018, BMVC 2018, Newcastle, UK, September 3-6, 2018*. BMVA Press, 53. <http://bmvc2018.org/contents/papers/0096.pdf>
- [46] Weipeng Xu, Avishek Chatterjee, Michael Zollhöfer, Helge Rhodin, Dushyant Mehta, Hans-Peter Seidel, and Christian Theobalt. 2018. MonoPerfCap: Human Performance Capture From Monocular Video. *ACM Trans. Graph.* 37, 2, Article 27 (May 2018), 15 pages. <https://doi.org/10.1145/3181973>
- [47] Sijie Yan, Yuanjun Xiong, and Dahua Lin. 2018. Spatial Temporal Graph Convolutional Networks for Skeleton-Based Action Recognition. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, Sheila A. McIlraith and Kilian Q. Weinberger (Eds.). AAAI Press, 7444–7452. <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17135>
- [48] Yongqi Zhang, Biao Xie, Haikun Huang, Elisa Ogawa, Tongjian You, and Lap-Fai Yu. 2019. Pose-Guided Level Design. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). ACM, New York, NY, USA, Article 554, 12 pages. <https://doi.org/10.1145/3290605.3300784>
- [49] Zhongyi Zhou, Yuki Tsubouchi, and Koji Yatani. 2019. Visualizing Out-of-Synchronization in Group Dancing. In *The Adjunct Publication of the 32nd Annual ACM Symposium on User Interface Software and Technology* (New Orleans, LA, USA) (UIST '19). Association for Computing Machinery, New York, NY, USA, 107–109. <https://doi.org/10.1145/3332167.3356888>
- [50] Wenlin Zhuang, Congyi Wang, Siyu Xia, Jinxiang Chai, and Yangang Wang. 2020. Music2Dance: DanceNet for Music-driven Dance Generation. [arXiv:2002.03761](https://arxiv.org/abs/2002.03761)

A EXAMPLES OF POSE SIMILARITY VISUALIZATION

Figure 11 provides three examples of side-by-side views with heatmap overlays alongside pose similarity score visualizations. Note that a darker red color in the 1D heatmap visualization indicates a lower pose similarity score.

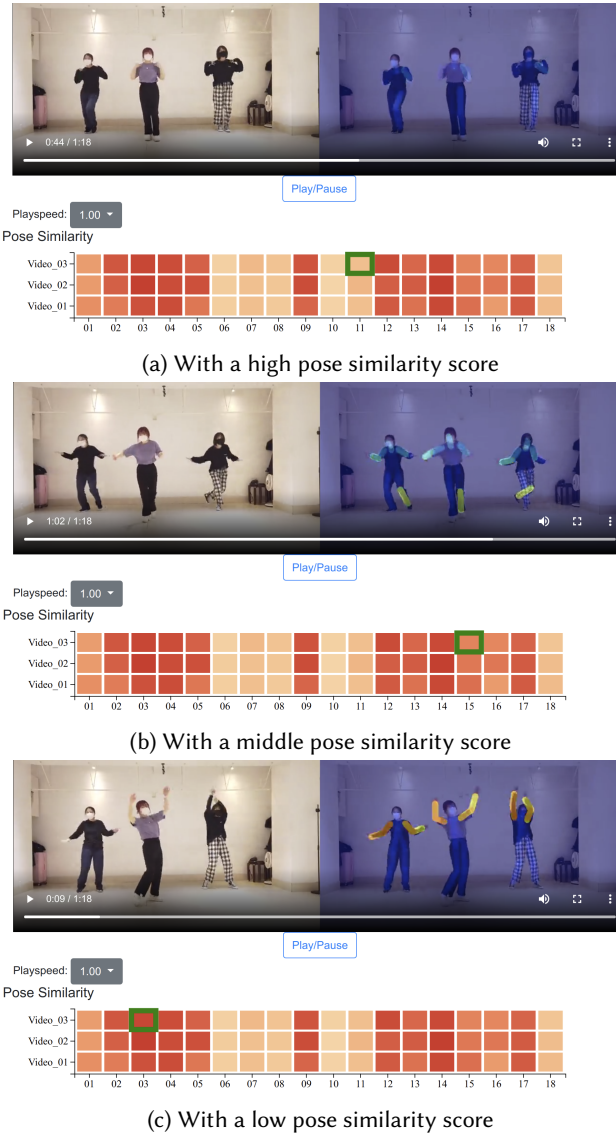


Fig. 11. Three examples of visualization related to pose similarity. These examples were chosen to illustrate cases where the pose similarity score was low, in the middle, and high.

