Vis Comput
DOI 10.1007/s00371-017-1452-z

ORIGINAL ARTICLE

Style-based motion analysis for dance composition

Andreas Aristidou[1,2] · Efstathios Stavrakis[1] · Margarita Papaefthimiou[3] · George Papagiannakis[4] · Yiorgos Chrysanthou[1]

Abstract Synthesizing human motions from existing motion capture data is the approach of choice in most applications requiring high- quality visual results. Usually to synthesize motion, short motion segments are concatenated into longer sequences by finding transitions at points where character poses are similar. If similarity is only a measure of posture correlation, without consideration for the stylistic variations of movement, the resulting motion might have unnatural discontinuities. Particularly prone to this problem are highly stylized motions, such as dance performances. This work presents a motion analysis framework, based on Laban Movement Analysis, that also accounts for stylistic variations of the movement. Implemented in the context of Motion Graphs, it is used to eliminate potentially problematic transitions and synthesize style-coherent animation, without requiring prior labeling of the data. The effectiveness of our method is demonstrated by synthesizing contemporary dance performances that include a variety of different emotional states. The algorithm is able to compose highly stylized motions that are reminiscent to dancing scenarios using only plausible movements from existing clips.

Electronic supplementary material The online version of this article (https://doi.org/10.1007/s00371-017-1452-z) contains supplementary material, which is available to authorized users.

✉ Andreas Aristidou
a.aristidou@ieee.org

1 University of Cyprus, Nicosia, Cyprus

2 Interdisciplinary Center Herzliya, Herzliya, Israel

3 Institute of Computer Science of the Foundation for Research and Technology Hellas, Heraklion, Greece

4 University of Crete, Rethymnon, Greece

Keywords Laban Movement Analysis · Motion Graphs · Motion style · Motion synthesis

1 Introduction

Motion capture (mocap) technology has advanced to the point that fine-grained digitization of human motion has become widespread and has found utility in entertainment, sciences, sports, education, etc. Despite the availability of large motion capture datasets, their analysis, processing and ultimately their reuse for synthesis of novel motions remain hard problems. Motion analysis consists of inquiring about the content of the different types of human actions (e.g., walking, dancing, running or jumping) and their stylistic variations (e.g., intention, expression). In contrast, synthesis is faced with the difficulty of generating *plausible* motion.

Early computational motion analysis and synthesis techniques focused on the quantitative characteristics of human motion (e.g., distances, velocities), while more recent research is increasingly interested in the qualitative attributes of motion. These attributes are directly linked to the personality of the individual, including his/her style, emotion, effort and the purpose of the action, reflecting its *nuance*.

Though current motion capture and animation technologies have made certain strides in simulating this subtle dynamic range of motion, the need for more sophisticated approaches is essential if lifelike human motion is to be produced. One challenging example for many motion analysis and synthesis algorithms is generating plausible dance animations. Dancing consists of complex and highly dynamic movements, including heterogenous and arbitrary poses that are difficult to analyze using traditional methods. Techniques such as Motion Graph (MG) [19] combine motion clips based on posture similarities to create longer motions. However,

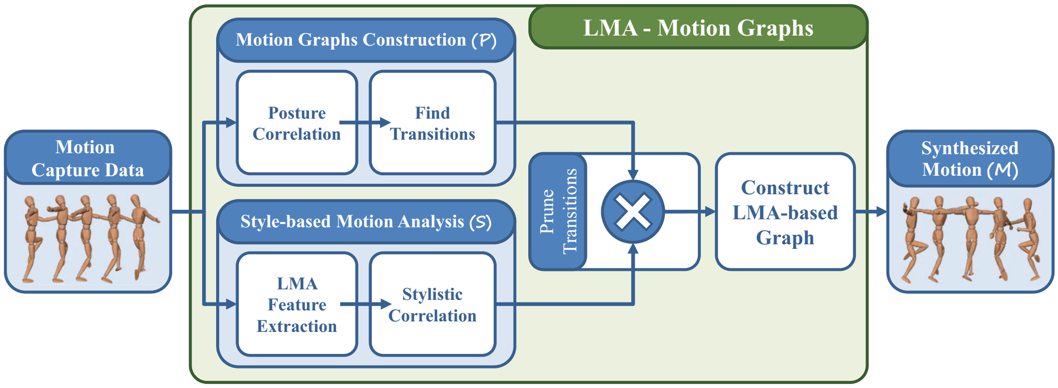

Fig. 1 An overview of our method; the transitions of the generated Motion Graph are pruned based on the similarity of the LMA-derived features between frames. The remaining transitions ensure that style is coherent

they disregard the stylistic qualities of motion and require special constraints (e.g., manual labeling) in order to splice motions together that have continuity in their style, as is the case with dancing motions. Ren et al. [34] observed that motions synthesized by utilizing incorrect transitions often have significant velocity discontinuities and are thus perceived as being unnatural; this is because the blending procedure attempts to smoothly merge two distant poses with different styles.

In this work, we propose an efficient method that can automatically extract motion characteristics for synthesis of dance animation. The method leverages knowledge from anatomy, kinesiology and psychology as that is incorporated in the Laban Movement Analysis (LMA) [21] framework. LMA draws on Rudolph Laban's theories and examines movement through four interrelated components: Body, Effort, Shape and Space. The LMA framework has been widely used to map the personal movement vocabulary and skills in different areas, including dance, choreography, theater and kinesiology.

Similarly, we express the four LMA components via a selection of 114 low-level metrics that are obtained by computationally processing motion captured data. These metrics capture qualities in motions that go beyond the body's postural configuration. We compute motion similarity by calculating the Pearson correlation coefficients of these LMA-derived features between frames of different motion sequences. We adopt the Motion Graph data structure to arrange our motions in a graph, but in addition to geometric posture similarities, we use the style-based LMA correlations to connect its nodes (i.e., the frames at which motion clip transitions can occur). The LMA-based Motion Graph (LMA MG) is utilized in two motion synthesis scenarios to demonstrate that dance animations generated with our method are plausible and their style is fairly consistent. In this work, we use emotion as a representative kind of style that is easily recognizable. In line with emotion research [35], we use acted data of various contemporary dance scenarios, each expressing a different emotional state (e.g., angry, excited). To verify our results, we devised a transition cost metric based on these ground truth emotion states, which is calculated over consecutive motion clips comprising a complete synthesized motion. Motions generated with our method have lower costs than motions generated using standard methods. In addition, we have conducted an online user survey in which participants preferred the resulting motions of our method over motions generated with plain Motion Graphs.

Figure 1 shows an overview of our methodology; after data acquisition, the procedure is divided into two individual processes. The first process calculates posture correlations so as to construct a Motion Graph and find the potential transitions between clips, while the second process encodes both the movement's quantitative and qualitative characteristics, based on principles drawn from Laban Movement Analysis, in order to find their stylistic correlations. Finally, these LMA correlations are used to prune transitions of the Motion Graph that are not stylistically coherent, leading to an LMA-based Motion Graph. This novel Motion Graph can be used to synthesize plausible dance animations. Note that our method does not produce dance motions that are choreographically correct neither does it attempt to compose a real dancing performance.

The main contributions of this work are:

- A novel style-coherent motion analysis algorithm based on Laban Movement Analysis principles.
- A framework for synthesizing style-coherent animation based on Laban Movement Analysis, without requiring prior labeling of the data.
- A style-preserving motion synthesis technique that can be readily used with the popular Motion Graphs algorithm.

2 Related work

Central to our approach is an LMA-based technique for identifying and linking instances of similar pose and style within motion sequences. The links manifest as edges on a Motion Graph [2,19,22] to provide style-coherent motion synthesis, targeted at dance performances. The literature on data-driven animation is quite broad [33]; here we will review some of the most relevant work relating to style, dance and LMA. Techniques in the area of indexing and classification also look at the relation and similarities between motions [9,14,18,20,28]. However, these techniques are primarily aiming to identify logically similar motions, while in this work we focus more on style rather than content.

Style There are a number of techniques that deal explicitly with the style of movement. Style transfer methods look at the issue of carrying over the style from one character motion onto another. Brand and Hertzmann [6] use machine learning techniques with a hidden Markov model to capture style. Hsu et al. [16] use iterative time warping to map from the input to the output sequence, while Shapiro et al. [37] use style components. Other approaches learn a parametric model that enables them to interpolate and extrapolate to new styles. For example, Torresani et al. [40] employed LMA to map into perceptual space, while Urtasun et al. [41] used PCA to encode style. Hartmann et al. [13] presented an augmentation to the GRETA agent architecture [12], where the authors described the gesture selection process, that allows for parametric control of the qualitative aspects of gestures. The authors present a computational model of gesture quality, drawing from psychology research, which allows behavior modification via gesture synthesis, thus creating an expressive Embodied Conversational Agent. Vasilescu [42] organized data as higher-order arrays, and then by using singular value decomposition (SVD), extracted motion style factors. Min et. al [26] looked at actions performed by multiple actors and in various styles to produce a generative model with two parameters, "identity" and "style"; *Motion Graphs++* were later introduced for semantic motion analysis and synthesis [25]. Müller and Röder [27] introduced Motion Templates, in which logically related motions were classified and retrieved from motion databases. Recently, considerable effort has been devoted to transferring style from one motion to another; for instance, Xia et al. [45] introduced local mixtures of autoregressive models to capture the relationship between styles of motions, while Yumer and Mitra [47] proposed a method for style transfer based on spectral analysis. Holden et al. [15] employed deep learning for animation synthesis; the authors learned motion manifolds, which are represented by the hidden units of a convolutional autoencoder, to synthesize the style of motion via interpolation. Aristidou et al. [5] learnt regression models to map motion onto a parametric space of emotion that was then leveraged to stylize motion by modifying selected features.

Laban Movement Analysis (LMA) The principles of LMA have been used in computer animation for over a decade. Chi et al. [10] introduced the EMOTE system, in which a set of parameters, inspired by the LMA Effort and Shape components, is presented. These parameters are used to synthesize gestures for motion parameterization and expression. Later, Zhao and Badler [48] used the EMOTE results to design a neural network for gesture animation. Many works in the literature used LMA to quantify the expressive content of gestures or to learn motion styles. Torresani et al. [40] learn a nonlinear mapping between animation parameters and movement styles in perceptual space. This mapping can then be utilized to synthesize stylistic variations from artificially generated examples using the LMA Effort factors. Luo and Neff [24] studied the relationship between posture and gesture for virtual characters using LMA components; in addition, Wakayama et al. [43] and Okajima et al. [30] used these components for motion retrieval. Similarly, Kapadia et al. [17] encoded structural, geometric and dynamic features of motion as keys, as part of an indexing process and for searching complex motions in large motion databases. These keys were then combined to specify search queries to retrieve motions. Shiratori et al. [38] used Laban theory for synthesizing dance motion matched to music, while Nakata et al. [29] used the LMA effort component to explain bodily expressions. Aristidou et al. [3] have presented an LMA-based framework to extract the quantitative and qualitative characteristics of acted contemporary dances, aiming to classify motion sequences with regard to the expressed emotion. Senecal et al. [36] used the Aristidou et al. framework to create social agents that recognize the emotion of users by mapping motion onto an emotion diagram using neural networks. More recently, Durupinar et al. [11] have presented PERFORM, an approach for controlling the personality of human motion. The authors applied the knowledge earned by studying the LMA system to create variation in the motion styles and satisfy user-assigned personality traits.

Dance A number of computational tools or symbolic systems for dance and choreography are available, such as DanceForms [7], and LabanDancer [44]. Shiratori et al. [38] synthesized dance sequences by taking into account the music features, while Oliveira et al. [31] relied on beat-segmented mocap to create a compressed representation of dance gestures which was then used to synthesize dance movements. In Li et al. [23], they use motion textures and apply them to dancing. Aristidou et al. [4] presented a method that uses motion qualities to assess the similarity between dancing motions. Alexiadis and Daras [1] designed a framework for automatic dance performance evaluation using motion capture data using marker-less motion capture. The authors represented the human motion data as sequences of pure quaternions and subsequently introduced a set of quaternionic vector-signal processing methodologies for dance motion evaluation and comparison purposes. A number of educational or gaming dance systems have also been proposed, e.g., [8,39,46], where motion comparisons are a core component; however, this is typically done using geometric posture similarities.

3 Motion analysis

Human motion analysis is particularly challenging, especially when movement qualities and stylistic characteristics are of high importance. The difficulty is even more pronounced when we aim at smooth transitions in motion composition using highly stylistic motions, such as dance movements. In this work, we used a motion analysis framework which is based on the LMA principles, aiming to identify those factors that describe the movement signature of the performer. LMA is a language for interpreting, describing, visualizing and notating human movement; it offers a holistic documentation of the human motion and it is divided into four components: (a) BODY, which describes the structural and physical characteristics of the human body, (b) EFFORT, which describes the intention and the dynamic quality of the movement, the texture, the feeling tone and how the energy is being used on each motion, (c) SHAPE, which analyzes the way the body changes shape during movement, and (d) SPACE, which describes the movement in relation with the environment.

In order to achieve style-coherent motion synthesis, as we describe later in Sect. 4, we utilize the LMA framework described by Aristidou et al. [3]; the authors proposed 27 basic spatiotemporal features for motion analysis, covering all LMA components. A list of the LMA-derived features is presented in Table 1.[1] We extend this LMA framework by adding some features that take into consideration the modes of interaction with oneself, others, and the environment, aiming at improving the stylistic coherence with regard to the SPACE component that has not been thoroughly investigated in [3]. Features f^{28} and f^{29} have been specifically chosen to enable us to prune motion transitions between motions that have large variations in terms of the relationship between the performer's own movement and the environment. In addition, in contrast to [3], the measurements for each side of the body are treated independently so as to increase the sensitivity of the system when comparing movements and improve the smoothness of the synthesized motion. These additional features are:

- The total volume (f^{28}) of the space covered by the performer in a time period (usually 30 frames), which describes the relationship between the performer and the environment. The total volume can be estimated by taking the union ($\cup$) of all the performer's volumes (as calculated through f^{19}); in this work it can be calculated by accumulating the bounding volume of all joints within this time period.
- The cumulative distribution (f^{29}) of the performer's movement, which is useful for estimating the performer's movement sphere. For instance, this feature enables distinguishing cases where the performer conveys a sense of central radiance or a delicate sense of peripheral boundary. This feature can be measured as the distance between the mean of the projection of the pelvis (root) on the ground and the current projection of the pelvis over the time window.