B EXAMPLE FRAMES WITH HIGH PREDICTION ERRORS IN THE ANALYSIS OF POSE SIMILARITY

Figure 12 shows frames that caused the five largest pose similarity prediction errors with each OPS quantification method. The top figure in each column shows the frame that exhibited the largest error with the corresponding method. This result suggests that one major cause of large pose similarity prediction error was occlusion of body parts. Other error attributes include lack of hand pose detection (the bottom of Figure 12a) and lack of 3D pose detection support (the fourth figure in Figure 12d). Different from our learning-based method in which errors are caused by our dependent pose detection algorithms, we observed that the Simple Addition method makes high-error predictions even when the pose detection is correct. Such hard-to-interpret errors may confuse users and developers in practice.

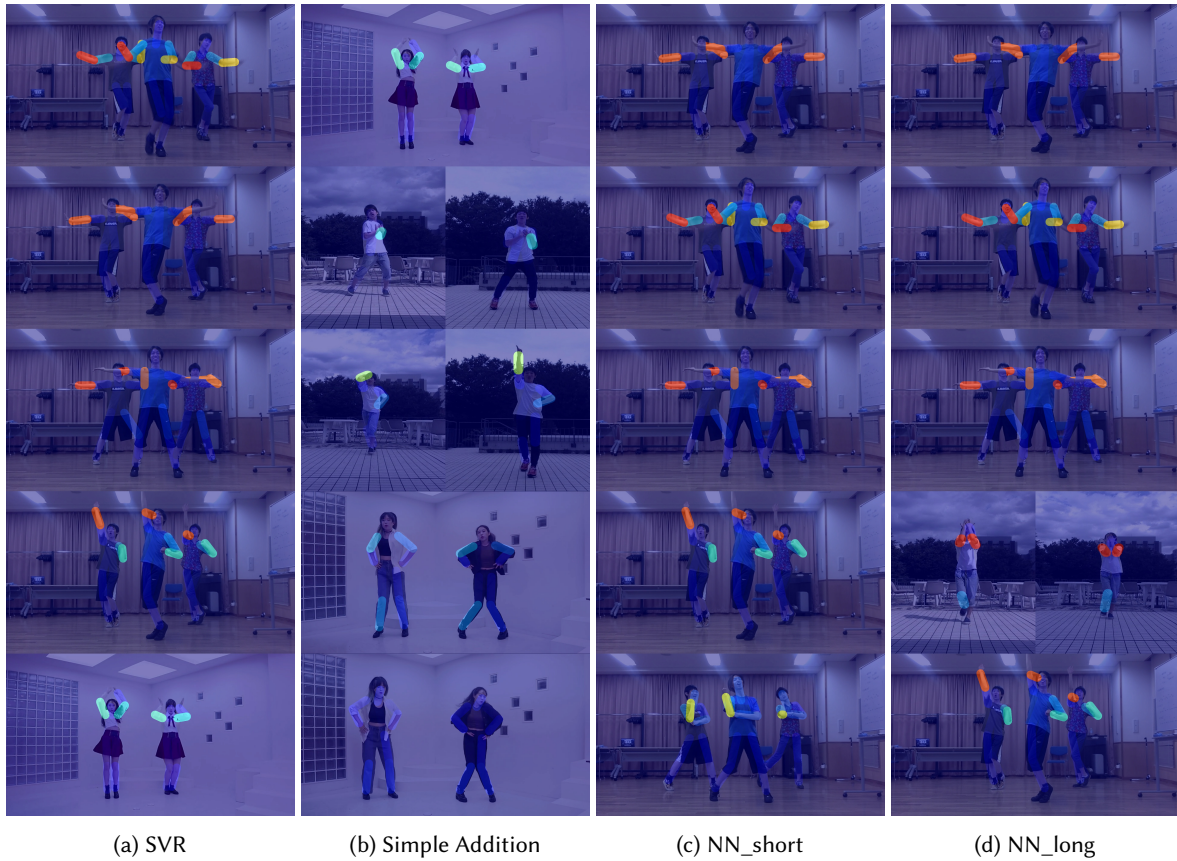


Fig. 12. Frames that caused the five largest pose similarity prediction errors with each OPS quantification method. The top figure in each column shows the frame that exhibited the largest error with the corresponding method.