It is important to note that, in this work, the term *cumulative distribution* is related to the character's movement in space, rather than modeling the spatial tensions (peripheral, central) in a body centric manner (e.g., arm movements relative to the core), as commonly used in the literature.

4 Motion synthesis

This section presents our novel Motion Graph framework (LMA MG) for synthesizing style-coherent animation based on the LMA components. The new synthesized animation must maintain the realism of the original data and have plausible continuity. Motion Graphs, in general, achieve good connectivity and smooth transitions; however, it is possible that edges connect clips with different motion styles, since the links are formed subject only to matching their body geometries. The proposed LMA-based algorithm filters the graph to prune these redundant transitions, preserving only those with a highly similar motion style in order to increase in the naturalness of the generated motion.

As shown in Fig. 1, the proposed framework is divided into two main processes: (a) a Motion Graph construction process and (b) a style-based motion analysis process. The former process represents the posture correlation $\mathcal{P}$ of input motions, while the latter yields their style correlations. The framework combines the two correlations to compute an overall correlation ($\mathcal{M} = \mathcal{P}+\mathcal{S}$) between motion clips. More specifically, $\mathcal{P}$ computes posture correlations between motion clips and constructs a Motion Graph. For the posture correlation

[1] The volume features ($f^{19}-f^{23}$), apart from describing the LMA Shape component, as given in [3], could also give intimations of the Space component, as they additionally reveal the character's kinesphere.

Table 1 The measurements used in our implementation to compute the LMA-derived features

	Features		Measurements			
	f^i	Description	$f^i_{\max}$	$f^i_{\min}$	f^i_{σ}	f^i_{μ}
Body	f^1	Left foot-hip distance	φ_1	φ_2	φ_3	φ_4
		Right foot-hip distance	φ_5	φ_6	φ_7	φ_8
	f^2	Left hand-shoulder distance	φ_9	φ_{10}	φ_{11}	φ_{12}
		Right hand-shoulder distance	φ_{13}	φ_{14}	φ_{15}	φ_{16}
	f^3	Hands distance	φ_{17}	φ_{18}	φ_{19}	φ_{20}
	f^4	Left hand-head distance	φ_{21}	φ_{22}	φ_{23}	φ_{24}
		Right hand-head distance	φ_{25}	φ_{26}	φ_{27}	φ_{28}
	f^5	Hip-ground distance	φ_{29}	φ_{30}	φ_{31}	φ_{32}
	f^6	Hip-ground minus feet-hip	φ_{33}	φ_{34}	φ_{35}	φ_{36}
	f^7	Centroid-ground distance	φ_{37}	φ_{38}	φ_{39}	φ_{40}
	f^8	Centroid-pelvis distance	φ_{41}	φ_{42}	φ_{43}	φ_{44}
	f^9	Gait size	φ_{45}	φ_{46}	φ_{47}	φ_{48}
Effort	f^{10}	Head orientation	φ_{49}	φ_{50}	φ_{51}	φ_{52}
	f^{11}	Deceleration peaks				φ_{53}
	f^{12}	Pelvis velocity	φ_{54}		φ_{55}	φ_{56}
	f^{13}	Left-hand velocity	φ_{57}		φ_{58}	φ_{59}
		Right-hand velocity	φ_{60}		φ_{61}	φ_{62}
	f^{14}	Left foot velocity	φ_{63}		φ_{64}	φ_{65}
		Right foot velocity	φ_{66}		φ_{67}	φ_{68}
	f^{15}	Pelvis acceleration	φ_{69}		φ_{70}	
	f^{16}	Left-hand acceleration	φ_{71}		φ_{72}	
		Right-hand acceleration	φ_{73}		φ_{74}	
	f^{17}	Left foot acceleration	φ_{75}		φ_{76}	
		Right foot acceleration	φ_{77}		φ_{78}	
	f^{18}	Jerk	φ_{79}		φ_{80}	
Shape	f^{19}	Volume (5 joints)	φ_{81}	φ_{82}	φ_{83}	φ_{84}
	f^{20}	Volume (upper body)	φ_{85}	φ_{86}	φ_{87}	φ_{88}
	f^{21}	Volume (lower body)	φ_{89}	φ_{90}	φ_{91}	φ_{92}
	f^{22}	Volume (left side)	φ_{93}	φ_{94}	φ_{95}	φ_{96}
	f^{23}	Volume (right side)	φ_{97}	φ_{98}	φ_{99}	φ_{100}
	f^{24}	Torso height	φ_{101}	φ_{102}	φ_{103}	φ_{104}
	f^{25}	Hands level				$\varphi_{105}-\varphi_{107}$
Space	f^{26}	Total distance				φ_{108}
	f^{27}	Total area				φ_{109}
	f^{28}	Total volume				φ_{110}
	f^{29}	Cumulative distribution	φ_{111}	φ_{112}	φ_{113}	φ_{114}

between two motion clips, a 35-frames[2] moving window with 1-frame step is used on each clip (our motion data are sampled at 30 fps), while each time window is compared against all time windows of the other clip. The posture correlation is calculated by taking differences of body geometries of the performer, defined as a set of points that represent the joint angles, similarly to [19]. The distance d_j between two time windows anchored at the center (frame j) is the sum of their pair pose distances over the time window; the distance between two poses equals the weighted sum of the squared distances between the corresponding points $\mathbf{p_i}$ and $\mathbf{p'_i}$, as shown in the following equation:

$$\min_{\theta, x_0, z_0} \sum_i w_i \|\mathbf{p_i} - \mathbf{T}_{\theta,\mathbf{x_0},\mathbf{z_0}} \mathbf{p'_i}\|^2 \tag{1}$$

[2] previous Motion Graph implementations have suggested using a shorter window, however we set it at 35 to make it comparable to the LMA window.

for all points i, where w_i are the weight coefficients and $\mathbf{T}_{\theta,\mathbf{x_0},\mathbf{z_0}}$ is the linear transformation that rotates a point $\mathbf{p}$ about the vertical axis by θ degrees and then translates it by (x_0, z_0) that brings the root of one pose to align with the other. The distances d_j are then normalized using the formula $\hat{d}_j = (d_j - d_{\min})/(d_{\max} - d_{\min})$, where $d_{\min}$ and $d_{\max}$ are the min and max distance values over all the data. Short distances indicate a high correlation ($\mathrm{cor}_{\max} = 1$) between poses, and vice versa longer distances indicate a low correlation ($\mathrm{cor}_{\min} = 0$). The local maxima of pose correlations between two motions indicate which pairs of frames are more opportune, compared to their neighbors, for transitioning between two motions. These pairs of poses can then be selected given a user-defined correlation threshold T_{MG}. In our experiments, we found empirically that a T_{MG} set at 90% provides a reasonable amount of highly correlated poses. Having computed the local maxima of pose correlations, a Motion Graph can be constructed. The graph contains all input poses as nodes, which are interconnected at the pairs of frames identified as being highly correlated, i.e., candidate transition frames. The edges of the graph are transition motion clips that can be generated using motion interpolation between the two frames.

The second process, $\mathcal{S}$, encodes both the movement's quantitative and qualitative style characteristics, based on principles drawn from Laban Movement Analysis. To extract the proposed LMA-derived features and measure the observations, each of the motion clip's frame is filtered with a 35-frames moving window, anchored at the center. Note that for the stylistic coherence we measure motion qualities represented by distances in joint positions, similarly to [3]. We use a window stepping of 1, but this can be increased to speed up computation at the expense of accuracy. We compute, for each time window, the minimum, maximum, mean and standard deviation of the 29 basic features f^i and derive 114 feature measurements (ϕ_is).

For each window of a motion clip, a correlation metric is computed to the corresponding window of the other clip which provides an association between the time windows of the two motions. The correlation metric measures the Pearson's linear correlation coefficient [32],

$$r_{\phi^A,\phi^B} = \frac{1}{N-1}\sum_{i=1}^{N}\left(\frac{\phi_i^A - \mu_{\phi^A}}{\sigma_{\phi^A}}\right)\left(\frac{\phi_i^B - \mu_{\phi^B}}{\sigma_{\phi^B}}\right) \tag{2}$$

where $N = 114$ is the number of feature measurements at anchored frame i, μ_{ϕ^A} and σ_{ϕ^A} are the mean and standard deviation for the window of the motion clip A, respectively, and μ_{ϕ^B} and σ_{ϕ^B} are the mean and standard deviation for the window of the motion clip B. Similarly, the style-correlation metric is normalized to take values between 0 and 1. To evaluate the correlation between two performances, each of

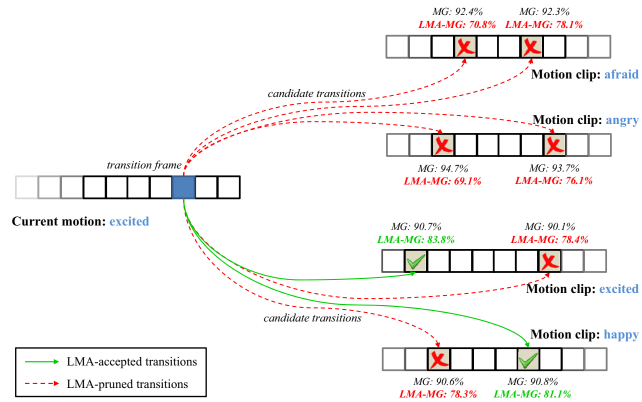

Fig. 2 In this diagram, a node (i.e., motion clip) of the constructed graph and a subset of its transitions to other nodes is shown. In this example, using the standard Motion Graph algorithm, only transitions to highly similar frames of other motion clips are considered by supplying a similarity threshold of 90%. The LMA-based algorithm prunes a large number of the MG transitions keeping only those that better preserve style-continuity

the four LMA components has been assessed separately for each window, returning a Pearson's linear correlation coefficient for each LMA component (in this case, N corresponds to the number of the derived feature measurement for each LMA component); the overall evaluation for a window is a weighted sum of all its LMA components. The weights are user-defined and provide a mechanism to control the importance of each LMA component when comparing motions. For example, the weights used for each of the four LMA components can be set to 25% to weight them equally. The overall correlations computed in each window are then filtered to reduce noise with a 1D Gaussian function with mean $\mu = 0$ and variance $\sigma^2 = 1$. These correlations provide an estimate of the relevance between the windows of the two performances based on the LMA components. Two windows (or frames which anchored in the center, as in our case) are considered stylistically similar if their overall Pearson's linear correlation coefficient is larger than a user-specified threshold; empirical findings show that a good LMA correlation threshold (T_{LMA}) which gives balance between good connectivity with high correlation is 80%.

Finally, these LMA correlations are used to prune transitions of the Motion Graph that are not stylistically coherent, leading to an LMA-based Motion Graph. After pruning the graph, individual frames may still have multiple transitions to other motions, which may have both different posture and LMA correlations. Selecting one of these transitions depends on the application at hand. For instance, one may choose the highest posture correlation, the highest LMA correlation, or a weighted sum of the aforementioned correlations. An example illustration is given in Fig. 2; Motion Graph links a node from a dance motion, where the performer is acting as being *excited*, to numerous motions annotated with a variety of different feelings. Our LMA-based algorithm allows transitions only to motions that satisfy a stylistic similarity threshold of 80%, which in this example are nodes from motions annotated as *happy* and *excited*.

In this work, we compute both the Motion Graph and style-based analysis for all time windows to allow fine-tuning between pose and style correlation at all frames. More specifically, we want to give users the opportunity to decrease the $\mathcal{P}$ threshold when building the graph and increase the $\mathcal{S}$, in order to make the stylistic coherency more important. Nevertheless, optimizations could be applied to save computational time; for instance, instead of computing the LMA-derived features for all time windows, it is possible to compute the LMA similarity only at the pose-based transitions of the generated Motion Graph.

5 Synthesizing plausible dances

In this section, we present 2 example cases of style-coherent dance animation synthesis using the proposed algorithmic framework.

5.1 Data acquisition and processing

We acquired and used real motion capture data of contemporary dance performances. In the data acquisition phase, we used an 8-camera PhaseSpace Impulse X2 motion capture system. Three actors, who were all professional dancers, were asked to prepare and perform 6 contemporary dance choreographies, all with different music. Each performer was recorded independently, so as to avoid influencing the other actors. A total of 18 contemporary dance performances, of around 2700–3600 frames each, were motion captured at a rate of 30 frames per second, approximately 54,000 frames of motion in total. To enable uniform processing of all acquired motion capture data, we retargeted motions to a single 3D skeleton with standard human proportions and body structure.

The algorithms were implemented in C# within the Unity3D game engine, and all processing was performed on a computer running Microsoft Windows 10 with a quad-core Intel i7-4700MQ CPU clocked at 2.4 GHz and 8 GB main memory. These 18 performances were processed in our motion analysis and synthesis framework to construct an LMA MG, preserving only transitions between frames that had above 80% LMA similarity. Using this graph, we have performed two different motion synthesis tasks to obtain our experimental results, presented in Sect. 6.

Our motion analysis relies on offline preprocessing, but motion synthesis is computationally inexpensive and can be performed in real time. Table 2 reports an indicative breakdown of execution time of each processing step using our implementation, which was not specifically optimized. The preprocessing times are for all 18 performances. The most expensive step is the Pearson correlation performed on the LMA-derived features, which takes up to around 72.8% of all preprocessing time. This is expectable since features of all frames in the dataset must be correlated with all features of all other frames; note that this cost could be drastically reduced if we compute the LMA only at the MG edges. The next most expensive part is the posture correlations which takes up around 25.8% of the total time. The LMA-derived feature extraction and the graph construction are trivial and take only 0.8 and 0.6%.

Table 2 Execution time of the main parts of our implementation

Step	Time (mins)	Time (%)
Posture correlations	92	25.8
LMA feature extraction	3	0.8
LMA correlations	259	72.8
Graph construction	2	0.6
Total time	356	

Is this the same data as in 5.1? If so, then why are the numbers different? Maybe it is better to move this two sentences to the previous section to avoid the confusion

5.2 Emotion state distance metric

For every motion sequence, dancers were instructed to act an emotional state for 90–120 s (2700–3600 frames). The acted emotions that were motion captured were *excited*, *happy*, *relaxed*, *sad*, *angry* and *afraid*, one performance for each emotion separately, by each actor. Emotions were selected based on the Russell's circumplex model [35] (RCM) of affect, shown in Fig. 3, which is a model of affective states. Multidimensional scaling of the states distributes emotions in a two-dimensional circular space, with arousal and valence as its dimensions. The particular emotions used in this work were carefully selected to span all 4 quadrants of the RCM.

The projection of emotions in a 2D space provides a set of emotional state distances, as shown in Table 3. As it can be clearly observed, *excited* and *happy* have a low Euclidean distance, while *angry* and *happy* have a higher distance. This can be interpreted as the cost of shifting from one emotional state to another. Since our data were intentionally acted according to selected emotional states, our dataset can be considered to be labeled. However, note that we use this metric exclusively for evaluating and comparing the ability of the MG and the proposed LMA MG to produce animations that are plausible, but we do not use the emotional states as semantic constraints while traversing the graph. This eliminates the requirement for an annotated dataset of any form.

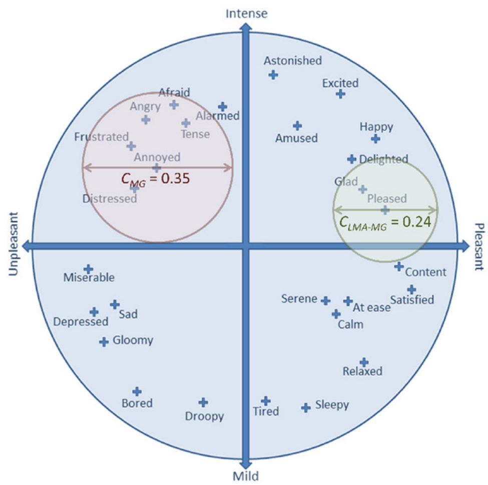

Fig. 3 Multidimensional scaling applied to the Russell circumplex model of affect. Arousal is represented by the vertical axis and valence by the horizontal axis. In our calculations, the diameter of the circle is equal to 1. The red circle shows the average cost of a transition for Motion Graph ($\mathcal{C}_{\text{MG}}$) scaled on the RCM diagram (centered at the emotion *annoyed*), and the green circle the average cost of a transition for our LMA-inspired method ($\mathcal{C}_{\text{LMAMG}}$), centered at the emotion *pleased*

Table 3 Euclidean distances between the emotions in Russell's circumplex model

	Afraid	Angry	Excited	Happy	Relaxed	Bored
Afraid	0	0.08	0.36	0.44	0.68	0.45
Angry	0.08	0	0.43	0.07	0.69	0.39
Excited	0.36	0.43	0	0.13	0.58	0.67
Happy	0.44	0.51	0.13	0	0.48	0.67
Relaxed	0.68	0.69	0.58	0.48	0	0.54
Bored	0.45	0.39	0.67	0.67	0.54	0

When synthesizing motions using our dataset, e.g., when making transitions between the frames of multiple motions over time, we can accumulate the emotional state switching cost using the Euclidean distance between the emotions in Russell's circumplex model. More specifically, the mean cost of transition $\mathcal{C}$ can be calculated as $\mathcal{C} = \frac{1}{k}\sum_k s_i \times s_j$, where k is the number of transitions and $s_i \times s_j$ is the Euclidean distance between the emotions s_i and s_j, as given in Table 3. For example, if we explicitly synthesize a motion that transitions from a *happy* to a *relaxed* motion, then the cost would be 0.48. If a transition occurs between a *happy* motion to another *happy* motion, the cost would be 0. Therefore, we could anticipate that the higher the cost of a motion, the more transitions are likely to have happened to other motions that were not of similar emotion, which potentially influences the motion style.

5.3 Synthesis based on highest similarity

With the MG and the LMA MG pre-computed, we were able to synthesize animation sequences in real time. In this first synthesis task, we generated two motions, one for each of the algorithms (MG and LMA MG). A random frame of a random motion was selected from the entire dataset of 18 contemporary dance performances. Both graphs were traversed starting from that same frame. Every 80–100 frames, each algorithm was forced to select the most similar transition frame according to the set threshold (MG threshold 90% and LMA MG threshold 80%). We have generated 3 different animations, each of 1000 transitions and for each algorithm. These animations cannot be readily compared one-to-one since each graph has different connectivity and the respective transitions are more likely to generate a completely different sequence of motions. Instead, we accumulated the emotion state distance for all transitions (3 runs $\times$ 1000 transitions) and computed the mean cost for each algorithm. The cost for the MG was $\mathcal{C}_{\text{MG}} = 0.35$ with a standard deviation of 0.0046, while the cost for LMA MG was $\mathcal{C}_{\text{LMAMG}} = 0.24$ and the standard deviation for the values of the three runs was 0.0035. Figure 3 shows the average costs, scaled on the RCM diagram, for each method. In order to assess the usefulness and the contribution of the two newly introduced features (f^{28} and f^{29}), we ran the same example, but this time using only the 27 original features. The mean cost was $\mathcal{C}_{\text{LMAMG}*} = 0.26$ with a standard deviation of 0.0036, requiring about 1% less computational time. A lower cost means that overall the animation states followed were closer, with respect to the RCM distances, which also suggests that the style was more coherent for the LMA MG than for the MG. In both cases, repeating the synthesis task multiple times yields very similar results for each algorithm, as it can be observed by the respective standard deviations.

5.4 Synthesis by example

A second synthesis task was performed based on a synthesis-by-example scenario. We asked the professional dancers to perform an additional contemporary dance and intentionally alternate between multiple emotional states (in this order): *relaxed*, *sad*, *angry*, *happy*. Each emotion was acted for 30 s before the dancer moved on to the next one, without stopping or pausing. The task is to synthesize a new dance using motion data from the database that has an emotion that is as similar as possible to the one expressed by the dancer. In real time, the synthesis task (for both MG and our LMA-based MG) began at the first frame of the example motion. As in the previous case, we again forced the motion synthesis algorithms to transition to another motion every 80–100 frames and keep consistency with the emotion of the dancer. Approximately 39 transitions were performed. We then calculated

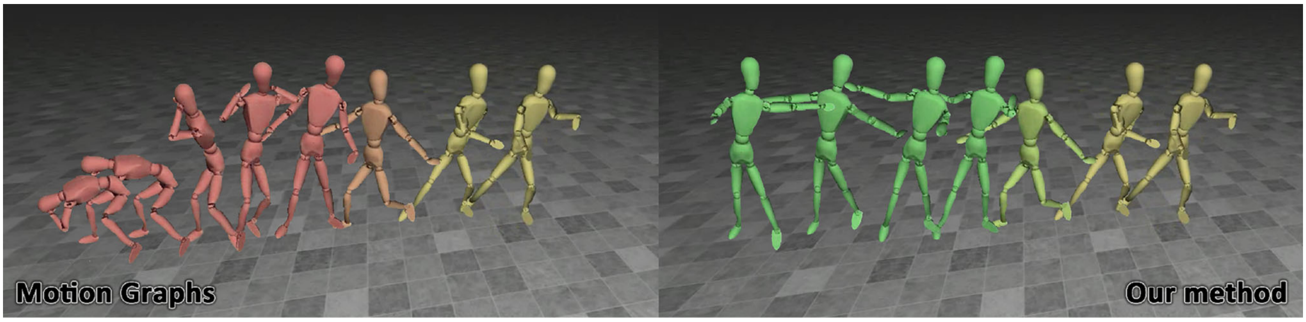

Fig. 4 A sequence of selected key frames from a synthesized dance using MG is shown on the left, while on the right the synthesized motion using the proposed LMA MG is shown. In this example, the character uses the highest possible similarity to select a motion to transition to. MG uses body posture similarity between frames, while our LMA MG method uses a style-coherent similarity metric. On the left, it can be observed that although the character transitions to a similar posture, he quickly crouches to the floor. In contrast, our method selects a motion to transition to in which the character remains standing and has similar motion style overtime

the average transition cost; the transition cost is computed by measuring the distance between the current emotional state against the expected and accumulate that cost over all transitions. For the MG, the average divergence of that was 0.4024, while for LMA MG was 0.3421. Again, when the original 27 features were used, the average divergence for LMA MG* was 0.3592. The readers are encouraged to check the results in the accompanying video.

6 Results and discussion

We have synthesized a variety of dance animation sequences using the framework and the procedures described in the previous sections. The constructed LMA-based Motion Graph by default satisfies posture correlation; in our implementation, we select the transition with the highest LMA correlation. Although the MG algorithm may encourage transition to frames of other motions where body posture is highly similar, in contrast LMA MG selects those transitions that motion style is more coherent, despite body posture being less similar. Figure 4 outlines this fundamental difference between Motion Graphs and LMA Motion Graphs transitioning. We have also applied our method to other dances, including Eastern Mediterranean folk dancing, and we have managed to produce plausible animations; Fig. 5 demonstrates an example utilizing LMA Motion Graphs with Cypriot folk dance motion clips.

By pruning the Motion Graph transitions based on the LMA similarity metric, keeping only transitions that are stylistically more coherent, the number of transitions was drastically reduced. For instance, in one of the motions obtained from the highest similarity motion synthesis, in 1000 frames the Motion Graphs technique with a 90% similarity threshold identified 3544 transitions. The same motion processed in our LMA framework had 579 transitions, a reduction of 84%. For some tasks, low connectivity may be considered a disadvantage, but it must be noted that the number of transitions remaining is very similar in style, and therefore, synthesized motions will still be plausible. In Fig. 6 presents a transition subgraph from one motion to all others. We count present the percentage of transition based on the emotional labeling of the motions. The transitions generated when connecting motions using the MG technique seem to be equally distributed across all other motions. When applying the LMA MG algorithm, the transitions generated are biased toward motions that in our dataset were annotated with emotions of the same RCM quadrant, e.g., *angry* and *afraid*. Note that emotions here are used only for verification. Both MG and the LMA MG are agnostic of the emotion state labeling.

Fig. 5 A sequence of selected key frames from a synthesized dance based on Eastern Mediterranean folk dancing motions, using our style-cohered method. From left to right, the first dance (grayscale) transitions to the second dance (color) while both the bodily and stylistic characteristics are satisfied, achieving good continuity in motion and a plausible movement

Our method can synthesize motion by calculating both posture and style correlations. In contrast to others, we utilize a variety of stylistic features to prune the Motion Graph and choose the transition with the highest similarity that satisfies both posture and stylistic constraints. For instance, Müller et al. [28] describe a number of relational features that study extensively the geometry of the pose; several measurements have been applied that compute the rhythm and

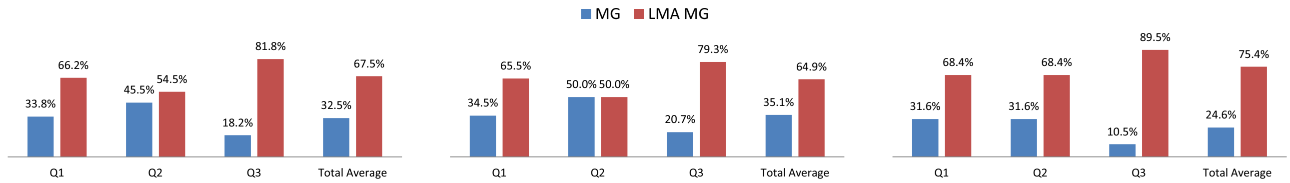

(a) All Participants **(b)** Non-expert Participants **(c)** Expert Participants

Fig. 7 Histogram of responses on the preference of 77 participants between motions generated with MG and LMA MG (our method). (a) Reports the responses of all participants, (b) the responses of non-expert participants are presented, while (c) describes the responses of expert participants are shown. In all three cases, participants preferred the results of our method, irrespective of being experts in dancing, or not

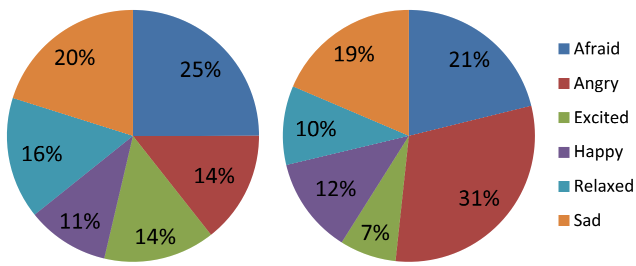

(a) MG **(b)** LMA MG

Fig. 6 In this figure, the frequency of transitions from a reference motion (acted motion: *angry*) to all other motions in the dataset is shown. Transitions are grouped by the acted emotion of the destination motion. Our LMA MG contains a higher percentage of edges from the reference "angry" motion to destination motions with similar emotions, e.g., connected to "angry" destination motions (31%), when compared to MG (14%)

the relations between body and pose points. These measurements add extra constraints so as to restrict the selection of the transition within the graph, increasing its quality and accuracy. Nevertheless, they emphasized more on the body relations without considering the effort required to perform the movement, neither the relation between the performer and the environment. Our method comprises these tighter posture criteria within the BODY and SHAPE components and additionally considers stylistic variations within the EFFORT and SPACE components. Conversely, Kapadia et al. [17] proposed a relatively low number of LMA parameters to evaluate short motion segments. Our LMA style-based method is suitable for longer, more complex, heterogenous and highly stylistic movements, such as dances, in which distinct movements are difficult to isolate, compare and interconnect.

6.1 User evaluation

The results obtained by the two motion synthesis tasks described in Sect. 5 are particularly challenging to evaluate. The emotional state distance provides only a quantitative measure of the synthesized motion and acts like a general purpose error metric. However, the quality of human motion and its plausibility cannot be captured by these metrics. Therefore, we have conducted an online survey to identify whether humans preferred the animations produced using Motion Graphs or the LMA-based Motion Graphs. We asked participants to watch 2 pairs of synthesized dance animations, shown side by side and state which one from each pair has a more coherent dancing performance. One animation was generated with the standard MG algorithm and the other with our LMA MG method for each pair. Both pairs of motions were generated according to the highest similarity synthesis described in Sect. 5.3.

We accumulated the participants' responses as Q1 for the first pair and Q2 for the second. In a third question (Q3), we used a motion captured animation, which was not part of our original dataset, in which actors performed a series or predefined emotional states for 30 s each, i.e., *relaxed*, *sad*, *angry*, *happy*. We then used this motion in an example-based motion synthesis scenario, in which MG was traversed using a high posture similarity threshold. Similarly, the LMA MG was traversed using the LMA style threshold instead. The two motions generated were played along the example motion, and participants were asked to identify which of the two dancers better accompanies the example dancer. Participants were not aware of the method used to generate any of the motions, or even that they were synthesized by algorithms. We also asked other questions, e.g., how experienced they consider themselves in dancing (in a Likert scale of 1–10), their age and gender.

We collected the responses of 77 participants (36 males and 41 females; 11 dance experts and 66 non-experts). Overall, 67.5% of all participants preferred the motions generated by our proposed LMA MG technique versus 32.5% who preferred those of the MG. We have divided participants in two groups according to their reported level of dancing expertise, to identify how experienced and non-experienced participants judged the generated motions. 64.9% of non-experienced participants (who declared a less than 7/10 experience) preferred the LMA MG generated motions versus 35.1% preferring those generated with MG. Similarly, 75.4% of experienced participants (who declared equal or higher than 7/10 experience) reported preferring animations generated with our method versus 24.6% who preferred those of MG. Their responses as percentages for each of the Q1, Q2, Q3 are presented in the histograms of Fig. 7.

6.2 Limitations

It is true that the emotional and stylistic attributes of human behavior are subjective and may depend on the performer's skill, experience, momentary feelings, as well as external factors, such as the environment where the performance happens. During data acquisition, we came across of a number of challenges that may affect the quality of the captured expression. For instance, the mocap suit has markers attached on every limb giving the feeling of restriction or reduced motion to the performer, while the size of the mocap laboratory restricts the movements of the performer to a limited space. In an attempt to reduce the influence of these factors, we allowed the performers 5–10 min for warming up to be familiarize with the outfit and the environment. Moreover, the actors danced with music of their choice and had the required time to prepare the scenario of their performance.

In addition, we would like to acknowledge the difficulty to fully abduct the LMA components using a number of discrete low-level parameters. Human motion is complex, and it consists of stylistic elements that define various states, including the emotion, creativity and intention of performer, while the way they are expressed is subjective based on the perception of the performer. The proposed framework gives a good approximation of the mapping between measurable features and the LMA components.

7 Conclusions and future work

In this paper, we have presented a novel style analysis algorithm based on principles drawn from the movement analysis theories of Laban. We have also shown how it can be integrated in a Motion Graph for synthesizing style-coherent animation given unlabeled motion data. The proposed MG takes into consideration not only body posture, but also style, including the required effort, shape and interaction of the performer with the environment. The results demonstrate the effectiveness of our method to extract qualitative and quantitative characteristics of the movement and find style correlations between different dance performances.

The algorithm as presented above traverses the graph using only a local search. It makes the transition decision without having a long-term goal; a future direction could be to adapt the proposed framework to make global search of the graph to prevent cases where each transition might have a small deviation, but after a longer sequence we might end up in motion clips very dissimilar from the initial one. In addition, the proposed style-coherent motion analysis algorithm could be used independently to find stylistic correlation between motion clips, or for motion clustering in large databases. In the future, we will apply the proposed framework in motion libraries with more dances, larger variety of styles and different performers; we would also like to apply and adapt the LMA-based framework to generic motion data, such as walking, running and jumping.

Acknowledgements This work is co-financed by the European Regional Development Fund and the Republic of Cyprus through the Research Promotion Foundation Under Contract DIDAKTOR/0311/73

References

1. Alexiadis, D.S., Daras, P.: Quaternionic signal processing techniques for automatic evaluation of dance performances from mocap data. IEEE Trans. Multimedia **99**, 1–16 (2014)
2. Arikan, O., Forsyth, D.A.: Interactive motion generation from examples. ACM Trans. Gr. **21**(3), 483–490 (2002)
3. Aristidou, A., Charalambous, P., Chrysanthou, Y.: Emotion analysis and classification: understanding the performers emotions using the LMA entities. Comput. Gr. Forum **34**(6), 262–276 (2015)
4. Aristidou, A., Stavrakis, E., Charalambous, P., Chrysanthou, Y., Himona, S.L.: Folk dance evaluation using Laban Movement Analysis. ACM J. Comput. Cult. Herit. **8**(4), 20:1–20:19 (2015)
5. Aristidou, A., Zeng, Q., Stavrakis, E., Yin, K., Cohen-or, D., Chrysanthou, Y., Chen, B.: Emotion control of unstructured dance movements. In: Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA '17. Eurographics Association, Aire-la-Ville, Switzerland (2017)
6. Brand, M., Hertzmann, A.: Style machines. In: Proceedings of SIGGRAPH '00, pp. 183–192. ACM Press/Addison-Wesley Publishing Co., New York (2000)
7. Calvert, T., Wilke, W., Ryman, R., Fox, I.: Applications of computers to dance. IEEE Comput. Gr. Appl. **25**(2), 6–12 (2005)
8. Chan, J., Leung, H., Tang, J., Komura, T.: A virtual reality dance training system using motion capture technology. IEEE Trans. Learn. Technol. **4**(2), 187–195 (2011)
9. Chao, M.W., Lin, C.H., Assa, J., Lee, T.Y.: Human motion retrieval from hand-drawn sketch. IEEE Trans. Visual. Comput. Gr. **18**(5), 729–740 (2012)
10. Chi, D., Costa, M., Zhao, L., Badler, N.: The emote model for effort and shape. In: Proceedings of SIGGRAPH '00, pp. 173–182. ACM, New York (2000)
11. Durupinar, F., Kapadia, M., Deutsch, S., Neff, M., Badler, N.I.: PERFORM: perceptual approach for adding OCEAN personality to human motion using Laban movement analysis. ACM Trans. Gr. **36**(1), 6:1–6:16 (2016)
12. Hartmann, B., Mancini, M., Pelachaud, C.: Formational parameters and adaptive prototype instantiation for mpeg-4 compliant gesture synthesis. In: Proceedings of the Computer Animation, CA '02, pp. 111–120. IEEE Computer Society, Washington, DC, USA (2002)
13. Hartmann, B., Mancini, M., Pelachaud, C.: Implementing expressive gesture synthesis for embodied conversational agents. In: Proceedings of GW '05, pp. 188–199. Springer, Berlin (2006)
14. Hartmann, S., Trunz, E., Krüger, B., Klein, R., Hullin, M.B.: Efficient multi-constrained optimization for example-based synthesis. Vis. Comput./Proc. Comput. Gr. Int. (CGI 2015) **31**(6–8), 893–904 (2015)
15. Holden, D., Saito, J., Komura, T.: A deep learning framework for character motion synthesis and editing. ACM Trans. Gr. **35**(4), 138:1–138:11 (2016)

16. Hsu, E., Pulli, K., Popović, J.: Style translation for human motion. ACM Trans. Gr. **24**(3), 1082–1089 (2005)
17. Kapadia, M., Chiang, I.k., Thomas, T., Badler, N.I., Kider Jr., J.T.: Efficient motion retrieval in large motion databases. In: Proceedings of I3D '13, pp. 19–28. ACM, New York (2013)
18. Kovar, L., Gleicher, M.: Automated extraction and parameterization of motions in large data sets. ACM Trans. Gr. **23**(3), 559–568 (2004)
19. Kovar, L., Gleicher, M., Pighin, F.: Motion graphs. ACM Trans. Gr. **21**(3), 473–482 (2002)
20. Krüger, B., Tautges, J., Weber, A., Zinke, A.: Fast local and global similarity searches in large motion capture databases. In: Proceedings of SCA '10, pp. 1–10 (2010)
21. Laban, R., Ullmann, L.: The Mastery of Movement, 4th edn. Dance Books Ltd, Binsted (2011)
22. Lee, J., Chai, J., Reitsma, P.S.A., Hodgins, J.K., Pollard, N.S.: Interactive control of avatars animated with human motion data. ACM Trans. Gr. **21**(3), 491–500 (2002)
23. Li, Y., Wang, T., Shum, H.Y.: Motion texture: a two-level statistical model for character motion synthesis. ACM Trans. Gr. **21**(3), 465–472 (2002)
24. Luo, P., Neff, M.: A perceptual study of the relationship between posture and gesture for virtual characters. In: Motion in Games, pp. 254–265 (2012)
25. Min, J., Chai, J.: Motion graphs++: a compact generative model for semantic motion analysis and synthesis. ACM Trans. Gr. **31**(6), 153:1–153:12 (2012)
26. Min, J., Liu, H., Chai, J.: Synthesis and editing of personalized stylistic human motion. In: Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, I3D '10, pp. 39–46. ACM, NY (2010)
27. Müller, M., Röder, T.: Motion templates for automatic classification and retrieval of motion capture data. In: Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA '06, 137–146 (2006)
28. Müller, M., Röder, T., Clausen, M.: Efficient content-based retrieval of motion capture data. ACM Trans. Gr. **24**(3), 677–685 (2005)
29. Nakata, T., Mori, T., Sato, T.: Analysis of impression of robot bodily expression. J. Robot. Mechatron. **14**(1), 27–36 (2002)
30. Okajima, S., Wakayama, Y., Okada, Y.: Human motion retrieval system based on LMA features using IEC method. In: Innovations in Intelligent Machines, pp. 117–130 (2012)
31. Oliveira, J.L., Naveda, L.A., Gouyon, F., Leman, M., Reis, L.P.: Synthesis of variable dancing styles based on a compact spatiotemporal representation of dance. In: Proceedings of IROS '10 (2010)
32. Pearson, K.: Notes on the history of correlation. Biometrika **13**(1), 25–45 (1920)
33. Pejsa, T., Pandzic, I.S.: State of the art in example-based motion synthesis for virtual characters in interactive applications. Comput. Gr. Forum **29**(1), 202–226 (2010)
34. Ren, L., Patrick, A., Efros, A.A., Hodgins, J.K., Rehg, J.M.: A data-driven approach to quantifying natural human motion. ACM Trans. Gr. **24**(3), 1090–1097 (2005)
35. Russell, J.A.: A circumplex model of affect. J. Personal. Soc. Psychol. **39**, 1161–1178 (1980)
36. Senecal, S., Cuel, L., Aristidou, A., Magnenat-Thalman, N.: Continuous body emotion recognition system during theater performances. Comput. Anim. Virtual Worlds **27**(3–4), 311–320 (2016)
37. Shapiro, A., Cao, Y., Faloutsos, P.: Style components. In: Proceedings of Graphics Interface, GI '06, pp. 33–39. Canadian Information Processing Society, Toronto, Canada (2006)
38. Shiratori, T., Nakazawa, A., Ikeuchi, K.: Dancing-to-music character animation. Comput. Gr. Forum **25**(3), 449–458 (2006)
39. Tang, J.K.T., Chan, J.C.P., Leung, H.: Interactive dancing game with real-time recognition of continuous dance moves from 3d human motion capture. In: Proceedings of ICUIMC '11, pp. 50:1–50:9. ACM, New York (2011)
40. Torresani, L., Hackney, P., Bregler, C.: Learning motion style synthesis from perceptual observations. In: Proceedings of NIPS'06, pp. 1393–1400 (2006)
41. Urtasun, R., Glardon, P., Boulic, R., Thalmann, D., Fua, P.: Style-based motion synthesis. Comput. Gr. Forum **23**(4), 799–812 (2004)
42. Vasilescu, M.A.O.: Human motion signatures: analysis, synthesis, recognition. In: Proceedings of the International Conference on Pattern Recognition, ICPR '02, pp. 456–460. IEEE, Washington, DC (2002)
43. Wakayama, Y., Okajima, S., Takano, S., Okada, Y.: IEC-based motion retrieval system using Laban movement analysis. In: Proceedings of KES'10, pp. 251–260 (2010)
44. Wilke, L., Calvert, T., Ryman, R., Fox, I.: From dance notation to human animation: the LabanDancer project: motion capture and retrieval. Comput. Anim. Virtual Worlds **16**(3–4), 201–211 (2005)
45. Xia, S., Wang, C., Chai, J., Hodgins, J.: Realtime style transfer for unlabeled heterogeneous human motion. ACM Trans. Gr. **34**(4), 119:1–119:10 (2015)
46. Yang, Y., Leung, H., Yue, L., Deng, L.: Generating a two-phase lesson for guiding beginners to learn basic dance movements. Comput. Educ. **61**, 1–20 (2013)
47. Yumer, M.E., Mitra, N.J.: Spectral style transfer for human motion between independent actions. ACM Trans. Graph. **35**(4), 137 (2016)
48. Zhao, L., Badler, N.I.: Acquiring and validating motion qualities from live limb gestures. Gr. Models **67**(1), 1–16 (2005)

Andreas Aristidou is a senior post-doc researcher at the Graphics Virtual Reality Lab, University of Cyprus. He had been a Cambridge European Trust fellow at the University of Cambridge, where he obtained his PhD (2011). Andreas has a BSc in Informatics and Telecommunications from the National and Kapodistrian University of Athens (2005), and he is an honor graduate of Kings College London (2006), where he obtained his MSc degree in Mobile and Personal Communications. His main interests are focused on 3D motion analysis and classification, motion synthesis, human animation, and involve motion capture, inverse kinematics, and applications of Conformal Geometric Algebra in graphics.

Efstathios Stavrakis is currently a senior researcher and algorithms engineer at Algolysis Ltd. He was previously a postdoctoral researcher at the University of Cyprus (CY), the Glasgow School of Art (UK), INRIA (FR) and the Vienna University of Technology (AT). He received his Ph.D. (2009) from the Vienna University of Technology (AT) and his M.Sc. (2001) and B.Sc. (2000) from the University of Teesside (UK). His main research interests are in the areas of 3D computer graphics, games and animation, as well as image processing, psychophysics and visual perception in graphical applications.

Margarita Papaefthimiou is a PhD candidate at the University of Crete in the area of Computer Graphics, and she is also working as a Marie Curie Early Stage Researcher(ESR) for the Initial Training Network on Digital Cultural Heritage (ITN-DCH) fellowship project at FORTH-ICS. She has a BA in Computer Science with a specialization in Computer Systems and Networks from the University of Cyprus and an MSc degree in the field of Computer Games and Interactive Technologies from the University of Cyprus in collaboration with Cyprus University of Technology. Her primary interests lie in virtual/augmented reality, real-time rendering and computer animation.

George Papagiannakis is an Associate Professor at the Computer Science department of the University of Crete and Research Fellow at the Computer Vision and Robotics Laboratory in the Institute of Computer Science of the Foundation for Research and Technology Hellas, Heraklion, Greece. His research and development interests are confined in the areas of VR/AR character simulation. In 2011, he has been awarded with a Marie-Curie Intra-European Fellowship for Career Development from the European Commission's Research Executive Agency. He is Principal Investigator in the Marie-Curie ITN-DCH and ViMM EU Projects. In 2016, he was Conference Chair of Computer Graphics International.

Yiorgos Chrysanthou is an Associate Professor at the Computer Science Department of the University of Cyprus, where he is heading the Graphics and Hypermedia lab. He was educated in the UK (BSc and PhD from Queen Mary and Westfield College) and worked for several years as a research fellow and a lecturer at University College London. He has been a Visiting Researcher at the University of California at Berkeley, USA (1992) and at Tel-Aviv University, Israel (1997). His research interests are in the general area of Computer Graphics, Virtual/Augmented Reality and applications.